

University of Southampton Research Repository
ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Aeronautics, Astronautics and Computational Engineering

Fractionated Satellites: A Systems Engineering Analysis

by

Benjamin Samuel Schwarz

Thesis for the degree of Doctor of Philosophy

JULY 2014

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Aeronautics, Astronautics and Computational Engineering

Thesis for the degree of Doctor of Philosophy

FRACTIONATED SATELLITES: A SYSTEMS ENGINEERING ANALYSIS

Benjamin Samuel Schwarz

The current method of operating space-based assets involves the design and launch of large, monolithic spacecraft. These spacecraft are not responsive to failures or changes in mission requirements, as both require the launch of a completely new spacecraft. The concept of fractionated spacecraft was introduced in 2006 by the US Defence Advanced Research Projects Agency (DARPA) as a way of designing and operating space systems which is more responsive than current methods. The fractionated concept involves the decomposition of the traditional monolithic satellite into a number of free-flying spacecraft connected wirelessly. The free-flying satellites each carry a subset of the required subsystems and share their fractionated resources to achieve the mission objectives. Since 2006, studies of this concept have focused on analyses of the value and utility provided by fractionated spacecraft or on design studies of specific systems.

This thesis addresses two key questions: Firstly, is fractionation a mission enabling technology? Secondly, can design rules and guidelines be developed for fractionated satellites to allow continuity of measurements to be maintained and launched mass minimised? The first question is addressed by studying a payload and system for measuring coastal salinity from space. This provides an initial opportunity to assess the fractionated concept from a systems engineering point of view. The second question is addressed by undertaking a more general analysis of fractionated architectures, building on the knowledge gained in development of the fractionated coastal salinity measurement system. A computer model was developed to simulate the lifetime of different fractionated architectures when subjected to subsystem failures. A local search algorithm was used to find fractionated architectures which gave the best compromise between mass launched versus the operational time over a 50 year lifetime. The results from the application of this model showed that architectures that are highly fractionated, containing several homogeneous satellites, best achieved this compromise. These findings provide a contrast to the heterogeneously fractionated architectures proposed by DARPA. When a fractionated architecture is first implemented, the technology required to fractionate all the spacecraft subsystems may not be available. Consequently, the key to the implementation of these first fractionated architectures will be to ensure that there is redundancy in the fractionated subsystems spread across the architecture.

Table of contents

Table of contents	iii
List of tables	v
List of figures.....	vi
DECLARATION OF AUTHORSHIP.....	x
Acknowledgments.....	xii
Nomenclature	xiii
Acronyms	xvi
1. Introduction	1
1.1. What is a fractionated satellite?	1
1.1.1. Why use a fractionated approach?.....	2
1.1.2. A fractionated response to failed spacecraft.....	3
1.1.3. Value	4
1.1.4. The Pleiades fractionated system	6
1.2. Conclusions and research problem.....	6
1.2.1. The research problem	7
1.2.2. Structure of the thesis.....	9
2. Fractionation: A mission enabling technology?.....	10
2.1. Sea surface salinity and its measurement	10
2.1.1. Why measure coastal salinity?.....	10
2.1.2. Ocean salinity measurement by remote sensing.....	11
2.1.3. Aperture synthesis techniques for remote sensing	14
2.1.4. Other ocean salinity measurement systems.....	20
2.1.5. Conclusions	21
2.2. Design drivers and requirements.....	21
2.3. Method	22
2.3.1. Extending the Doppler Radiometer concept	23
2.4. Results.....	24
2.5. A fractionated coastal salinity measurement system.....	35
2.5.1. Requirements and system architecture.....	36
2.5.2. Summary and conclusions	47
3. Developing design principles for fractionated space systems.....	51
3.1. Introduction	51

3.2.	Model structure	53
3.3.	The fractionation of spacecraft subsystems	53
3.3.1.	Communications	54
3.3.2.	On-board data handling	55
3.3.3.	Power	56
3.3.4.	Attitude determination and control system	56
3.3.5.	Orbital control system	58
3.3.6.	Mass and power values used in the model.....	58
3.4.	Architecture evaluation	60
3.4.1.	Fractionated architectures.....	60
3.4.2.	Architecture simulation	62
3.4.3.	Outputs from the architecture evaluation.....	68
3.4.4.	Analysis on the number of Monte-Carlo runs required.....	70
3.5.	Search algorithm	73
3.6.	Testing, debugging and validation	76
3.7.	How the model was used and the data analysed	77
3.7.1.	Sensitivity analysis	77
3.7.2.	How the model was used.....	86
3.7.3.	Analysis method.....	89
4.	Results and data analysis	90
4.1.	Data exploration	90
4.1.1.	Scenario 1.....	90
4.1.2.	Scenario 2.....	103
4.1.3.	Summary	113
4.2.	Results.....	114
4.2.1.	Results and data analysis	120
4.2.2.	Comparison between the three replacement periods	135
4.3.	Discussion.....	139
5.	Conclusions	144
6.	References	147
	Appendix A: Journal paper and conference abstracts	155

List of tables

Table 1 Spatial resolution achievable and spatial frequency coverage for different combinations of antennas mounted on a central hub and with free flying elements	28
Table 2 Mass and power properties for the ACS components	39
Table 3 Mass and power properties for the communications equipment	41
Table 4 Mass and power values for the AODS equipment	46
Table 5 Mass budgets for the satellites in the coastal salinity measurement system	49
Table 6 Power budgets for the satellites in the coastal salinity measurement system	50
Table 7 Mass and power properties for the communications equipment	55
Table 8 Mass and power properties for the ACS components	57
Table 9 Mass and power values for the AODS equipment	58
Table 10 Mass and power values used in the model, derived from SSTL data sheets	59
Table 11 Example matrix format of a fractionated architecture with the payload, OBDH, ADCS and power subsystems fractionated.....	61
Table 12 Heterogeneous fractionated architecture	62
Table 13 2-mixed Weibull function parameters	65
Table 14 Number of Monte-Carlo simulations, run-time and mean standard deviation.....	70
Table 15 Example matrix format of a fractionated architecture with the Payload, OBDH, ADCS and Power subsystems fractionated.....	73
Table 16 Total number of generations evaluated, run time, number of generations evaluated at each mutation level, percentage of generations with viable architectures and maximum fitness found for the number of generations sensitivity analysis	80
Table 17 Time-step and total probability of failure for the probability density functions generated using the time-steps given	81
Table 18 Time-step, model run time and average standard deviation for the time-step sensitivity analysis.....	83
Table 19 Matrix format for a single monolithic satellite	86
Table 20 Matrix format for a six satellite monolithic architecture	87
Table 21 Matrix format for a six satellite architecture which is fully fractionated with a single subsystem on each satellite.....	87
Table 22 Matrix format for a single satellite which is fully fractionated	87
Table 23 Matrix format for a fully fractionated architecture of six satellites with a full complement of subsystems on each	88
Table 24 Matrix format for a fully fractionated architecture of three satellites with a full complement of subsystems on each	88
Table 25 Table describing all scenarios explored using the model.....	89
Table 26 Spacecraft subsystems and associated reliability curve categories.....	115
Table 27 Shape and scale parameters for the reliability curves used [96]	115
Table 28 Subsystem failure probabilities	119

List of figures

Figure 1 Comparison of monolithic, pure and hybrid fractionated architectures [11]	4
Figure 2 The evolution of a hybrid architecture [11]	4
Figure 3 Brightness temperature vs. SST for various salinity levels at 1.43GHz [24]	12
Figure 4 Brightness temperature vs. SST for various salinity levels at 2.65GHz [24]	13
Figure 5 Artists impression of SMOS [32]	15
Figure 6 SMOS field of view showing aliased and un-aliased regions and the Earth-sky horizon [34]	
.....	16
Figure 7 Doppler Radiometer observation scenario where x is the along track direction, y is the across track direction and z is directed upwards from the ground, V_a is the speed of the antenna array, H_0 is the altitude, and θ is the angle which the antennas are inclined to the nadir, in the y direction (“look -angle”) [30]	17
Figure 8 An example of a spatial frequency diagram for an Interferometer with 16 antennas arranged in a “Y” formation (five per arm, and one in the centre)	19
Figure 9 An example of a spatial frequency diagram for a Doppler Radiometer with 16 antennas arranged in a “Y” formation (five per arm, and one in the centre)	19
Figure 10 Early calibrated brightness temperature image from SMOS [41]	20
Figure 11 Free flying antenna formation. Each point represents an individual antenna	24
Figure 12 Free-flying antenna formation as a Doppler Radiometer spatial frequency diagram.	
Spatial frequencies sampled as curves	25
Figure 13 21 antenna formation: central hub (blue square) with 5 free-flying antennas per arm spaced 3 m apart (red circles).....	26
Figure 14 Doppler Radiometer, 21 antenna central hub with 5 free-flying antennas per arm spaced 3 m apart.....	27
Figure 15 Geometry of the pendulum formation [48].....	29
Figure 16 Free-flying antenna motion of a formation consisting of a 21 antenna central hub with 5 free-flying antennas per arm spaced 3 m apart	30
Figure 17 Projected circular orbit relative motion for a 20 m diameter formation	31
Figure 18 Spatial frequency diagram for a circular formation Doppler Radiometer, with 36 antennas in a single 15 m radius ring	32
Figure 19 Spatial frequency diagram for a circular formation Doppler Radiometer, with two concentric rings of 15 m and 7.5 m Radius with 36 antennas per ring	33
Figure 20 Spatial frequency diagram for a circular formation Doppler Radiometer, with three concentric rings of 15 m, 7.5 m and 2 m radius with 36 antennas per ring	34
Figure 21 Circular formation, with three concentric rings of 15 m, 7.5 m and 2 m radius with 36 spacecraft equally spaced around the circumference of each ring.....	35
Figure 22 Architecture block diagram for the fractionated coastal salinity system	37
Figure 23 Primary payload data flow diagram. Each box represents a separate satellite. The blue lines represent the flow of the data collected by the individual antennas, while the red lines show the flow of the correlated images.	40
Figure 24 The centralised architecture formation control loop	43
Figure 25 The distributed architecture formation control loop.	44
Figure 26 Star network (left) and mesh network (right) topologies.....	45

Figure 27 Failure rate curve for scenario 1. On the x-axis is the age of the subsystem and on the y-axis is the failure rate in failures per year.....	64
Figure 28 2-Mixed Weibull probability density function	65
Figure 29 Replacement flow diagram	67
Figure 30 Graph showing the fitness values for a typical run of the model	69
Figure 31 Model run with 50 Monte-Carlo simulations per architecture with 1 standard deviation error bars	71
Figure 32 Model run with 100 Monte-Carlo simulations per architecture with 1 standard deviation error bars	71
Figure 33 Model run with 300 Monte-Carlo simulations per architecture with 1 standard deviation error bars	72
Figure 34 Model run with 1000 Monte-Carlo simulations per architecture with 1 standard deviation error bars	72
Figure 35 Block diagram showing search algorithm process	75
Figure 36 Model run with 100 architectures explored and 1 standard deviation error bars.....	78
Figure 37 Model run with 500 architectures explored and 1 standard deviation error bars.....	79
Figure 38 Model run with 1000 architectures explored and 1 standard deviation error bars.....	79
Figure 39 Model run with 3000 architectures explored and 1 standard deviation error bars.....	80
Figure 40 2-Mixed Weibull probability density function generated with a time-step of 0.05 years	81
Figure 41 2-Mixed Weibull probability density function generated with a time-step of 0.5 years..	82
Figure 42 2-Mixed Weibull probability density function generated with a time-step of 1 year	82
Figure 43 2-Mixed Weibull probability density function generated with a time-step of 5 years....	83
Figure 44 Model run with a time step of 0.05 years and 1 standard deviation error bars	84
Figure 45 Model run with a time step of 0.1 years and 1 standard deviation error bars	84
Figure 46 Model run with a time step of 0.5 years and 1 standard deviation error bars	85
Figure 47 Model run with a time step of 1 year and 1 standard deviation error bars	85
Figure 48 Model run with a time step of 5 years and 1 standard deviation error bars	86
Figure 49 Failure rate curve for scenario 1. On the x-axis is the age of the subsystem and on the y-axis is the failure rate in failures per year.....	91
Figure 50 Probability distribution for $P(t_{op})$ for scenario 1	92
Figure 51 Probability distribution for $P(N_{satellites})$ for scenario 1	93
Figure 52 Histogram showing conditional probability distributions for $P(N_{satellites} t_{op} \geq 90\%)$	93
Figure 53 Histogram showing conditional probability distributions for $P(N_{satellites} t_{op} \leq 80\%)$	94
Figure 54 Probability distribution for $P(N_{subsystems})$ for scenario 1	95
Figure 55 Histogram showing conditional probability distributions for $P(N_{subsystems} t_{op} \geq 90\%)$	95
Figure 56 Histogram showing conditional probability distributions for $P(N_{subsystems} t_{op} \leq 80\%)$	96
Figure 57 D versus t_{op} for scenario 1.....	97
Figure 58 Probability distribution for $P(F)$ for scenario 1	97
Figure 59 Histogram showing conditional probability distributions for $P(F top \geq 90\%)$	98
Figure 60 Histogram showing conditional probability distributions for $P(F top \leq 80\%)$	98
Figure 61 Probability distribution for $P(n_f)$ for scenario 1	99
Figure 62 Histogram showing conditional probability distributions for $P(n_f top \geq 90\%)$	99
Figure 63 Histogram showing conditional probability distributions for $P(n_f t_{op} \leq 80\%)$	100
Figure 64 a_f versus t_{op} for scenario 1	100
Figure 65 Probability distribution for $P(M)$ for scenario 1.....	101

Figure 66 Histogram showing conditional probability distributions for $P(M \text{top} \geq 90\%)$	102
Figure 67 Histogram showing conditional probability distributions for $P(M \text{top} \geq 90\%)$	102
Figure 68 2-Mixed Weibull probability density function	103
Figure 69 Probability distribution for $P(\text{Fitness})$ for scenario 2	104
Figure 70 Probability distribution for $P(N_{\text{satellites}})$ for scenario 2	105
Figure 71 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} \text{Fitness} \geq 1.002)$	105
Figure 72 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} \text{Fitness} \leq 1.001)$	106
Figure 73 Probability distribution for $P(N_{\text{subsystems}})$ for scenario 2	106
Figure 74 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} \text{Fitness} \geq 1.002)$	107
Figure 75 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} \text{Fitness} \leq 1.001)$	107
Figure 76 D versus Fitness for scenario 2.....	108
Figure 77 Probability distribution for $P(F)$ for scenario 2	108
Figure 78 Histogram showing conditional probability distributions for $P(F \text{Fitness} \geq 1.002)$	109
Figure 79 Histogram showing conditional probability distributions for $P(F \text{Fitness} \leq 1.001)$	109
Figure 80 F versus fitness for the top architecture found by each model run in scenario 2	110
Figure 81 E versus fitness for the top architecture found by each model run in scenario 2	111
Figure 82 E versus the difference between n_f and n_{fs} for the top architecture from each model run in scenario 2	111
Figure 83 Probability distribution for $P(n_f)$ for scenario 2	112
Figure 84 Histogram showing conditional probability distributions for $P(n_f \text{Fitness} \geq 1.002)$	112
Figure 85 Histogram showing conditional probability distributions for $P(n_f \text{Fitness} \leq 1.001)$	113
Figure 86 a_f versus fitness for scenario 1.....	113
Figure 87 2-Mixed Weibull probability density function, Category 1.....	115
Figure 88 2-Mixed Weibull probability density function, Category 2	116
Figure 89 2-Mixed Weibull probability density function, Category 4	116
Figure 90 2-Mixed Weibull probability density function, Category 6	117
Figure 91 2-Mixed Weibull probability density function, Category 7	117
Figure 92 2-Mixed Weibull probability density function, Category 8	118
Figure 93 2-Mixed Weibull probability density function, Category 10	118
Figure 94 2-Mixed Weibull probability density function, Category 11.....	119
Figure 95 Probability distribution for $P(\text{Fitness})$ for the 0.7 year replacement time scenario	120
Figure 96 Probability distribution for $P(N_{\text{satellites}})$ for the 0.7 year replacement time scenario	121
Figure 97 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} \text{Fitness} \geq 1.0006)$	121
Figure 98 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} \text{Fitness} \leq 1.0002)$	122
Figure 99 Probability distribution for $P(N_{\text{subsystems}})$ for the 0.7 year replacement time scenario....	122
Figure 100 Histogram showing conditional probability distributions for $P(N_{\text{subsystems}} \text{Fitness} \geq 1.0006)$	123
Figure 101 Histogram showing conditional probability distributions for $P(N_{\text{subsystems}} \text{Fitness} \leq 1.0002)$	123

Figure 102 D versus fitness for the 0.7 year replacement time scenario	124
Figure 103 Probability distribution for $P(F)$ for the 0.7 year replacement time scenario	124
Figure 104 Histogram showing conditional probability distributions for $P(F \text{Fitness} \geq 1.0006)$...	125
Figure 105 Histogram showing conditional probability distributions for $P(F \text{Fitness} \leq 1.0002)$...	125
Figure 106 External redundancy versus fitness for the top 15 architectures for the 0.7 year replacement time scenario	126
Figure 107 External redundancy versus the difference between n_f and n_{fs} for the top 15 architectures for the 0.7 year replacement time scenario	126
Figure 108 Number of satellites versus total launched mass and percentage operational time for $D = 6$ and $N_{subsystems} = 6$	127
Figure 109 Number of satellites versus fitness for $D = 6$ and $N_{subsystems} = 6$ with the optimum numbers of satellites circled in blue	128
Figure 110 Histogram for operational times recorded in 0.05 and 0.7 year replacement time scenarios	129
Figure 111 Histogram for total launched masses recorded in 0.05 and 0.7 year replacement time scenarios	130
Figure 112 Number of satellites vs total launched mass and percentage operational time for $D = 6$ and $N_{subsystems} = 6$ using a 0.05 year replacement period	131
Figure 113 Number of satellites vs fitness for $D = 6$ and $N_{subsystems} = 6$ using a 0.05 year replacement period with the optimum numbers of satellites circled in blue	132
Figure 114 Histogram for operational times recorded in 0.05, 0.7 and 1.5 year replacement time scenarios	133
Figure 115 Histogram for total launched masses recorded in 0.05, 0.7 and 1.5 year replacement time scenarios	133
Figure 116 Number of satellites vs total launched mass and percentage operational time for $D = 6$ and $N_{subsystems} = 6$ using a 1.5 year replacement period	134
Figure 117 Number of satellites vs fitness for $D = 6$ and $N_{subsystems} = 6$ using a 1.5 year replacement period with the optimum numbers of satellites circled in blue	135
Figure 118 Number of failures requiring a new satellite comparison	136
Figure 119 Average age at failure, 0.05 year replacement period	137
Figure 120 Degree of fractionation vs fitness, 0.05 year replacement period for all architectures and those that experienced 15 or more failures. Error bars show 1 standard deviation	138
Figure 121 Degree of fractionation vs number of failures requiring a new satellite comparison. Error bars show 1 standard deviation.....	138

DECLARATION OF AUTHORSHIP

I, Benjamin Samuel Schwarz declare that this thesis and the work presented in it are my own and have been generated by me as the result of my own original research.

Fractionated Satellites: A Systems Engineering Analysis

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:

Schwarz, B. S., Tatnall, A. R. L., & Lewis, H. G. (2012). Coastal salinity measurement using a Doppler Radiometer. *Advances in Space Research*, 50(8), 1138–1149.
doi:10.1016/j.asr.2012.01.020

Signed:

Date:.....

In loving memory of Sita, Issy and Harry

For Annette

Acknowledgments

I would like to thank many people for their invaluable advice, assistance and support throughout my PhD research. Firstly, my supervisors Dr. Adrian Tatnall and Dr. Hugh Lewis for steering me through my PhD. Your wisdom, expertise, advice and assistance have been invaluable, and you have always found time to help when it was needed.

Dr. Adriano Camps and Dr. Hyuk Park from the Universitat Politècnica de Catalunya for their advice and help with understanding the Doppler Radiometer concept presented in chapter 2.

My PhD colleagues and friends Simon George and Adam White for making the process of studying for a PhD much more enjoyable! You have always been there to share a joke, discuss the latest happening in the world of Space or to act as a sounding board for research ideas.

My oldest and dearest friend Alon Ratner, for a great portion of our 22 year friendship we have shared a passion for science and the pursuit of knowledge. Even though we are now in very different fields, may our shared enjoyment of science long continue!

Mom, Dad, Daniel and Adam your continued love, support and inspiration has been a pillar of strength for me during my PhD.

My wonderful wife, Sarah you have borne the brunt of the ups and downs of my PhD more than anyone, I could not have made it through this four year period without your love and support.

Finally, I would like to thank my Grandmother, Annette and Grandfather, Harry. You bought me my first telescope and ignited my passion for space which has not diminished in 13 years. I would not be on this path now if it was not for you.

Nomenclature

Chapter 2:

α_0	Initial in-plane phase angle
β_0	Initial out-of-plane phase angle
λ	Wavelength of measurement
$\Delta\theta$	Difference in true anomaly between the two bodies
$\Delta\Omega$	Difference in right ascension of the ascending subsystem between the two bodies
θ	Look angle
η	The efficiency of the solar array system
A_{SA}	Solar array area
a	Semi-major axis
c_i	PCO constants
D	Diameter of the antenna
D_x	Separation, in meters between pairs of antennas in the x direction
D_y	Separation, in meters between pairs of antennas in the y direction
d_{ref}	Diameter of the circular formation
F_{sun}	Solar flux
H_0	Altitude
i	Inclination of the orbits of the bodies
L	Distance from the Earth to the instrument
L_x	Antenna footprint length in x direction
L_y	Antenna footprint length in y direction
n	Number of satellites in the network
P	Power required by the spacecraft
r	Spatial resolution
T_b	Brightness temperature

t	Time
u	Horizontal spatial frequencies
v	Vertical spatial frequencies
V_a	Speed of the antenna array
x	Along track direction
y	Across track direction
\bar{y}	Centre line of the footprint along the y axis
z	Vertical Direction

Chapters 3 and 4:

α	Mixture parameter
β	Shape parameter
η	Solar array efficiency
θ	Scale parameter
A_{SA}	Solar array area
a	Age of the subsystem in years
a_f	Average age of the subsystems at failure
a_{\max}	Maximum subsystem age
b_i	One of the following six model parameters: $N_{satellites}$, $N_{subsystems}$, D , F , n_f and a_f
D	Degree of fractionation
E	External redundancy
F	Fractionated subsystems per satellite
F_{sun}	Solar flux at the Earth
k	Parameter to be normalised
\bar{k}	Mean of k
M	Total launched mass
$M_{arch_{ij}}$	Architecture matrix, i is the column number and j is the row number

$M_{Fractionated}$	The matrix formed by the fractionated columns only of the architecture matrix
N_{F_k}	Normalised fitness value of parameter k
N_{rand}	A random number
$N_{Satellites}$	Number of satellites
$N_{Subsystems}$	Number of subsystems per satellite
n_f	Number of failures
n_{fs}	Number of failures requiring a new satellite to be launched
n_{SL}	The number of satellites launched
P	The power required by the spacecraft when in sunlight
P_D	Probability density function
$P_{failure}(t)$	Probability of failure at time t
$P(b_i S)$	Probability of a particular value of parameter b_i occurring given the observation of a fitness value or operational time
R	Overall failure rate function
R_{EOL}	End of life wear-out rate
R_e	Early wear-out failure rate
R_r	Random failure rate
$r(t)$	Reliability of a satellite at time t
S	Optimisation parameter (Fitness or operational time)
T	Time-step size
t	Time
t_{op}	Operational time
$Var(k)$	Variance of k

Acronyms

ACS	Attitude Control Subsystem
ADCS	Attitude Determination and Control System
AMSR-E	Advanced Microwave Scanning Radiometer – Earth Observing System
AODS	Attitude and Orbit Determination Subsystem
ASCAT	Advanced Scatterometer
CFF	Coulomb Formation Flying
COTS	Commercial, Off The Shelf
DARPA	Defence Advanced Research Projects Agency
EO	Earth Observation
EP	Electric Propulsion
ESA	European Space Agency
FDIR	Fault Detection, Isolation and Recovery
GA	Genetic Algorithm
GPS	Global Positioning System
IBS	Ion Beam Shepherd
LEO	Low Earth Orbit
MAUT	Multi-Attribute Utility Theory
MC	Monte-Carlo
MIRAS	Microwave Imaging Radiometer using Aperture Synthesis
MISAR	Motion Induced Synthetic Aperture Radiometer
MLE	Maximum Likelihood Estimation
NASA	National Aeronautics and Space Administration
OBDH	On Board Data Handling
OCS	Orbit Control Subsystem
PCDU	Power Conditioning and Distribution Unit
PCO	Projected Circular Orbit

PISAR	Position Induced Synthetic Aperture Radiometer
PSU	Practical Salinity Unit
RPS	Relative Position Sensors
SLFMR	Scanning Low Frequency Microwave Radiometer
SMOS	Soil Moisture and Ocean Salinity
SPOT	Système Pour l'Observation de la Terre
SSS	Sea Surface Salinity
SST	Sea Surface Temperature
SSTL	Surrey Satellite Technology Limited
TT&C	Telemetry, Tracking and Command
USAF	United States Air Force
VCDM	Value-Centric Design Methodologies
XEUS	X-ray Evolving Universe Spectroscopy

1. Introduction

The current method of operating assets in orbit around the Earth involves the use of single spacecraft, often, with masses in excess of 1000 kg [1]. On-board these spacecraft are the subsystems and payloads required to achieve the mission objectives. A failure in a subsystem or payload, or the depletion of resources (such as propellant), results in the capabilities of the spacecraft being lost and a new one having to be launched. Any increase in capability or change in mission objectives can only be achieved with the launch of a new satellite. Launch costs can reach \$20,000 per kg (financial year 2000 values) depending on the launch vehicle used [2]. The launch of large replacement spacecraft is therefore very expensive and, since 2000, has led to an increase in the use of “small” satellites [1]. These spacecraft have a mass of less than 500 kg, and typically rely on the use of commercial off-the-shelf components (COTS) [1], [3]. However, the use of small satellites does not deal with the problem of major subsystem failures disabling the spacecraft. There is also evidence that small satellites are less reliable than their larger counterparts [4].

The concept of fractionated spacecraft was introduced by Owen Brown and Paul Eremenko of the US Defence Advanced Research Projects Agency (DARPA) in 2006 as part of their System F6 (Future, Fast, Flexible, Free Flying, Fractionated Spacecraft united by Information eXchange) program [5]. The concept represented a new approach to space mission design, and addressed the issues raised above. The main aim of the program was to demonstrate the responsiveness of the fractionated concept and its technological feasibility [5]. In their 2013 Global Horizons report, the United States Air Force (USAF) listed fractionated satellites as one of three game changing space technologies. Furthermore, it recommended that fractionated constellations be pursued, and that the space acquisition process be redefined in accordance with fractionated satellites (and low cost launch opportunities) to reduce the cost of advanced technologies by a factor of ten [6].

1.1. What is a fractionated satellite?

A fractionated satellite consists of a number of distinct free flying modules that have specific functions (such as a high speed downlink or a payload), and are connected wirelessly. The modules therefore function as a single entity and can deliver the same capability as a single satellite with the same functions on-board [5]. The concept of a cluster of free flying spacecraft functioning as a single entity was identified well before the System F6 program was launched. In

1984, Molette et al. compared a cluster of free-flying spacecraft with a modular spacecraft assembled on orbit for a large Geostationary communications system [7]. Their cluster consisted of the communications payload distributed between several free-flying satellites which were otherwise entirely monolithic. The authors concluded that a modular system that was physically assembled on orbit would be cheaper and more flexible than a cluster of spacecraft. Wireless technologies were more immature at the time of this study, even for terrestrial purposes, and this limitation would have had an influence on the ability to decompose the system across the cluster and on the conclusions drawn.

Brown and Eremenko described two different types of fractionation: *heterogeneous fractionation*, which is the use of functionally dissimilar modules, and *homogeneous fractionation*, which is the use of functionally similar modules. The modules can be connected through data links (including those exchanging attitude and position information), the transmission of power or even the sharing of forces and torques [5]. Orbital Sciences subsequently developed a fractionated system concept, called Pleiades, under a DARPA contract [8]. They argued here that the key innovation for fractionation is *information integration*, rather than physical decomposition. The information integration aspect distinguishes fractionation from clusters of satellites, such as described by Molette et al., or constellations of satellites such as the Disaster Monitoring Constellation [7], [9].

1.1.1. Why use a fractionated approach?

Brown proposed in 2004 that breaking up a satellite into functional blocks, which are launched separately and then assembled back into a complete satellite once in orbit, gives the system an “enhanced capability of on-orbit reconfiguration, upgrade and replacement” [10]. It also enables the launch of much larger space systems which could not fit onto today’s launch vehicles. Risk is reduced by splitting the spacecraft over several launches, which means that if there is a launch failure the entire system is not lost. It also provides the flexibility to build up the system in orbit incrementally, allowing different modules to be developed at different rates.

In Brown and Eremenko’s paper that introduced the concept of fractionated spacecraft, they presented a broad definition of responsive space as “the capability to respond rapidly to uncertainty” and conclude that fractionation enables responsiveness [5]. The sources of uncertainty in the lifetime of a spacecraft (from the initial study to the end of life) were assessed in this paper as:

- *Technical uncertainty*: Rather than having to replace the whole spacecraft in the event of a failure, a new module can simply be launched and added to the wireless network.
- *Environmental uncertainty*: Having multiple satellites in the system spreads the risk from single event effects, debris and other risks to the system, and only the affected modules have to be replaced.
- *Launch uncertainty*: Modules can be launched separately on different launch vehicles so that the entire system is not lost due to a single launch failure.
- *Demand uncertainty*: As the demand for the system changes it is very easy to upgrade or add new modules at any stage of the program.
- *Requirements uncertainty*: Payloads (and other modules) are mechanically decoupled from the system, so if one requirement changes the impact does not ripple across the whole system.
- *Funding uncertainty*: Incremental development strategies mean that cost is deferred to future funding cycles and that the total system cost can be reduced due to the discounting of future year investments.

They further argued that focusing solely on requirements and minimising cost results in the (incorrect) conclusion that the lifetime and capability of satellites must be increased continually. The outcome is an increase in the mass of satellites, posing more challenging integration, testing and launch problems and the use of multi-functional payloads imposing more complex pointing and isolation requirements for the spacecraft. This has produced systems that are less responsive to the uncertainties described above [5].

1.1.2. A fractionated response to failed spacecraft

Two possible strategies were described by Brown et al. for replacing failed spacecraft [11]. The first is a pure fractionated architecture, where the initial system is made up of free-flying satellites connected by a wireless network. If a satellite fails (or is upgraded), a new free-flying spacecraft is added to the network to replace the failed one. The second option is a hybrid of the monolithic and fractionated architectures. In this approach, the initial state is a single spacecraft, but with a wireless network connecting the individual subsystems, rather than a traditional wire harness. If a subsystem or component fails, a new, free-flying spacecraft carrying just the replacement subsystem is flown into the vicinity of the spacecraft and joins the wireless network. Figure 1 gives an overview of the monolithic, pure fractionated and hybrid architectures, and Figure 2 illustrates the evolution of the Hybrid architecture as described by Brown et al. [11].

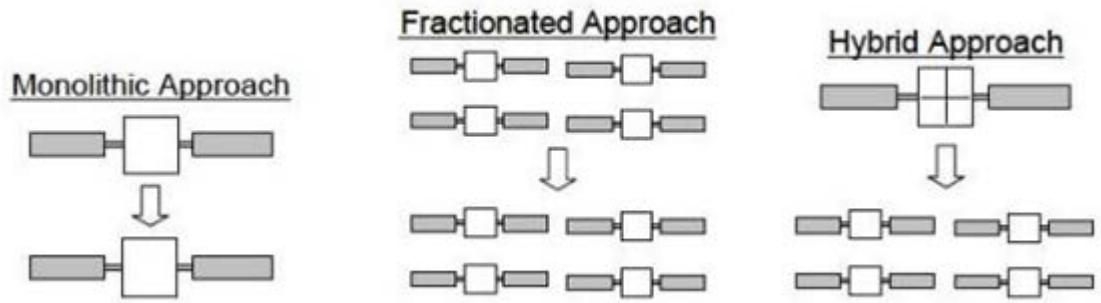


Figure 1 Comparison of monolithic, pure and hybrid fractionated architectures [11]

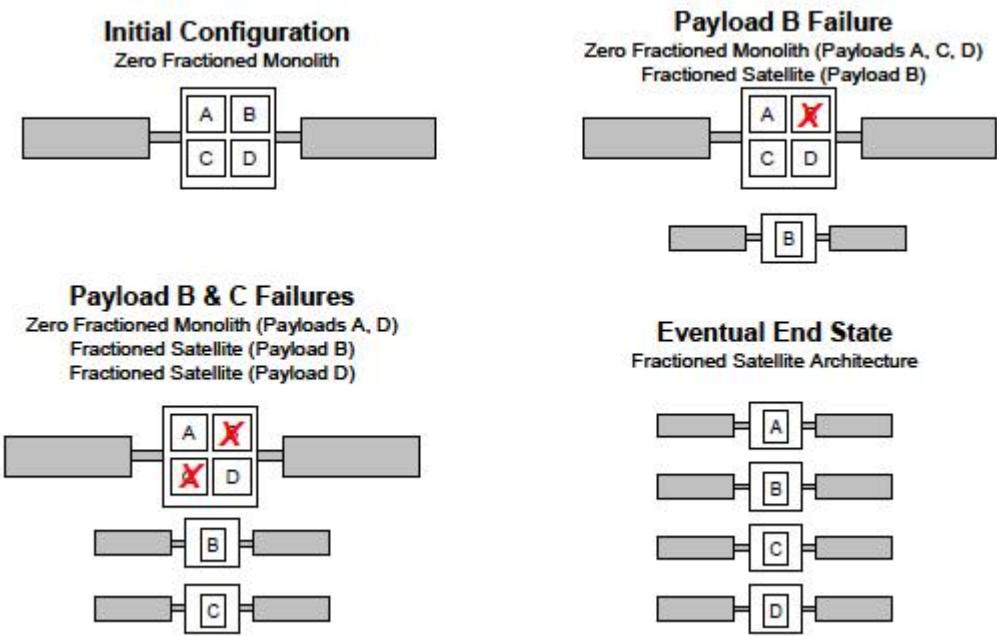


Figure 2 The evolution of a hybrid architecture [11]

1.1.3. Value

The additional components needed to provide redundancy to combat uncertainty also increase the complexity of the spacecraft, which, in turn, increases the system size, cost and schedule [12]. Any change in requirements or funding will also mean that the entire system may have to be redesigned, as traditional, monolithic spacecraft cannot adapt to these changes easily during the design phase or at all during their operational period. To combat these problems, Brown and Eremenko suggested shifting the focus of spacecraft design from requirements to the capability of the spacecraft to respond to uncertainty [12].

To design for uncertainty and to optimise a fractionated architecture, Brown and Eremenko argued that a value-centric mind-set must be adopted [12]. They define value as “a measure—wholly apart from cost—that reflects the utility of a particular system to its owner or operator” [13]. The majority of the work described by Brown, Eremenko and Roberts [14] focused on quantifying the sources of value likely to be affected by the use of fractionation. The value-enhancing benefits of adopting fractionated architectures have been identified as follows [11]:

- Diversification of launch and on-orbit failure risk
- Reliability enhancement through emergent sharing of subsystems and resources
- Scalability in response to service demand fluctuations
- Upgradeability in response to technological obsolescence
- Incremental deployment of capability to orbit
- Graceful degradation of capability on-orbit
- Robustness in response to funding fluctuations and requirements changes
- Reduced integration and testing due to subsystem decoupling
- Production learning across multiple similar modules
- Enabling spacecraft to be launched on smaller launch vehicles with shorter timescales.

The justification put forward by DARPA for changing the design philosophy to a fractionated approach is one based on the economic benefits. However, to help inform the design process, the impact of fractionation on the system must be understood. In this thesis it is argued that, as the fractionated spacecraft concept is revolutionary, new operational concepts and design practices must be developed before a fractionated space system can be implemented.

Studies undertaken for DARPA have focused on non-technical parameters for evaluation, using what has been termed “Value-Centric Design Methodologies” (VCDM) and Multi-Attribute Utility Theory (MAUT) [13], [14]. These methods have been used to make the case for the development of fractionated spacecraft, by arguing that such systems improve the utility of space-based assets and that maximising the “value” should be the objective, not simply minimising monetary cost.

1.1.4. The Pleiades fractionated system

The Pleiades fractionated system developed by Orbital Sciences is the only DARPA-funded study to produce a concept design for a fractionated space system [8]. Here, the functionalities of the spacecraft subsystems are spread over multiple, heterogeneous spacecraft, each carrying a unique mission payload or resource (data processor, data recorder, high bandwidth downlink etc.) [8]. The design of the Pleiades system is based on a Low Earth Orbit Earth science mission, utilising three optical imagers at various stages of development and with varying operational lifetimes. The images are co-registered to enhance the value of the scientific return, and image processing capability and large amounts of data storage are required on board. The final system is composed of seven spacecraft, put into orbit by three launches. The first launch would carry two spacecraft with the most mature sensor on one and a data recorder and downlink on the other. This would allow the first segment of the system to execute part of the mission. To represent a change in requirements, midway through the development, a new scientific objective was introduced for the mission. Consequently, the second launch would carry three spacecraft, one carrying a second imager, which was previously too immature to fly, a second carrying a computer for data processing capability and a third carrying a payload to make the measurements for the new requirements. Finally, the third launch would add a further two spacecraft, carrying the final imager and a new data recorder with downlink providing additional resources to allow the architecture to grow further should there be the need for additional payloads [8].

The capabilities that were considered fractionatable were the mission payload, Telemetry, Tracking and Command (TT&C) communication, high bandwidth downlink (for payload data), data storage and mission data processing. Other functions, such as attitude control, thermal and power management and propulsion, were not considered as fractionatable without significant technological breakthroughs [8].

1.2. Conclusions and research problem

Two main types of studies on fractionated satellites have been performed, those that assess the value of fractionation to the user (primarily led and directed by DARPA) and those that propose specific fractionated systems, as in Molette et al. or the Orbital Sciences studies [7], [8]. The studies described above were conducted using conventional design practices and, may not necessarily produce optimal fractionated systems. A wider, more general and technical study of the fractionated concept is required to re-evaluate the design principles used. It has also yet to be determined which missions are best suited to fractionation, or indeed if any missions are enabled by the use of fractionation.

1.2.1. The research problem

In this research, Brown and Eremenko's definition of a fractionated space system is expanded to consist of a number of free flying satellites each housing a subset of the spacecraft subsystems. These satellites share their fractionated resources to achieve the mission objectives. There is an important distinction to be made between cluster and constellation missions and fractionated missions: both use physical decomposition and separation of the space assets across several free-flying satellites, but fractionated systems take the decomposition one step further with the information integration techniques which enable the free-flying spacecraft to interact and behave as a single entity. The different architectures, described by Brown and Eremenko, Orbital Sciences and other studies are only a small fraction of the possible architectures that could be implemented. There are in fact a wide range of architectures between these solutions that use the principles of physical decomposition and information integration.

Traditional spacecraft system design methods may not produce optimal fractionated designs, as these methods have been developed over 60 years to produce optimal *monolithic* spacecraft. Therefore new design principles and methods must be developed specifically for fractionated spacecraft. In addition, the true value of fractionation may be in enabling missions that are unachievable with monolithic satellites and, for currently achievable missions, in extending their lifetimes to provide continuity of measurements for a much longer period of time [15]. Mass still has to be minimised, as launch costs will still be high without a major breakthrough in re-usable launch vehicles.

This leads to two key questions which are addressed in this thesis:

1. Is spacecraft fractionation a mission enabling technology?
2. What is the best way to design a fractionated system, such that a continuity of measurements (or other payload service) can be maintained over a period of time that is greater than a "traditional" mission lifetime (10 to 12 years) whilst minimising the mass launched over the operational lifetime?

Question 2 leads to the following further questions which will be addressed in this thesis:

- At what level of overall technological maturity should a fractionated architecture be implemented?
- Which subsystems should be fractionated?
- Is redundancy required across the fractionated system, not just within individual spacecraft within the system?

- What level of redundancy is required?

Here, mass and continuity of measurements are used as evaluation parameters as they are demonstrable *technical* parameters associated with a fractionated system. This allows the fractionated concept to be analysed from a technical point of view. The objective is to provide guidance to engineers designing fractionated systems, informing them about the key technical features of this concept.

For the purposes of this research, the questions posed above were explored with respect to Earth Observation (EO) mission objectives that can be achieved by the use of spacecraft in Low Earth Orbit (LEO). Other mission orbits require at least one orbit insertion manoeuvre to be performed, which imposes stringent propulsion requirements on the design. In addition, it is expected that many, or all satellites that will be employed in fractionated architectures will be “small” satellites. The “small” satellite philosophy relies on the use of COTS components. This allows the rapid development of spacecraft based on the experience of using these components in orbit, without them being designed for the space environment [3]. The majority of missions which use COTS components have been LEO EO missions, giving a substantial heritage to the components and subsystems in this environment, making them the most applicable to this study. Other mission type such as communications systems use more traditional, space hardened components and utilise lengthy design, analysis and testing phases making them less suitable for use in this research.

Question 1, the question of fractionation as a mission enabling technology was addressed by investigating a mission with the objective of measuring coastal salinity. Such a mission can be achieved using an instrument which uses fractionated principles. This analysis provided an initial exploration of the design of a fractionated system for a specific application. This allowed some of the key issues surrounding fractionation to be understood and fed into the work that followed addressing question 2. This was a more generic study and the question was addressed by examining an Earth Observation mission which utilised a single instrument. This mission was chosen due to the availability of data on the instruments used. Examples of this mission type are the Système Pour l'Observation de la Terre (SPOT) series of satellites [16] and the upcoming Sentinel-1 Satellite [17]. The payload had only one primary objective: to produce narrow-swath, high-resolution images of the Earth. This limited the requirements on the system and helped simplify the analysis required.

1.2.2. Structure of the thesis

Chapter 2 assesses fractionation as a potential mission enabling technology. In this chapter it is proposed that a system designed to measure sea surface salinity in coastal regions is one which is enabled and will benefit from the use of fractionation. Such a system requires multiple payload antennas hosted on separate, free-flying spacecraft to achieve the mission. The signals from these antennas are then combined by a central processing unit to produce the final image. This measurement system uses the two key concepts of fractionation, physical decompositions (of the payload antennas) and information integration (of the signals from the antennas). Sections 2.1 to 2.4 address the measurement of coastal salinity by remote sensing and a proof of concept sizing of the instrument. An exploration of the associated fractionated system is presented in section 2.5. Chapter 3 builds on the lessons learned from the development of the fractionated coastal salinity payload to answer the second question and subsequent sub-questions given in section 1.2.1. The development of a model which simulates the 50-year lifetime of generic fractionated architectures when subjected to random failures is presented. A local search algorithm is used to find architectures which maximise operational time while minimising mass launched. Chapter 4 presents the results generated from the use of the model described in chapter 3, and is broken down into three sections. Section 4.1 describes the initial data exploration using two, simpler versions of the computer model to gain some understanding of how it behaves. Section 4.2 describes the results from the final version of the model under different simulation conditions. Section 4.3 presents a discussion of the outcomes of the simulation and optimisation work. Chapter 5 gives the final conclusions of this thesis, based on the findings of the research presented.

2. Fractionation: A mission enabling technology?

The European Space Agency's (ESA) Soil Moisture and Ocean Salinity (SMOS) mission uses the fractionated principles of physical decomposition and information integration in order to generate a useful data product [18]. The Microwave Imaging Radiometer using Aperture Synthesis (MIRAS) instrument on the SMOS spacecraft consists of a number of radiometric antennas mounted on a single structure. The measurements from these antennas are combined to produce a brightness temperature image using a principle known as *Aperture Synthesis* [18]. SMOS takes measurements of Sea Surface Salinity (SSS) in the open ocean where a high radiometric resolution is the priority. SSS measurements in coastal waters require spatial resolution to be prioritised over radiometric resolution. To achieve an improvement in spatial resolution over what is currently achieved by the SMOS spacecraft, a larger aperture has to be synthesised. To achieve this increase in aperture size, more antennas are required, mounted on free flying satellites. The use of a payload that utilises information integration and requires the decomposition of the payload onto separate spacecraft made the measurement of SSS in coastal waters a good example of a mission enabled by satellite fractionation, as this allows the subsystem architecture to be shifted away from satellites hosting the payloads making them smaller and more functional. The objective of the research presented in this chapter was to firstly conduct a proof of concept study for the instrument and then to explore a fractionated satellite bus to support the operation of the payload.

2.1. Sea surface salinity and its measurement

2.1.1. Why measure coastal salinity?

The salinity content of coastal waters can have a strong impact on biological and physical processes in these regions [19]. Biological processes can include biological effects on ecosystems and their function (e.g. oyster disease, nursery grounds, coastal wetlands, and corals), development of harmful algal blooms and the survival of invasive species [19]. In estuaries, the measurement of salinity is important for monitoring drinking water quality and agricultural uses of water [19]. Freshwater run off from the land can affect the strength of estuarine and coastal currents, buoyancy fluxes, the internal wave regime and the air-sea gas exchange. In turn, these physical processes can cause problems such as the movement of estuarine salt wedges altering the sediment disposition patterns and hence dredging requirements [19]. The effect that salinity can have on water density (up to 90% of total horizontal gradients [20]) can affect a ship's draft and therefore its commercial load capacity [19]. The density gradients can lead to currents of up

to two knots at the surface, which can impact upon divers, drifting debris and small boats. Search and rescue operations are also affected, where accurate predictions of surface currents are needed and *in situ* measurements will not provide the information quickly enough. In addition, the density variations can create acoustical dead zones, rendering submarines effectively invisible, and acoustic sensors could be inoperable [20].

Regular, global measurements of salinity at a spatial resolution compatible with these coastal regions could be used for a number of applications that have a fundamental bearing on the lifestyle of the world's population, as well as providing key data for climate change models [21]. Coastal salinity is characterised by large and variable salinity levels on relatively small scales [20]. This is in contrast to open ocean salinity measurement, which is characterised by small changes in salinity over large spatial scales [22]. Woody et al. recommend that *in situ* coastal salinity measurements should range in location from river inflows to shelf boundaries (100-200km offshore) [19].

2.1.2. Ocean salinity measurement by remote sensing

Traditional oceanographic methods of measuring salinity involve collecting water samples from a ship and testing them for their salinity content. While this technique has the advantage of very accurate measurements at a variety of water depths, it only provides point measurements of salinity, at a single moment in time [19]. Remotely sensed measurements, on the other hand, can provide data that cover a very wide area. By combining these measurements with other data such as ocean colour or radar derived surface currents, details of linear and non-linear bio-geophysical processes can be evaluated [19]. In addition, if the measurements are taken from a satellite system in LEO, global measurements can be made at regular intervals over a period of many years [23]. This allows long-term dynamic models to be developed using the data collected, and long-term trends can be studied over wide areas [19]. However, the main limitation of remotely sensed data is that it cannot give measurements through the water column [23].

SSS is measured from space by passive microwave instruments. The key attribute is the brightness temperature (T_b). Figure 3 and Figure 4 show brightness temperature as a function of salinity and Sea Surface Temperature (SST) at 1.43 GHz and 2.65 GHz, respectively [24]. The unit of measurement of salinity is the Practical Salinity Unit (PSU), where 1 PSU = 1g of salt per kg of water. In coastal waters, salinity values can range from 0 – 45 PSU [19]. At a given SST, the range of brightness temperatures corresponding to the range of salinity values is much higher for the 1.43GHz case, indicating that the sensitivity of brightness temperature to salinity is much greater at 1.43GHz than at 2.65GHz.

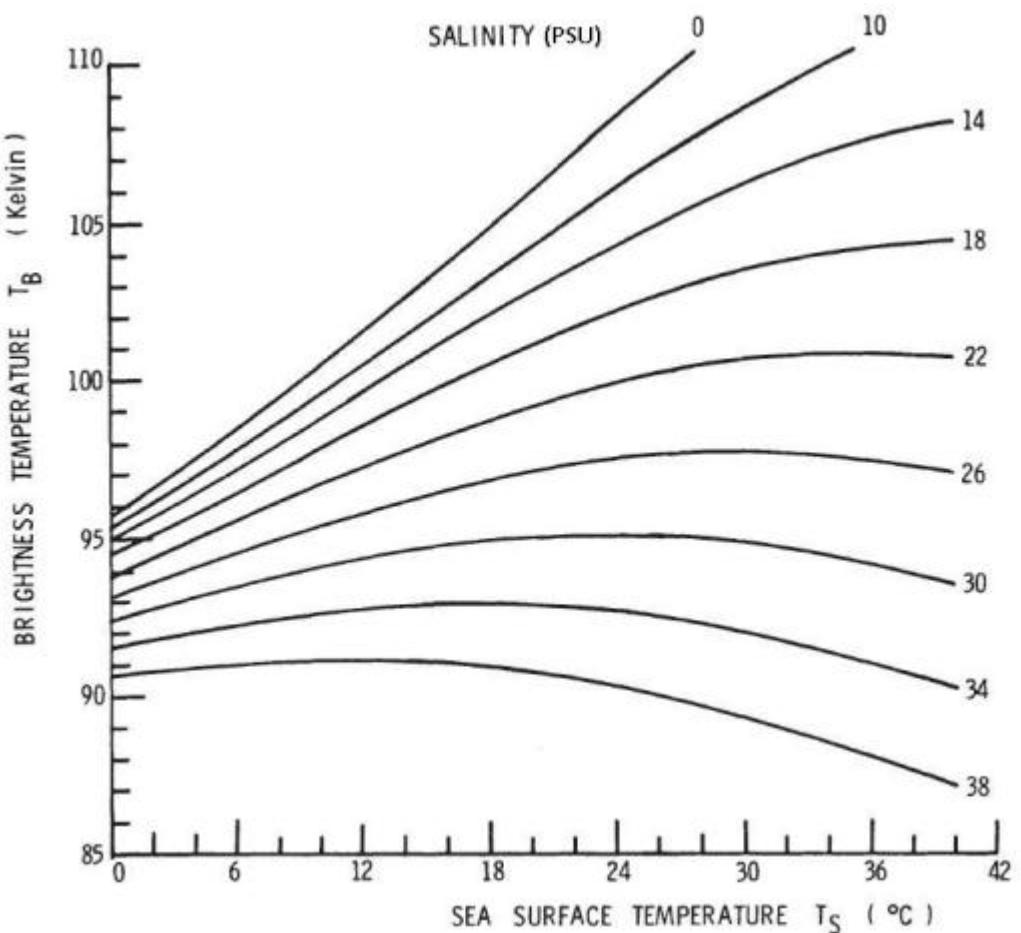


Figure 3 Brightness temperature vs. SST for various salinity levels at 1.43GHz [24]

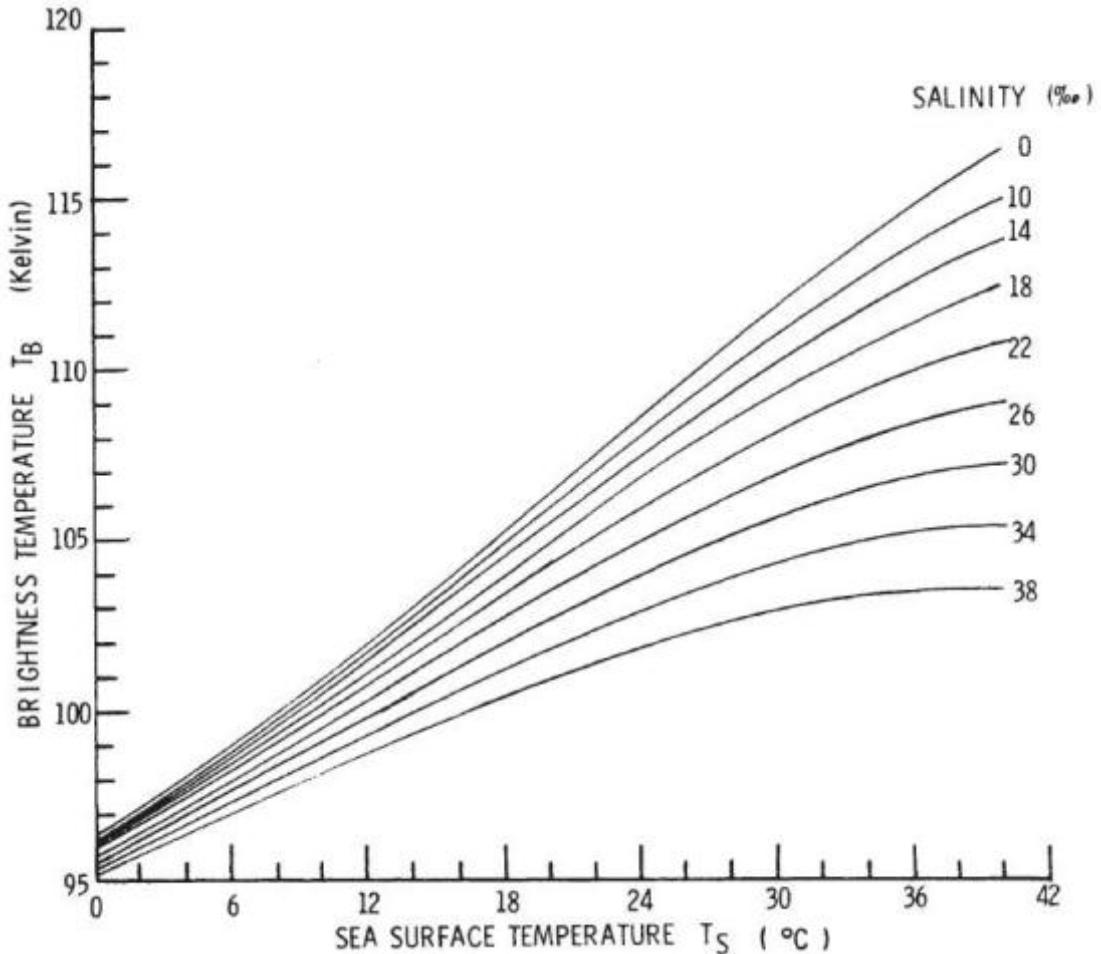


Figure 4 Brightness temperature vs. SST for various salinity levels at 2.65GHz [24]

It is for this reason that the most common frequency employed in existing radiometers is 1.43 GHz (L-Band). This gives the best possible sensitivity to salinity [18] whilst reducing the effects of atmospheric disturbances [25]. This frequency is also protected for use by Radio Astronomers. However, after the launch of SMOS, ESA found that there were several sources of interference coming from the Earth at this frequency, which were affecting the images generated [26]. Even at this frequency, the sensitivity of brightness temperature to salinity is still low [18].

The brightness temperature measured at L-Band frequencies represents the contribution of several factors, including the polarisation of the observation, the incidence angle of the observation, the SST, the surface roughness and the SSS [27]. Polarisation and incidence angle of observation are easily controlled parameters through the instrument and platform design. In order to account for the effects of SST and surface roughness, secondary instruments are required to provide measurements of these parameters [28]. In addition to these parameters, Faraday rotation in the Ionosphere affects the polarisation vector. This can be mitigated by avoiding noon-

midnight orbits and making horizontally and vertically polarised measurements. The effects of sun glint can be avoided by choosing an orbit which avoids high sun angles, and Galactic noise can be mitigated by using sky maps to identify where galactic core radiation reflected off the sea surface will be in radiometer view [28].

The spatial resolution of the image generated is approximated by,

$$r \approx \frac{\lambda l}{D} \quad 1,$$

where λ is the wavelength of measurement, l is the distance from the Earth to the instrument and D is the diameter of the antenna [29]. Therefore to achieve a resolution of 5000 m from an altitude of 800 km using an L-Band instrument, the antenna has to be about 35 m in diameter (an area of 962 m²). Under the same conditions, 1000 m resolution would require an antenna diameter of about 171 m. Even if this size of antenna could be successfully folded to fit under the fairing of any of today's rockets, its mass would be prohibitively large for any launch vehicle. To overcome this, an aperture synthesis technique is required in which a much larger antenna size is synthesised from a number of smaller antennas. These techniques will be discussed in the next section.

2.1.3. Aperture synthesis techniques for remote sensing

In aperture synthesis, an image is generated by taking the Fourier Transform of the brightness temperature measured between pairs of antennas (baselines) [29]. Park and Kim distinguish between synthetic aperture techniques that use the motion of the platform for image reconstruction (Motion Induced Synthetic Aperture Radiometer, MISAR) and those that do not (Position Induced Synthetic Aperture Radiometer, PISAR) [30]. The MIRAS payload on SMOS is a PISAR type instrument (also known as an Interferometer). Both techniques utilise the antenna separations to generate the image.

The antenna pattern employed has a strong bearing on the image processing and on the pattern of the ground track. A "Y" configuration was chosen for the MIRAS instrument on-board the ESA mission SMOS (Figure 5), as this minimises the number of redundant baselines when the Fourier transforms are taken between the pairs of antennas [31].



Figure 5 Artists impression of SMOS [32]

This solution enables the number of antennas required for the instrument to be minimised, saving mass [33]. The spatial resolution of the synthesised array is then determined from the diameter of a circle that the outer antennas of the “Y” formation lie on. A limiting factor in aperture synthesis instruments is aliasing. This is ambiguity in the phase difference between two incoming microwaves, leading to an uncertainty in the location of their origin. It manifests itself in array imaging as spatial aliasing (Figure 6) [23]. In order to avoid aliasing in the image generated by an Interferometer there is an upper limit in the spacing of the individual antennas. This limit is given by [33],

$$\text{spacing} = \frac{\lambda}{\sqrt{3}} \quad 2.$$

On SMOS, the 69 antennas are spaced at a distance of 0.875λ (18.37cm), which is greater than the alias-free condition given by equation 2. This spacing was chosen as it represents a compromise between minimising the number of required antennas (array thinning) and ensuring an alias free swath [33], [34]. Several methods are employed to ensure that as much of the field of view is alias free, one of which is to orientate the spacecraft forward at an angle of about 32° to the nadir so that a large portion of the aliased area of the field of view is not pointing at the surface of the Earth [33] (Figure 6) [34].

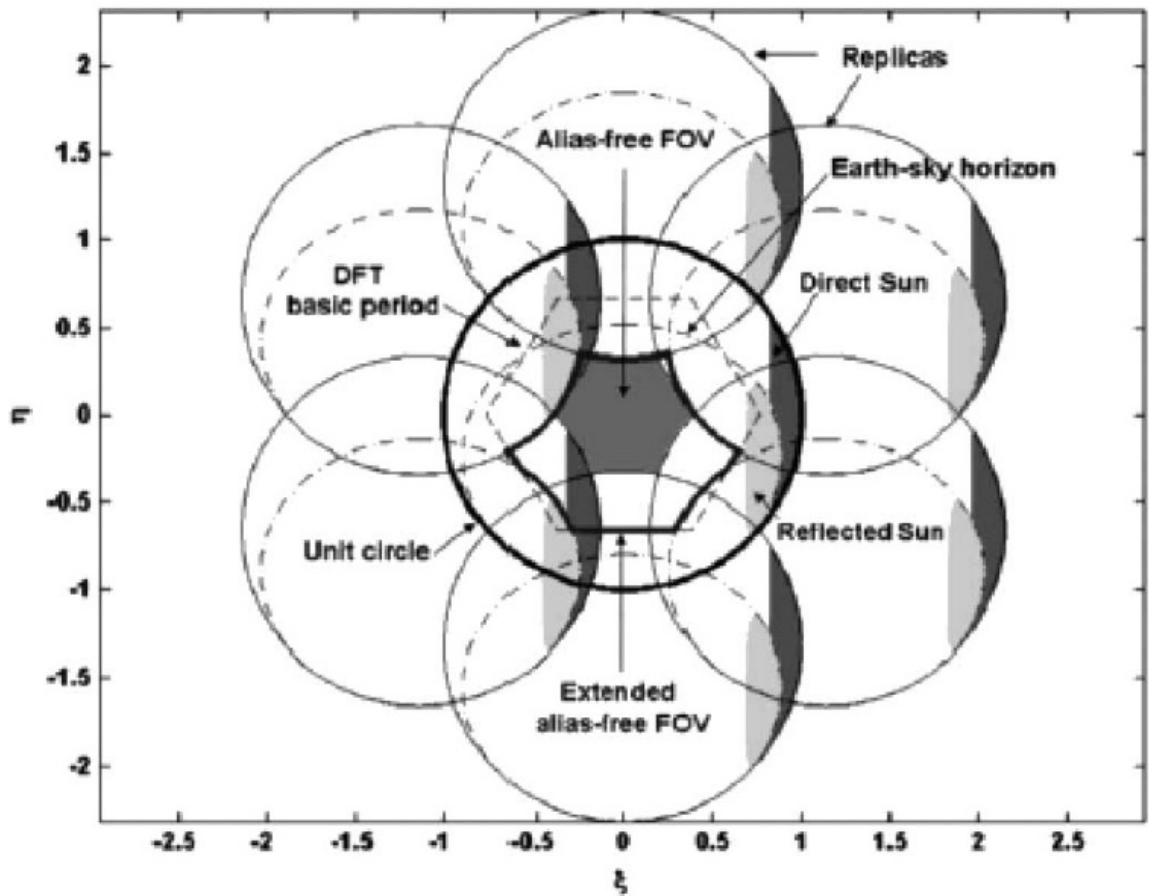


Figure 6 SMOS field of view showing aliased and un-aliased regions and the Earth-sky horizon [34]

A second method is to “average” the received data from several passes over a particular point in order to attempt to remove some of the alias effects and increase the radiometric sensitivity [18]. An increase in sensitivity is required for measuring open ocean salinity where the variation is between 33–38 PSU. The SMOS sensitivity is increased, using the averaging method, from 2.4 K to 0.1 K [29]. The main drawback of averaging is that instead of the theoretical resolution of 35 km, the final salinity data product has a spatial resolution of 200 km [27].

MISAR instruments were first proposed by Camps and Swift (2001) [35] under the name Doppler Radiometers. The concept was extended in a series of papers by a team at Gwangju Institute of Science and Technology, South Korea [30], [36], [37]. The major advantage that the Doppler Radiometer has over PISAR instruments in measuring coastal salinity, is that it enables a reduction in the number of antennas required to achieve an alias-free image [30]. This is due, in part, to the way that it samples the spatial frequencies in the scene. In the Doppler Radiometer-type instruments the spatial frequencies are sampled as *curves* across the (u, v) plane. In PISAR

instruments the spatial frequencies are sampled as *points* across the (*u*, *v*) plane. The sampling of the spatial frequencies as curves means that fewer baselines are required to cover the same area with a Doppler Radiometer than with an interferometer. For the interferometer, the spatial frequencies are calculated using,

$$u = \frac{D_x}{\lambda} \quad 3,$$

$$v = \frac{D_y}{\lambda} \quad 4.$$

where D_x and D_y are the separation, in meters, between pairs of antennas in the *x* and *y* directions [19]. Determining the spatial frequencies for the Doppler Radiometer requires knowledge of time and source position as well as antenna spacing. The observation scenario for the Doppler Radiometer is illustrated in Figure 7.

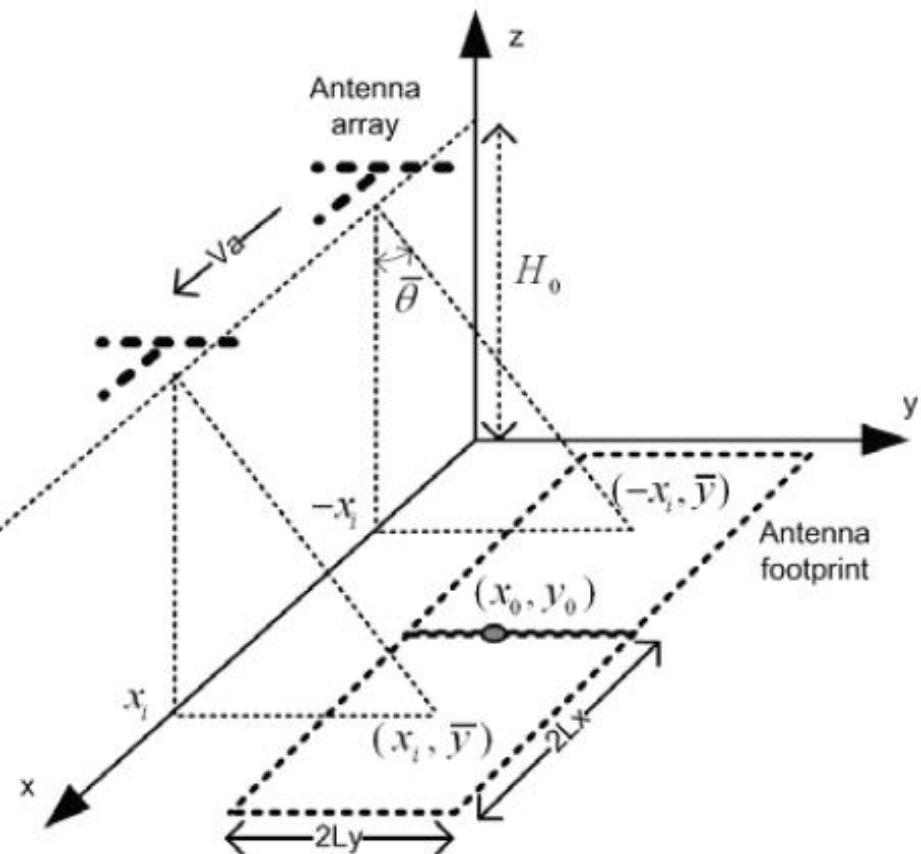


Figure 7 Doppler Radiometer observation scenario where *x* is the along track direction, *y* is the across track direction and *z* is directed upwards from the ground, V_a is the speed of the antenna array, H_0 is the altitude, and θ is the angle which the antennas are inclined to the nadir, in the *y* direction (“look -angle”)

[30]

The antenna footprint is $2L_y$ wide and $4L_x$ long. Measurements are taken from $-x_i$ to x_i , \bar{y} is the centre line of the footprint along the y axis and (x_0, y_0) is the brightness source under observation by the instrument [37]. Using this observation scenario, the spatial frequencies for the Doppler Radiometer are, [37]

$$u = \frac{\sqrt{x_0^2 + y_0^2 + H_0^2}}{\lambda} \frac{\partial \{r_k(t; x_0, y_0) - r_l(t; x_0, y_0)\}}{\partial x_0} \quad 5,$$

$$v = \frac{\sqrt{x_0^2 + y_0^2 + H_0^2}}{\lambda} \frac{\partial \{r_k(t; x_0, y_0) - r_l(t; x_0, y_0)\}}{\partial y_0} \quad 6.$$

The spatial frequencies are calculated for time steps from $-t_i$ at position $-x_i$, until time t_i at position x_i . This is the time taken for the brightness source to traverse the main beam of the antenna. Thus, the total length of the image is $2L_y$. In the observation scenario presented above, $x_0=0$ (i.e. the instrument is not looking forwards or backwards). In this case, the equations for the spatial frequencies simplify to, [37]

$$u = \frac{\sqrt{y_0^2 + H_0^2}}{\lambda} \frac{\{v_a t y_0 D_y - (y_0^2 + H_0^2) D_x\}}{(v_a^2 t^2 + y_0^2 + H_0^2)^{3/2}} \quad 7.$$

$$v = \frac{\sqrt{y_0^2 + H_0^2}}{\lambda} \frac{\{-v_a t y_0 D_x - (v_a^2 t^2 + H_0^2) D_y\}}{(v_a^2 t^2 + y_0^2 + H_0^2)^{3/2}} \quad 8.$$

Examples of Interferometer and Doppler Radiometer (u, v) plots for 16 antennas arranged in a “Y” formation (five per arm, and one in the centre) are shown in Figure 8 and Figure 9, calculated using 7 and 8. The largest curves in Figure 9 are formed by the longest baselines in the formation.

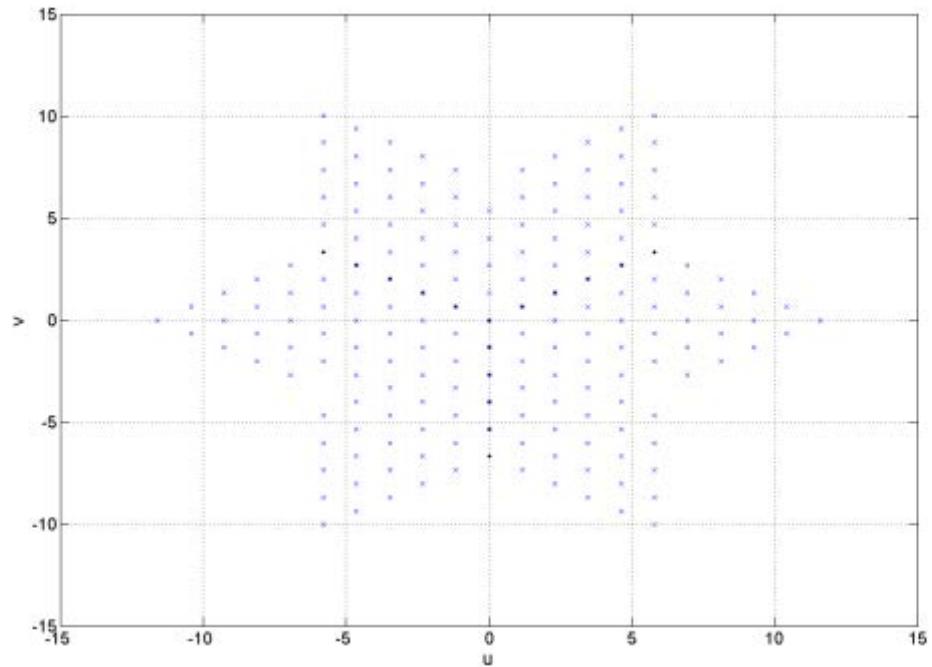


Figure 8 An example of a spatial frequency diagram for an Interferometer with 16 antennas arranged in a “Y” formation (five per arm, and one in the centre)

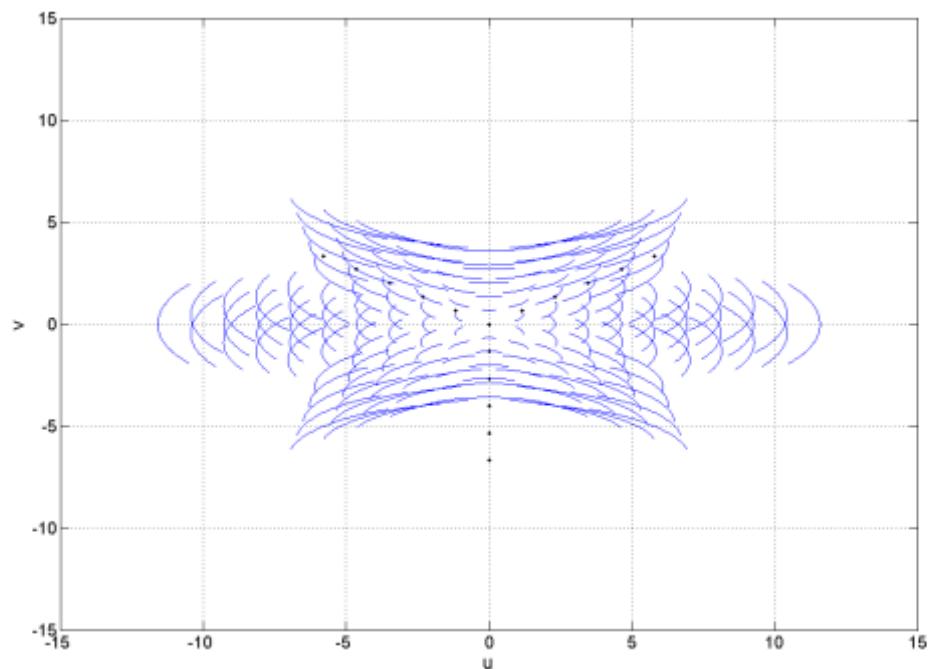


Figure 9 An example of a spatial frequency diagram for a Doppler Radiometer with 16 antennas arranged in a “Y” formation (five per arm, and one in the centre)

2.1.4. Other ocean salinity measurement systems

Remotely sensed salinity measurements have been made from space since the 1970's when an L-Band radiometer experiment was flown on Skylab. However due to the need for a very large antenna to make measurements from space with high spatial resolutions (of less than 50 km), remote sensing systems have largely been based on airborne platforms since then [38]. Airborne systems such as the Scanning Low Frequency Microwave Radiometer (SLFMR), have been used extensively for localised studies of salinity in coastal waters [39], [40]. The SLFMR is an L-Band radiometer, which can provide salinity measurements at an accuracy of about 1 PSU after calibration, at a spatial resolution of 0.5 km from an altitude of 1.5 km. An infrared radiometer is used in conjunction with the microwave radiometer to provide simultaneous SST measurements [39].

Dedicated salinity measurement missions have become achievable as synthetic aperture techniques have been developed and the possibilities for launching large antenna radiometers have increased. The first of these systems was the SMOS mission which was launched on the 2nd of November 2009 [32]. One of the first calibrated images returned by SMOS is shown in Figure 10 [41].

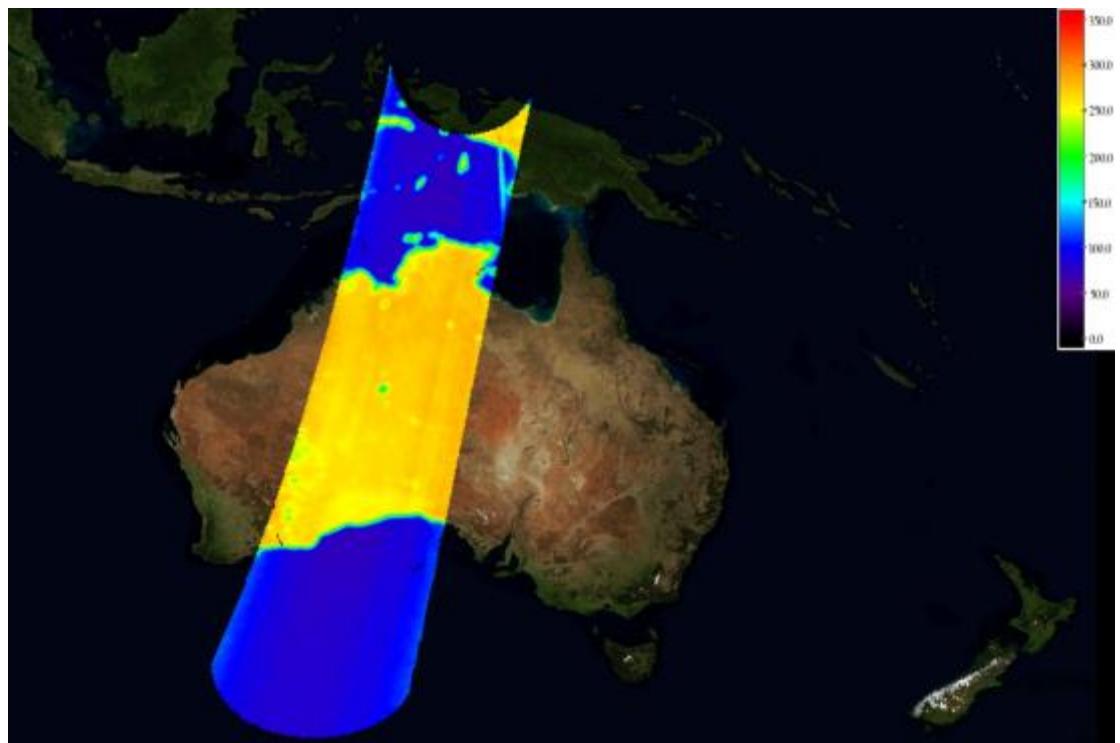


Figure 10 Early calibrated brightness temperature image from SMOS [41]

A second space borne instrument was a National Aeronautics and Space Administration (NASA) payload called Aquarius, which was launched in 2011 on the Argentinian satellite SAC-D [42], [43].

It incorporates a passive L-Band radiometer for salinity measurements and an active L-Band scatterometer for sea surface roughness measurements. It employs a real aperture antenna, 3 m in diameter, to make measurements at an accuracy of 0.2 PSU with a spatial resolution of 100 km.

2.1.5. Conclusions

Both SMOS and Aquarius produce salinity data products with spatial resolutions that are significantly poorer than those required for use in studying coastal waters. The spatial resolution these two systems can only provide one or two pixels to cover the 100-200 km extent of the coastal region. As explained in section 2.1.3, the spatial resolution of SMOS is sacrificed to gain an increased radiometric sensitivity, which is required for open ocean salinity measurements. As coastal salinity varies much more than open ocean salinity (0-45 PSU), such stringent radiometric resolution requirements are not required [19]. The trade-off has to be made between achieving a spatial resolution that is compatible with coastal waters, and a sufficient radiometric sensitivity to ensure that the salinity level is determined to an acceptable accuracy. Even simply extending SMOS to achieve a 10 km spatial resolution presents potential technical problems. Such a system would require 97 antennas per arm, and would increase the diameter of the SMOS instrument to over 33 m. This could present a significant challenge in terms of the mechanisms and structure required to deploy and support the instrument as well as increasing the satellite mass and volume.

Therefore, what is proposed is to utilise the Doppler Radiometer principles. Instead of mounting all the antennas onto a single platform, they will be mounted onto free flying satellites. As a single baseline of the Doppler Radiometer covers significantly more of the (u, v) plane than the interferometer, the array can be thinned significantly [30]. This will in turn relax the need for many satellites flying in a very close formation and allow significantly larger apertures to be synthesised.

2.2. Design drivers and requirements

Both the Interferometric and Doppler Radiometer techniques employ the two main principles of the fractionated satellite concept at a payload level. The instrument is decomposed physically into separate receiving antennas, and the formation of an image depends heavily on the integration of their received signals. Therefore it is proposed to investigate a space system, incorporating a passive microwave instrument, which is capable of making measurements of coastal salinity. This concept will utilise fractionated concepts at payload and system levels. To achieve an acceptable spatial resolution, it is anticipated that individual receiving antennas will have to be mounted on free-flying satellites. This will present additional challenges in quantifying the effects of the

relative motion on the imaging. In addition, this adds another significant parameter into the overall trade-off for the system design, as the formation of free-flying satellites needs to provide acceptable salinity measurements, whilst being easy to control, and avoiding possible collisions. The design drivers for this instrument are as follows:

- Spatial resolution must be suitable for coastal salinity measurement
- The propulsive requirements for formation control and maintenance should be minimised

There are three main requirements placed on the instrument: the spatial resolution, the radiometric sensitivity and the revisit time. One primary requirement on the instrument has been addressed in this work, the spatial resolution achieved. If a minimum of 40 pixels is required in an image from the coastline to the edge of the coastal waters (200 km offshore, as given in section 2.1.1), the minimum spatial resolution of coastal salinity measurements is approximately 5 km. The sea surface roughness and sea surface temperature errors must also be accounted for.

Other requirements and constraints were not addressed in the instrument because it was felt that it was beyond the scope of the research in terms of its relevance to fractionated systems, however they are given here for completeness. It is expected that the sea surface temperature of coastal waters could range from 0-38 °C, and that the salinity will range from 0-45 PSU [19]. To capture this range of salinities over the wide variation of sea surface temperatures, the radiometric sensitivity requirement is 1-2 K, over a range of 85-120 K, this corresponds approximately to 2-4 PSU, over a range of 0-45 PSU [44]. There are also a few constraints on the orbit needed to mitigate some of the sources of error: the Faraday rotation in the ionosphere and sun glint, detailed in section 2.1.3. These effects can be minimised by choosing a dawn-dusk type orbit [45].

2.3.Method

The initial work undertaken has involved the exploration of the Doppler Radiometer concept to assess its suitability for making salinity measurements in coastal regions. To provide a comparison, the interferometer concept was also investigated. The eventual output from this work was an instrument concept design based upon a number of free flying antennas. This formed the basis of the system design, which incorporated the fractionated concepts described in section 1.1. This is presented in section 2.5.

The studies into the Doppler Radiometer carried out by Camps and Swift [35] and Park and Kim [30], [36], [37] have produced instrument designs, but they have not focused on the measurement of any particular parameter. The research that has been carried out as a part of this

study, has involved applying the Doppler Radiometer to the coastal salinity measurement problem outlined above.

2.3.1. Extending the Doppler Radiometer concept

To achieve the spatial resolution required, an antenna diameter of approximately 35 m is needed when making measurements in the L-Band. A real aperture antenna of this size cannot be fitted into any currently available launch vehicle and the mass would be prohibitively large. Using an interferometer instrument with individual antennas mounted on a single spacecraft would require a heavy deployment mechanism to be used. As a result, the individual antennas were mounted on free-flying satellites. The resulting array would have to be much sparser than a SMOS-type instrument to ensure that the satellites will not collide. The Doppler Radiometer concept, as explored by Park and Kim [30], allows for a sparse array, as outlined in section 2.1.3. This means that the antennas can be further apart than if a SMOS-type interferometer were used.

2.3.1.1. Viewing scenario and assessment of antenna configurations

Equations for calculating the spatial frequencies for both instrument-types and the spatial resolution, are detailed in section 2.1.3. These were computed and plotted using MATLAB. For all of the Doppler Radiometer configurations investigated, the same observation scenario was utilised; the basic configuration of this is given in Figure 7. H_0 was assumed to be at 800 km. This altitude was chosen as a typical remote sensing orbit, giving sufficient orbital height to avoid significant atmospheric drag, but low enough to be capable of receiving sufficient microwave power to make measurements. V_a is the ground track speed of the sensor. x_0 is at the centre of the along-track axis, and so is set to zero. The y -coordinate of the microwave source, y_0 is dependent on the look-angle of the instrument θ ,

$$y_0 = H_0 \tan \theta \quad 9.$$

The beam-width of the instrument in the x -direction was assumed to be 30° and in the y direction 20° . These parameters were the same as those assumed in [30], [36], [37]. The frequency of observation was 1.43 GHz, corresponding to a wavelength of 0.21 m. This was used at the frequency most sensitive to salinity. The calculations were performed assuming a flat, planar geometry. That is, the effects of the curvature of the Earth were not taken into consideration as per [30], [36], [37]. The only exception to this was for the calculation of the ground-track speed.

When the spacing of antennas was determined, the diameter of each antenna was taken into account. The value assumed was 0.163 m, which corresponds to the diameter of the antennas

used in the MIRAS instrument aboard SMOS [46]. The alias-free spacing requirement is assumed to be taken from the edge of one antenna to the next.

The antenna configurations were assessed by producing (u, v) diagrams, and calculating the percentage of the frequencies sampled in the image using a pixel counting technique.

2.4. Results

The first formation evaluated consisted of a central hub, with four antennas per arm, spaced at the maximum alias-free distance, 18.37 cm. This was chosen as it was the formation studied during a project conducted at undergraduate level [47]. The rest of the formation was made up of 10 free-flying satellites per arm, spaced at 1 m intervals, with each satellite carrying two antennas spaced at the maximum alias-free distance, as shown in Figure 11.

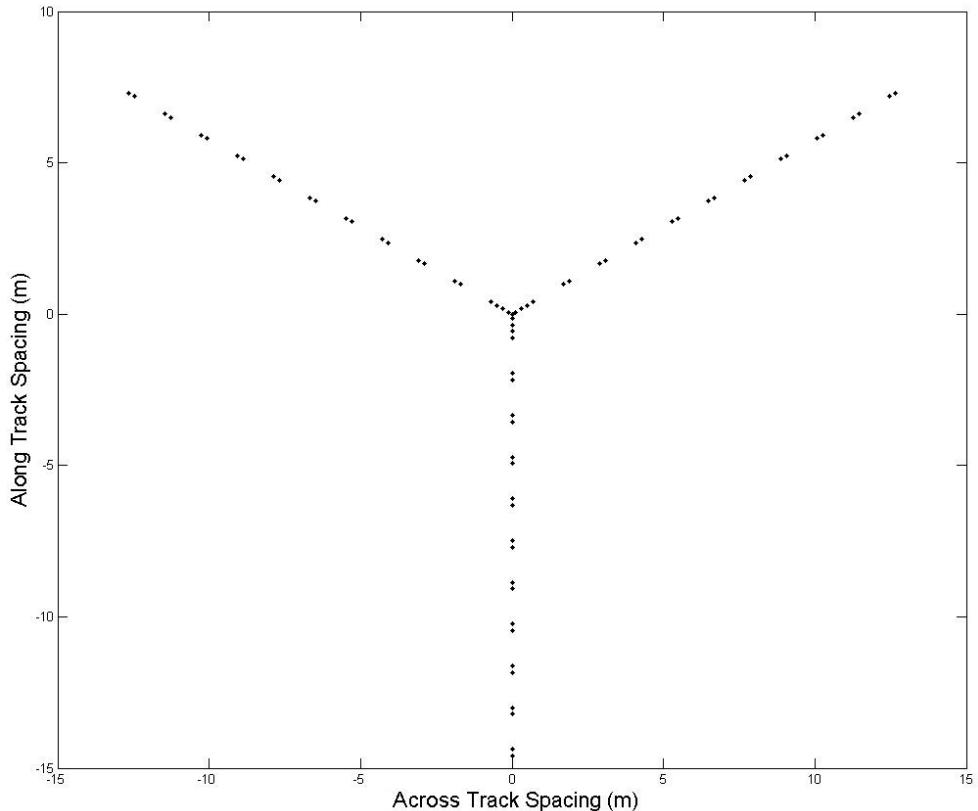


Figure 11 Free flying antenna formation. Each point represents an individual antenna

The formation shown in Figure 11 gives a maximum spatial resolution of 5.9 km. The spatial frequency diagram for a Doppler Radiometer using this antenna configuration is shown in Figure 12.

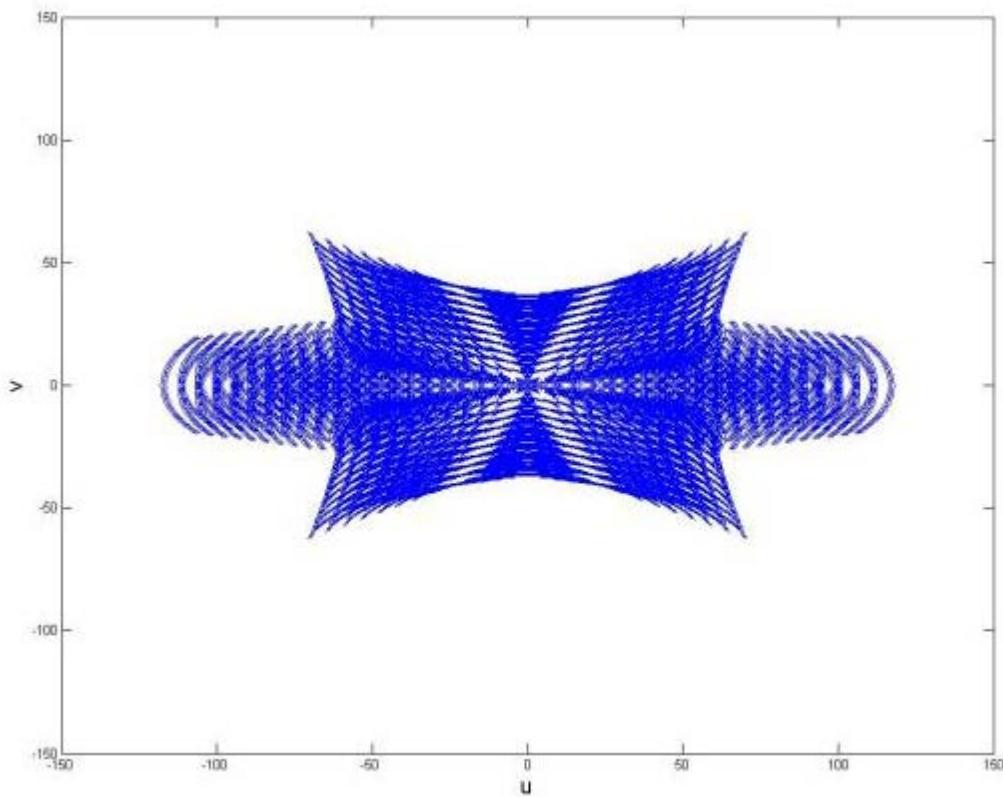


Figure 12 Free-flying antenna formation as a Doppler Radiometer spatial frequency diagram. Spatial frequencies sampled as curves

At the centre of this formation, an area of dense spatial frequency coverage was found. In addition, the baselines formed between the central hub and the other satellites, produce the denser areas of coverage along the $u = 0$ line and the diagonals of the diagram. In this formation, 40.25% of the spatial frequencies in the scene are sampled by the instrument, calculated using a pixel counting technique.

As a result of the dense spatial frequency coverage at the centre of the sampled scene, formations that consisted of a central hub and a number of antennas on free-flying platforms were further investigated. Instead of a small central hub, however, a hub of a similar size to that of the MIRAS instrument on SMOS was investigated. This enabled a larger portion of the coverage area to be covered by the baselines formed between antennas on the hub, and a larger area to be covered by the baselines formed between the hub and the outer antennas. If a "Y" formation with no central hub was employed, and the free-flying antennas are spaced at a significantly greater distance than the minimum alias-free distance, the spatial frequency coverage would be too sparse, giving an aliased image. A formation without a central hub would therefore present no advantage over a formation with a hub at the centre.

The central hub used consisted of a central “Y” shaped hub, with 21 antennas per arm, spaced at the maximum alias-free distance. On its own the central hub would provide a maximum spatial resolution of 14.37 km, and could sample 42% of the spatial frequencies in the scene. The free-flying antenna element of the formation consisted of five antennas per arm, assumed to be individually mounted on separate satellites. Three different separations of the free flying antennas were investigated (1 m, 3 m and 10 m) to investigate the effect of improving the spatial resolution (by increasing antenna separation) on the spatial frequency coverage achievable by the instrument. Figure 13 shows one of these formations, with the free-flying antennas spaced at 3 m. The spatial frequency diagram for this formation is shown in Figure 14. This configuration provided a spatial frequency coverage of 20.45% of frequencies in the scene. This is a sparser overall coverage than that provided by the formation shown in Figure 12, however the spatial frequency coverage in the centre of the diagram, is much denser, provided by the central hub.

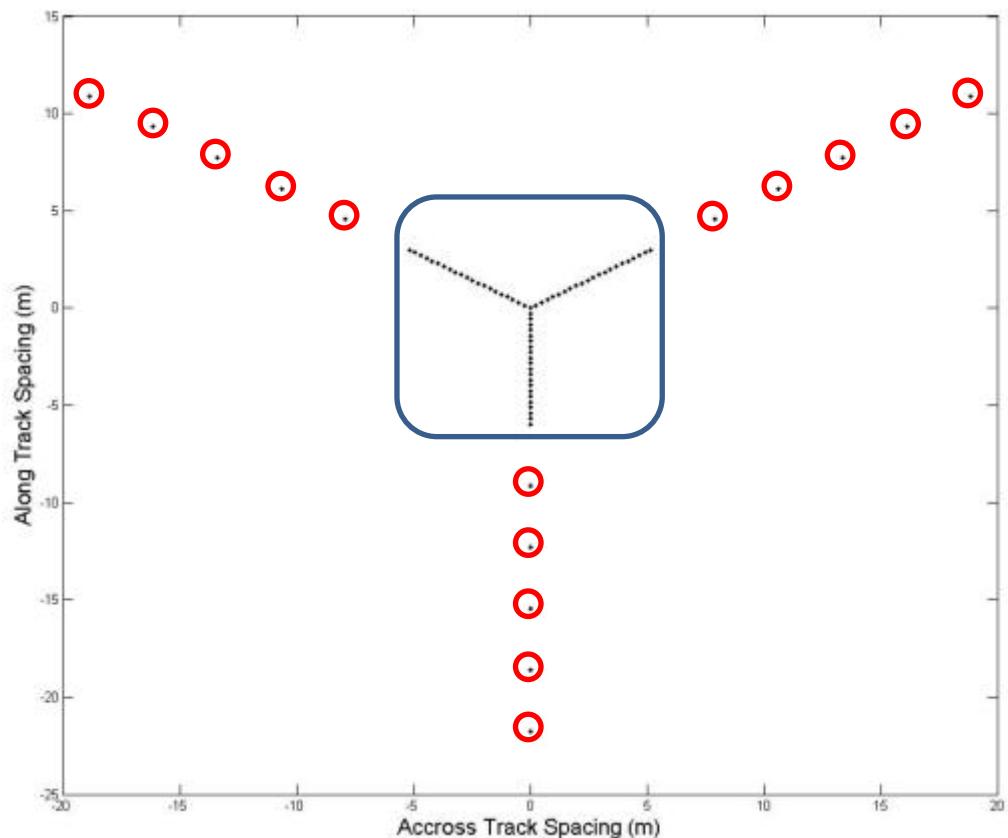


Figure 13 21 antenna formation: central hub (blue square) with 5 free-flying antennas per arm spaced 3 m apart (red circles)

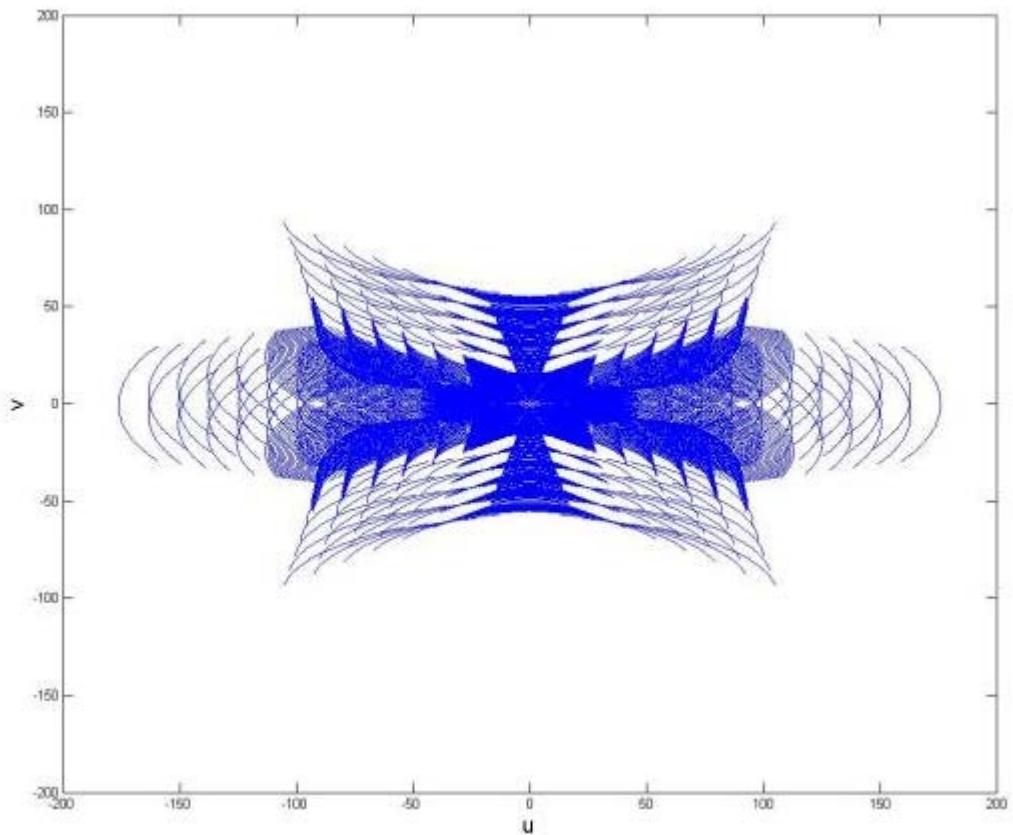


Figure 14 Doppler Radiometer, 21 antenna central hub with 5 free-flying antennas per arm spaced 3 m apart

Table 1 gives the achievable spatial resolution and the spatial frequency coverage for the central hub on its own, and for the central hub using free-flying antennas spaced at 1 m, 3 m and 10 m. This table shows that increasing the spacing of the free-flying antennas improved the achievable spatial resolution, but also had the effect of reducing the spatial frequency coverage, which would result in an aliased image.

Antenna Configuration	Achievable Spatial Resolution	Spatial Frequencies Sampled in Scene
Central hub only	14.37 km	42%
Central hub plus five free flying antennas per arm spaced at 1m apart	7.28 km	32%
Central hub plus five free flying antennas per arm spaced at 3m apart	3.94 km	20%
Central hub plus five free flying antennas per arm spaced at 10m apart	1.5 km	16%

Table 1 Spatial resolution achievable and spatial frequency coverage for different combinations of antennas mounted on a central hub and with free flying elements

One of the major factors affecting the imaging is the relative motion of the free-flying antennas. There are two aspects to this that were considered: the impact that this motion has on the imaging and, secondly, the requirement that the free-flying platforms must not collide with one another as the formation changes over the orbit.

In the “Y” configuration array, all the antennas need to be at the same orbital altitude to avoid any small velocity differences between the free-flying elements which would cause the formation to drift apart over time, assuming that the orbits are circular. The relative motion of the free-flying elements is very similar to the “pendulum” formation, in which the orbits of two spacecraft are separated by right ascension of the ascending subsystem and, if needed, true anomaly [48] (Figure 15).

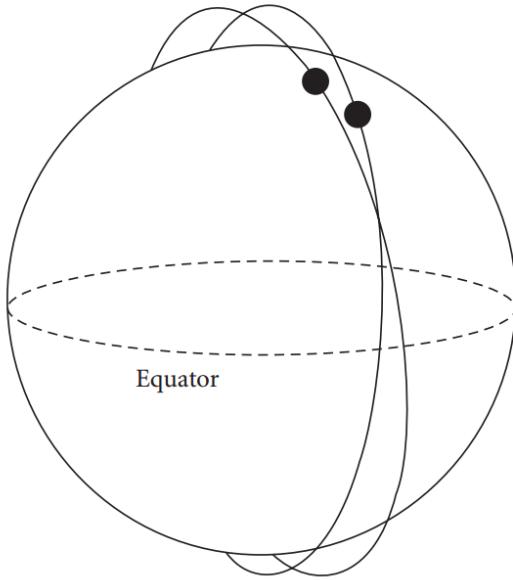


Figure 15 Geometry of the pendulum formation [48]

There can be no inclination difference if the J_2 perturbations on the orbits are to be the same. The relative position of one satellite with respect to the other when in the pendulum formation is given by

$$\begin{bmatrix} x(t) \\ y(t) \\ z(t) \end{bmatrix} \cong \begin{bmatrix} a \sin i \Delta\Omega \cos(nt) \\ a(\Delta\theta + \Delta\Omega \cos i) \\ 0 \end{bmatrix} \quad 10,$$

where, $x(t)$ is the along track separation between the two bodies as a function of time, $y(t)$ is the across track separation as a function of time, $z(t)$ is the altitude separation as a function of time, a is the semi major axis of the bodies, $\Delta\theta$ is the difference in true anomaly between the two bodies, $\Delta\Omega$ is the difference in right ascension of the ascending subsystem between the two bodies, i is the inclination of the orbits of the bodies, t is time round the orbit, and n is the mean motion of the orbit [48]. These were implemented in MATLAB to visualise the motion of "Y" formations, with the spacing of the free-flying elements used above (1 m, 3 m and 10 m), and to evaluate the suitability of this type of formation for use as a potential array formation. The relative motion of the formation used for Figure 14 is shown in Figure 16. At time $t = 0$ s the formation is over the equator, and is in its "nominal" configuration. This is an idealised situation in so far as the size of the satellite surrounding the antennas is not taken into account. The red components of the diagrams represent the central hub and the black components represent the free-flying elements.

Between time steps $t = 1490$ s and $t = 1550$ s the free flying elements from each arm cross over one another (at the poles). It is at this point that the risk of collisions between these elements arises. Only the configuration which did not experience these collisions was the one using the 10 m separation in the free flying elements. This collision risk made this formation unsuitable for use.

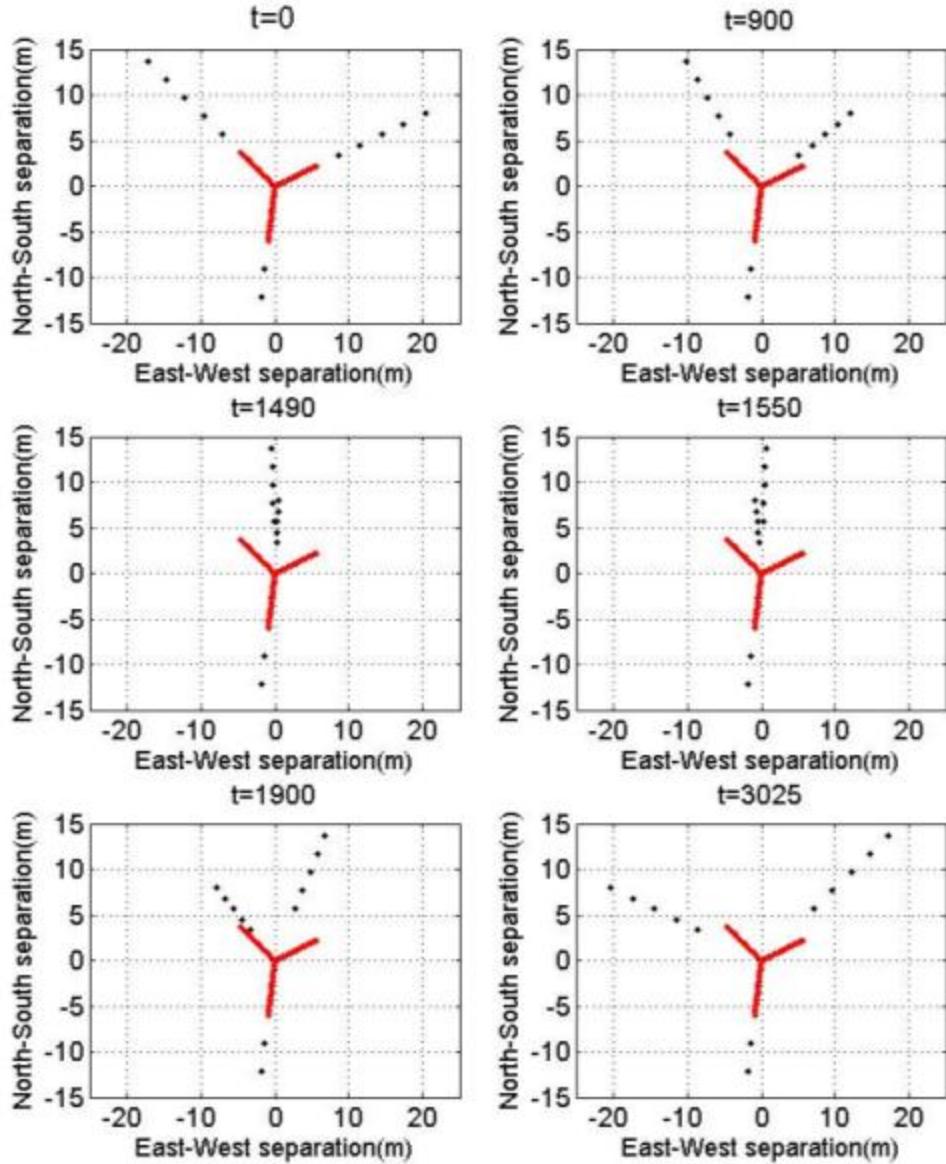


Figure 16 Free-flying antenna motion of a formation consisting of a 21 antenna central hub with 5 free-flying antennas per arm spaced 3 m apart

A circular formation was investigated as a formation that may naturally hold its shape over an orbit to avoid the collision problem identified above. In this formation, all the antennas are on free-flying platforms, arranged in a circular configuration. This kind of formation can be implemented as a “Projected Circular Orbit” (PCO) [49]. The relative motion of the free-flying spacecraft is given by,

$$\begin{bmatrix} x(t) \\ y(t) \\ z(t) \end{bmatrix} = \begin{bmatrix} c_1 \cos(nt + \alpha_0) \\ -2c_1 \sin(nt + \alpha_0) + c_2 \\ c_3 \cos(nt + \beta_0) \end{bmatrix} \quad 11,$$

where α_0 and β_0 are the initial in- and out-of-plane phase angles and n is the mean motion. In order to achieve a PCO orbit the following conditions must be met: $c_1 = \frac{d_{ref}}{2}$, $c_2 = d_{ref}$ and $c_3 = 0$ where d_{ref} is the diameter of the circular formation [49].

Figure 17 shows the evolution of a PCO formation consisting of 11 spacecraft on the circumference of a 20 m diameter ring, illustrated as 10 black data points and one that is red. The position of the red data point illustrates the relative movement of the spacecraft formation.

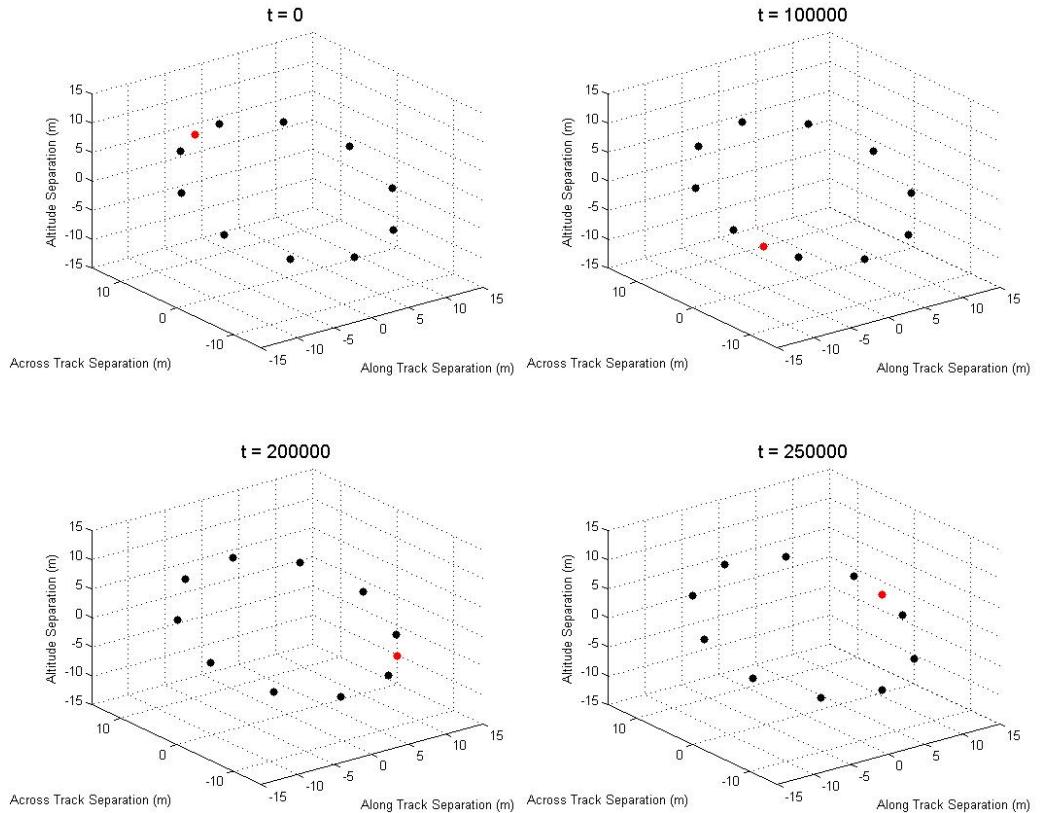


Figure 17 Projected circular orbit relative motion for a 20 m diameter formation

Figure 18 shows the spatial frequency coverage provided by a Doppler Radiometer which consists of a single 15 m radius ring with 36 antennas equally spaced around the circumference. This number of antennas was chosen arbitrarily to allow the performance of the system to be investigated. These are mounted on free-flying satellites using a PCO formation. The spatial resolution achieved by this formation is 5.7 km. This formation provides 19.4% coverage of the

spatial frequencies, approximately the same coverage as the "Y" configuration shown in Figure 14. The coverage is more uniform than the "Y" configuration, with the outer regions as densely covered as the inner ones. As a result, formations which had concentric rings of antennas were investigated.

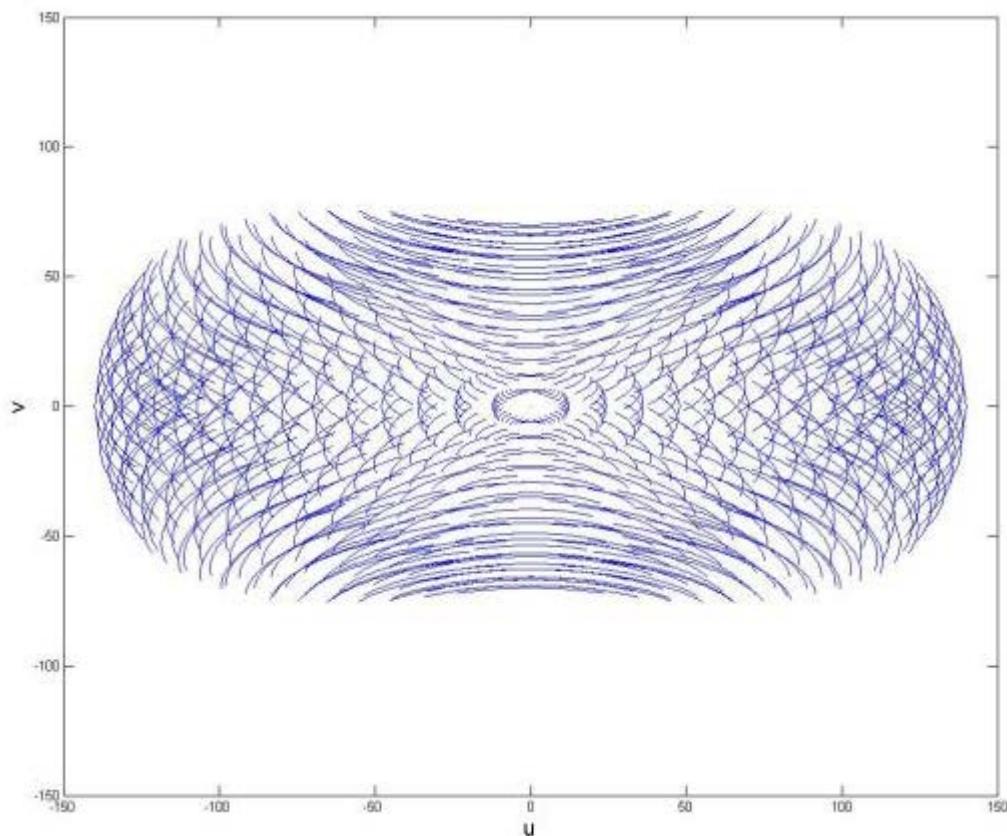


Figure 18 Spatial frequency diagram for a circular formation Doppler Radiometer, with 36 antennas in a single 15 m radius ring

Figure 19 shows the spatial frequency coverage for a formation consisting of two rings of antennas, a 15 m ring of 36 antennas around its circumference and an additional inner ring of radius 7.5 m, also with 36 free-flying antennas around its circumference. In this configuration, the level of spatial frequency coverage was 41.9%. The extra coverage is the result of baselines between antennas on the inner ring and the outer ring, and provided additional spatial frequency coverage between $u = \pm 50$ and $u = \pm 100$ and between $v = \pm 10$ and $v = \pm 50$.

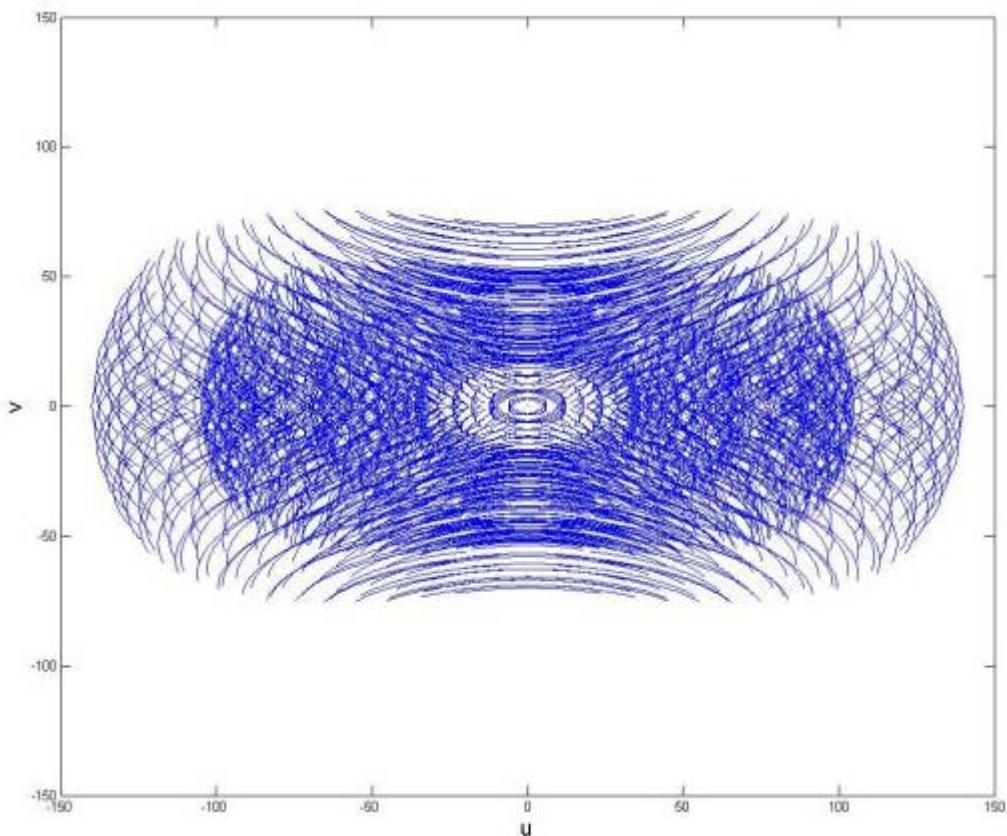


Figure 19 Spatial frequency diagram for a circular formation Doppler Radiometer, with two concentric rings of 15 m and 7.5 m Radius with 36 antennas per ring

The final iteration presented had a third ring, of radius 2 m with 36 antennas around the circumference, this time mounted on a central hub. The spatial frequency diagram for this configuration is shown in Figure 20. 49.1% of the spatial frequencies are sampled. This is higher than an equivalent sized "Y" shaped Doppler Radiometer, which can provide a coverage of only 20.45%. The formation used to generate Figure 20 requires 108 antennas in total, compared to 186 for an interferometer used to produce images at the same resolution using the maximum alias-free spacing.

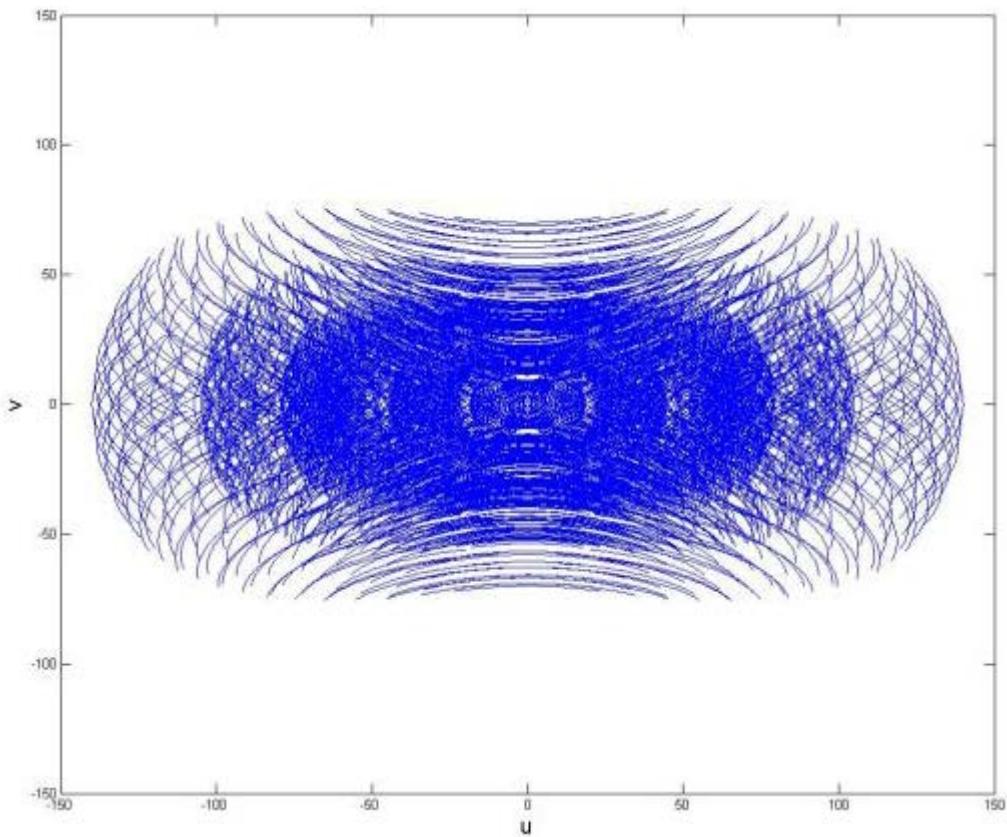


Figure 20 Spatial frequency diagram for a circular formation Doppler Radiometer, with three concentric rings of 15 m, 7.5 m and 2 m radius with 36 antennas per ring

A formation consisting of concentric rings of antennas provided a solution to an instrument that utilised free-flying elements. It allowed an alias-free image to be generated, and the PCO nature of the formation allowed for the formation shape to be maintained easily. Although the PCO formation can maintain its shape over many orbits, the formation would still be subject to small perturbations. The measurement would be robust to these perturbations as precise knowledge of the relative positions of each of the antennas is more important than their precise control. Improving the coverage by using more concentric rings of antennas showed diminishing returns. The improvement in spatial frequency coverage between 36 and 72 antennas was more than double (19.4% to 41.9%). However, adding another 36 antennas yielded a further improvement of less than 10%. This circular formation flies in a Projected Circular Orbit, as shown in Figure 21.

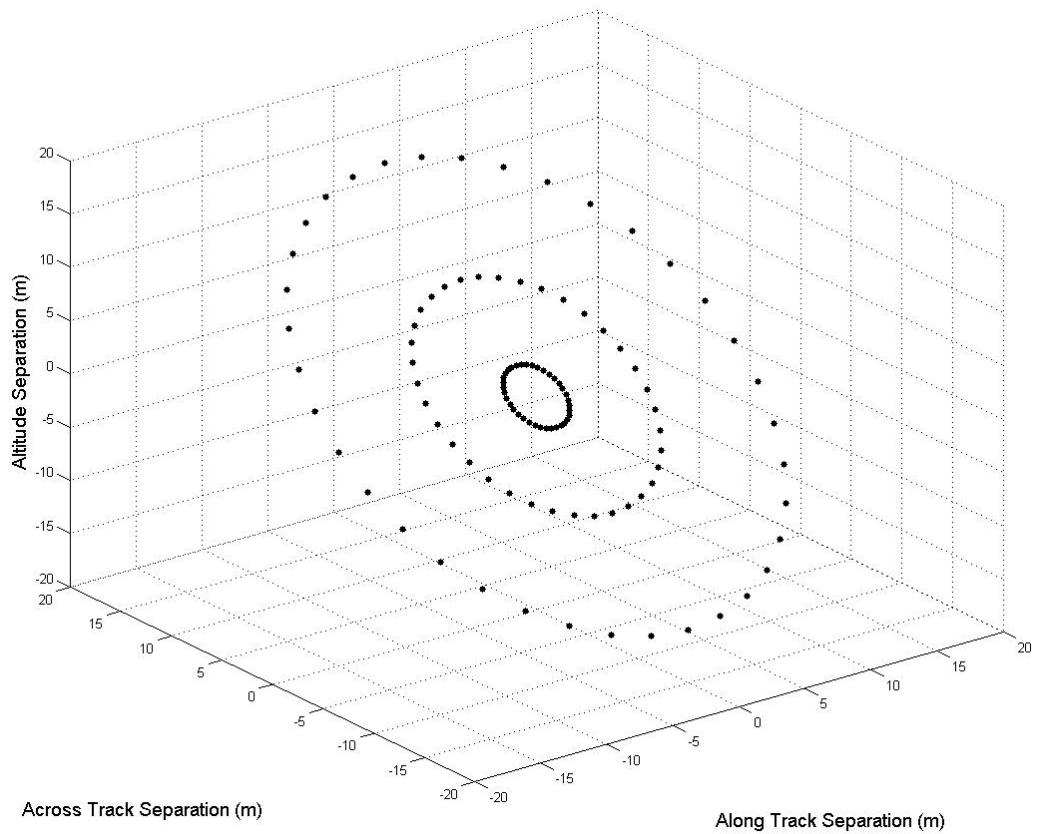


Figure 21 Circular formation, with three concentric rings of 15 m, 7.5 m and 2 m radius with 36 spacecraft equally spaced around the circumference of each ring

2.5. A fractionated coastal salinity measurement system

Presented in this section is a conceptual study into a coastal salinity measurement system using the fractionated payload described above in section 2.4. This concept shows how such a system can benefit from the use of fractionated satellite techniques and how the coastal salinity mission impacts upon the fractionated architecture. The measurement of coastal salinity is an example of a mission that is enabled by the use of fractionated spacecraft, as the type of payload required is impractical on a monolithic spacecraft. If each free-flying element was an independent satellite system, with little or no interaction with other elements, there would be a significant duplication of functionality throughout the formation. By fractionating the bus architecture used to support the payload, this duplication of functionality can be reduced. This will enable the individual payload elements to consume less power and reduce their mass [50].

2.5.1. Requirements and system architecture

The research presented here is primarily a system concept study. This section will describe the systems engineering aspects of the study; the study requirements and the system architecture employed.

The study requirements were as follows:

- The system should support a primary payload capable of making measurements of sea surface salinity in coastal waters with a spatial resolution of approximately 5 km.
- The system should support secondary payloads capable of making co-registered measurements of sea surface temperature and sea state to enable the retrieval of salinity measurements from the primary instrument
- The system should utilise fractionated subsystems, under the assumption of current technology capabilities.
- The fractionated subsystems must be capable of sharing their resources with the rest of the satellites in the architecture and be capable of fully supporting their operation.
- All the non-fractionated subsystems must be present on every satellite in the architecture.

The primary payload consisted of the Doppler Radiometer instrument concept presented in the previous section. 72 antennas were individually mounted on free flying satellites, with a further 36 mounted on a central hub satellite which also hosted a correlator unit, used to combine the signals received by the antennas. A scatterometer payload was mounted on a separate satellite to provide sea state measurements and a radiometer payload was also mounted on another satellite to provide SST measurements.

The design of the bus architecture considered the following subsystems: the ground-link communications, the on board data handling (OBDH) subsystem, attitude and orbit determination subsystem (AODS), the attitude control subsystem (ACS), the orbit control subsystem (OCS), and the power subsystem. Of these subsystems, it was assumed that the OBDH, AODS and communications subsystems were fractionated. These three functions are each uniquely hosted on free flying satellites. The remaining subsystems (ACS, OCS and power) are therefore present on every satellite in the architecture, as well as relative position sensors for attitude determination and inter-satellite communications links to facilitate the sharing of communications, OBDH, payload and AODS data. This architecture structure is summarised by the block diagram given in Figure 22.

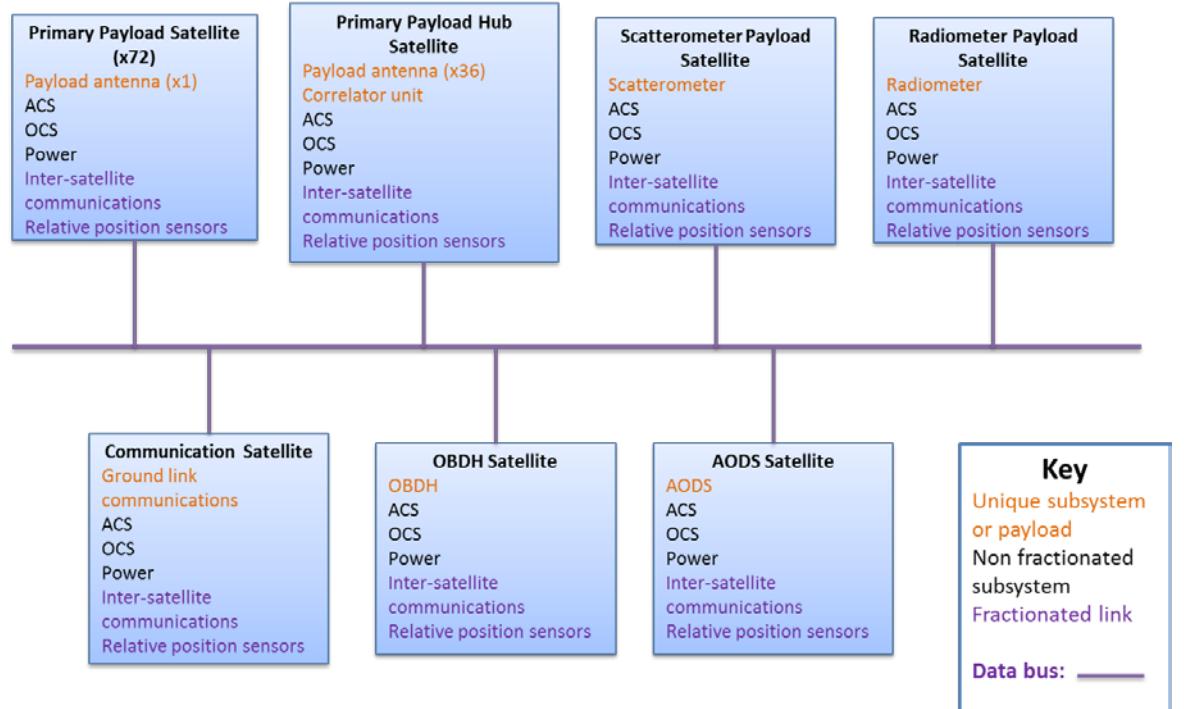


Figure 22 Architecture block diagram for the fractionated coastal salinity system

2.5.1.1. Payloads

The primary instrument was the three-ringed, circular formation concept presented in section 2.4 consisting of 108 individual antennas divided equally in three concentric rings. The three rings were at radii of 2 m, 7.5 m and 17 m and the entire formation is inclined at 45° to the nadir (Figure 21). The 17 m radius ring provided the required 5 km spatial resolution and the other antennas ensure that the scene was sufficiently sampled. The antennas in the outer two rings were each mounted on free-flying satellites, and the inner ring antennas were all mounted on a single, central hub satellite. Each antenna had a mass of 1 kg and consumed 2 W of power [46].

A correlator unit was required to combine the signals from the antennas on the hub and those from the free flying antennas to produce the brightness temperature maps. This unit was located on the central payload hub. The correlated image can then be passed onto the data handling architecture so that it can be downlinked to the ground. The correlator unit had a mass of 11.2 kg and consumed 34 W of power [34].

To retrieve the salinity from the brightness temperature maps produced by the primary payload, two additional parameters must be measured: SST and sea state [29]. These data could be obtained using instruments on other satellites, but it is better to use co-registered data generated at the same time as the salinity measurements under the same observation conditions. To make the SST measurements, a multi-channel, scanning, passive microwave radiometer was chosen,

similar to the Advanced Microwave Scanning Radiometer – Earth Observing System (AMSR-E) on-board the Aqua Satellite [51]. To achieve a comparable diffraction limited spatial resolution to that achieved by the primary payload, a reflector diameter of 4.5 m was required, when operating at 10.65GHz (calculated using equation 1). This instrument had a mass of 314 kg and consumed 350 W of power [52]. A scatterometer instrument, similar to the Advanced Scatterometer (ASCAT) instrument on the MetOp spacecraft was chosen to make measurements of the sea state [53]. The scatterometer had a mass of 260 kg and consumed 215 W of power [54]. Fractionation allowed these two secondary payloads to be much more easily integrated into the system, as they can be mounted on their own satellites. This allowed the SSS, SST and sea state measurements to be made concurrently, under the same conditions, rather than relying on other satellite systems to provide the measurements that may not be taken at the correct time or under the same observation conditions. The two secondary payloads are each mounted on individual satellite platforms, which fly outside the primary payload formation. As the relative positions of these two free-flying elements do not influence their performance, these spacecraft can fly 1-2 km away from the main formation so that they do not interfere with it.

2.5.1.2. *Satellite bus fractionation*

For this case-study, it was assumed that the system will be launched in three to five years' time, to represent the need to begin monitoring a key climate change variable quickly. This restricts the level of fractionation that can be achieved due to technological development constraints. The subsystems which are fractionated are: ground-link communications, OBDH and AODS. These subsystems were chosen because the technologies required to fractionate them have already been implemented in spacecraft in operation. The power system was not fractionated. Despite Earth based demonstrations and the recommendation of Brown and Eremenko that the development of this technology will increase the viability of satellite fractionation [5], Lafleur and Saleh indicated that it would not be suitable to provide the primary power source for a fractionated system [55]. There are, also a large number of unknowns in the technology required. In their design study, LoBosco et. al. did not consider the power system for fractionation but, instead, included this functionality as a technology demonstration only [8]. The ACS and OCS systems were also not considered for fractionation, as the technologies required are immature [5], [8].

Therefore, there were a number of satellite functions that were not fractionated for this study. These were: a primary and secondary power source (solar arrays and batteries), thermal control, basic computing and data storage required for the management of each individual satellite (depending on the OBDH architecture selected) and attitude and orbit control actuators.

The ACS subsystem on each satellite consisted of four reaction wheels (assumed to be SSTL Microwheel 100SP-M wheels for the smaller satellites and SSTL Smallwheel 200SP wheels for the satellites hosting the two larger secondary payloads) and three magnetorquers for momentum dumping (assumed to be SSTL MTR30 magnetorquers)[56]–[58]. Table 2 gives the mass and power consumption for these components.

Unit	Mass (kg)	Power consumption (W)	Reference
SSTL Microwheel 100SP-M	2 kg	2 W	[56]
SSTL Smallwheel 200SP	5.2 kg	16.3 W	[57]
SSTL MTR30	1.8 kg	1 W	[58]

Table 2 Mass and power properties for the ACS components

The OCS subsystem consisted of a SSTL Xenon propulsion system, which had a wet-mass of 20 kg and consumed 60 W of power [59]. The power subsystem was sized for each satellite individually. This consisted of a PCDU, a battery and solar arrays. The PCDU had a mass of 3.9 kg [60]. The battery which was sized to provide power to the bus subsystems in eclipse (i.e. the payload is not operated in eclipse). An 800 km circular orbit was assumed, which gave an eclipse time of 35.13 minutes [61]. The battery sizing process allowed one of five different battery sizes to be chosen. The masses of these batteries, taken from an ABSL Power Solutions Ltd. data sheet were 1.2 kg, 3.3 kg, 4.7 kg, and 15.9 kg, which corresponded to battery capacities of 130 W-h, 346 W-h, 486 W-h, 691 W-h and 1620 W-h [62]. The solar array area, A_{SA} was given by,

$$A_{SA} = \frac{P}{F_{sun} \times \eta} \quad [12],$$

where P is the power required by the spacecraft when in sunlight plus the power required to charge the battery (assuming a 30% depth of discharge). F_{sun} is the solar flux at the Earth (assumed to be 1350 W/m²). η is the efficiency of array system, assumed to be 0.2 and encapsulates the efficiency of the solar cells, the packing efficiency of the cells and the degradation factor [63]. It was assumed that the solar array would track the sun, to maintain a zero off-pointing angle. The mass of the solar array was given by multiplying the array area by 1.76 kg/m², the specific mass of a Spectrolab panel using triple junction cells [64].

The design of the fractionated bus subsystems and the satellites that host them will be presented in the following sections. These were designed in a heterogeneous way, with each fractionated subsystem housed on individual satellites [5].

2.5.1.2.1. The ground link communications satellite

The fractionated ground link communications took the form of a dedicated satellite carrying the ground-link hardware, including the transponders and antenna system. Communication occurs when the ground-link spacecraft is in view of the ground station.

To downlink payload data, the correlated image is passed from the central payload hub to the data storage satellite via an inter-satellite link. When the communications satellite is in view of the ground station, the images are transferred from the data storage satellite, again via an inter-satellite link, to the communications spacecraft where they can be transmitted to the ground (Figure 23).

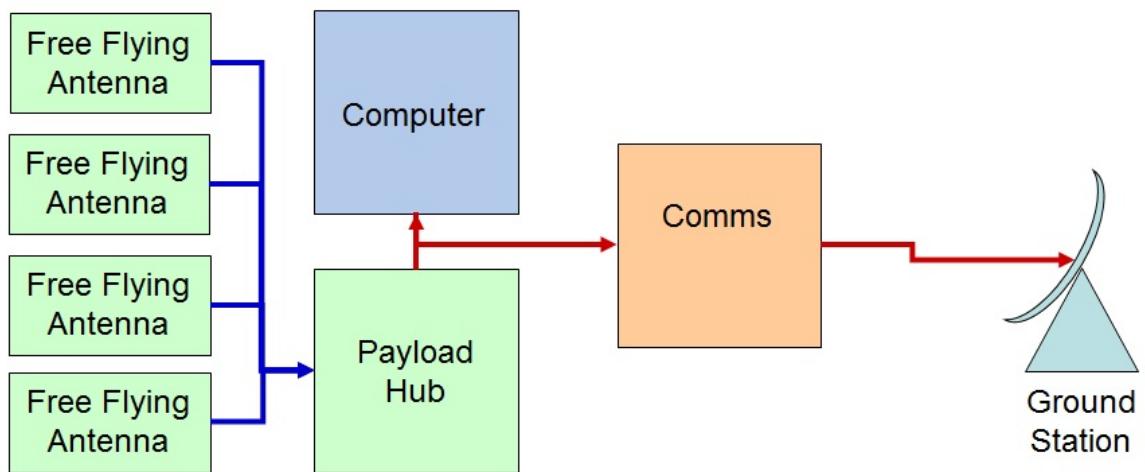


Figure 23 Primary payload data flow diagram. Each box represents a separate satellite. The blue lines represent the flow of the data collected by the individual antennas, while the red lines show the flow of the correlated images.

The technology required for fractionation of the ground-link communications is similar to that demonstrated by the Swedish PRISMA formation flying demonstration mission. Radio Frequency inter-satellite communication links was demonstrated between the two satellites, one of which had no dedicated downlink, had to communicate with the controllers through the other [65]. A similar architecture is used in the fractionated system to downlink secondary payload data and telemetry data, as well as the uplink of telecommands. Omni-directional Radio Frequency inter-satellite links would be required to avoid any conflict in pointing requirements between these links and the payload antennas. The dissemination of telecommands throughout the fractionated system and the collection of payload and telemetry data are handled by the OBDH network. The communications hardware took the form of S- and X-band units, for telemetry and payload data respectively. These were assumed to be units designed and developed by SSTL. Table 3 gives the mass and power values for the communications equipment.

Unit	Quantity	Mass (kg)	Power (W)	Reference
SSTL S-Band Downlink Transmitter	2	2 kg	38 W	[66]
SSTL S-Band Transmit Filter	2	0.09 kg	-	[67]
SSTL S-Band Uplink Receiver	2	1.3 kg	1.5 W	[68]
SSTL S-Band Receive Filter	2	0.09 kg	-	[69]
SSTL S-Band Patch Antenna	2	0.08 kg	-	[70]
SSTL S-Band High Power Amplifier	1	1.25 kg	33 W	[71]
SSTL X-Band Downlink Transmitter	2	3.25 kg	55 W	[72]
SSTL X-Band Antenna Pointing Mechanism with High Gain Antenna	2	2.7 kg	3.5 W	[73]

Table 3 Mass and power properties for the communications equipment

2.5.1.2.2. Fractionating the OBDH architecture

The OBDH subsystem on any spacecraft handles the data processing tasks for the architecture, monitors and maintains the health and safety of the architecture, and collects and disseminates data throughout the spacecraft. The OBDH hardware consisted of a SSTL OBC750 computer (1.5 kg, 10 W) and a SSTL High Speed Data Recorder (1 kg, 15 W) [74], [75]. The coastal salinity instrument requires a large amount of data to be collected from the payload antenna spacecraft and passed to the hub where they can be processed. This data transfer requires a sophisticated data handling architecture. The tight formation control requirements for the payload also impose a significant processing task on the OBDH subsystem. Therefore, the system benefited from a fractionated OBDH architecture containing a data bus that enables full communication between all spacecraft, and a processing architecture that allows flexible and robust formation control. Finally, as a consequence of using a fractionated system, an autonomous Fault Detection,

Isolation and Recovery (FDIR) strategy has to be implemented in the system to enable it to respond rapidly to a failure in a subsystem.

There were two options for the fractionation of the on-board computing: a centralised or a distributed architecture. Not all the processing tasks can be handled by a central computing spacecraft and, at least, some processing capability is required to monitor and regulate the basic functions of each spacecraft and provide safe-mode functions, regardless of the host spacecraft's primary function. As a result, the trade-off between centralised and distributed computing architecture was focused on the system-level processing tasks, rather than the spacecraft-level processing.

The main processing task was the formation control of the primary payload spacecraft. The formation of spacecraft must be maintained to ensure that the satellites do not collide and that the correct positions of the payload antennas are maintained. One way to assess the most appropriate architecture is to evaluate the flow of data in the system for the two options. This evaluation can be done by analysing the control loops involved in the task. In the centralised architecture, a main mission processor on-board a dedicated OBDH spacecraft evaluates the relative positions of all the other spacecraft, based on information provided to it from the relative position sensors on each of the other satellites. From this information all the necessary torques and manoeuvres that need to be applied by each spacecraft are calculated. This information is then disseminated to all the satellites which then apply them (Figure 24).

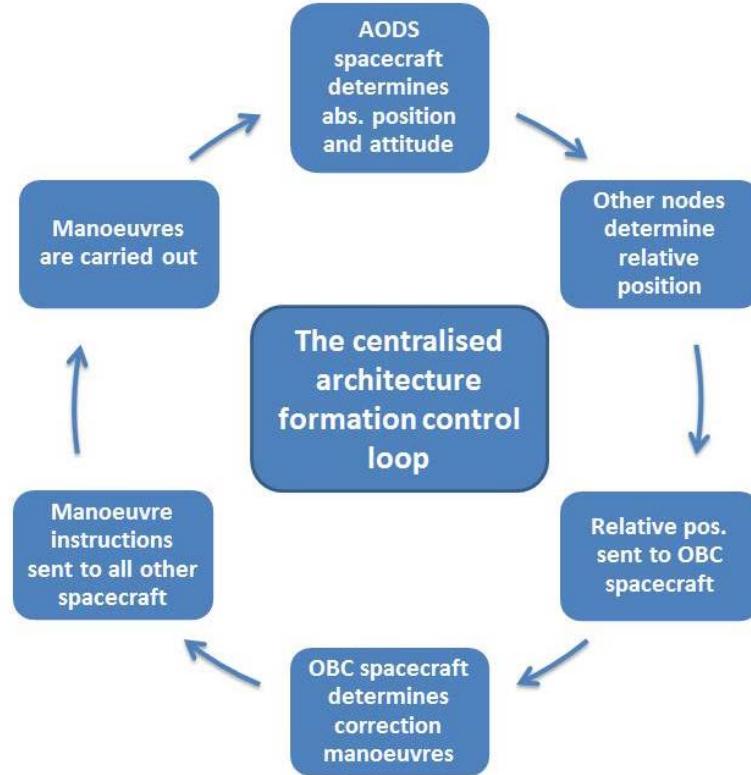


Figure 24 The centralised architecture formation control loop

In the distributed architecture, each spacecraft uses the absolute position information provided by the AODS subsystem and the relative position and attitude data it collects about its nearest neighbours to determine the correction manoeuvres required to maintain the formation, the pointing and the orbit of the satellite (Figure 25).

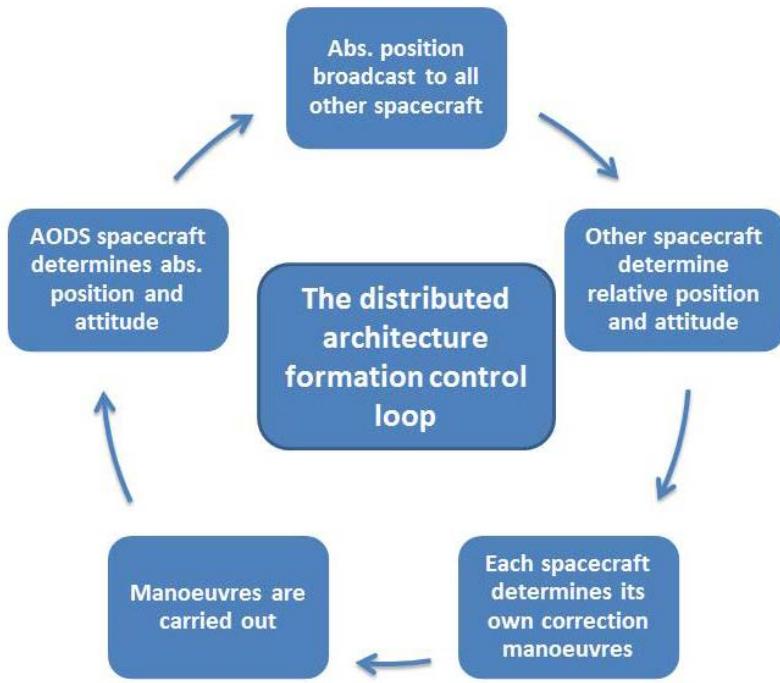


Figure 25 The distributed architecture formation control loop.

In the centralised architecture, there is considerable exchange of information between the OBDH spacecraft and all the other satellites and a heavy computational task that the OBDH subsystem must carry out before any correction torques or manoeuvres can be applied. This centralised process would be slow compared to the relatively simple process involved in the distributed architecture. Whilst the day-to-day formation maintenance required only minor corrections and low-level calculations, a debris avoidance scenario required the formation to disperse quickly and in a controlled manner or for the entire formation to move out the way of the debris. The requirement for a fast reaction time pushed the formation control architecture towards a distributed one rather than a centralised one. The distributed architecture is similar to a model of flocking birds [76], in which each bird uses its neighbours to determine where it should fly.

To fractionate the data handling architecture, the traditional wired data bus architecture, which links the subsystems together, was replaced with a wireless network architecture, which links the free-flying elements together. The use of wireless technologies on satellites has been demonstrated on-board the Dutch, Delfi-C3 CubeSat [77]. As with the processing architecture, the two main options were centralised or distributed architectures. In a centralised architecture, a single spacecraft acts as a server for the system and all communications from one OBDH subsystem to another go through this spacecraft. In the distributed architecture, every subsystem carries its own server such that each spacecraft can communicate with any other. In network

topology terms, these two architectures can be thought of as either a star network (centralised) or a mesh network (distributed) (Figure 26) [78].

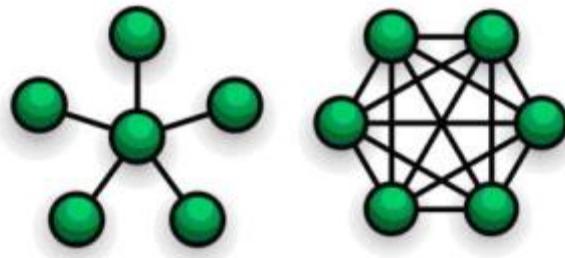


Figure 26 Star network (left) and mesh network (right) topologies

The main advantage of the centralised architecture is that it minimises the number of communications links that must be formed. The number of links is given by $n-1$, where n is the number of satellites in the network. In contrast, the distributed network requires $n(n-1)$ links, but increases the flexibility and reliability of the data architecture, as the central server spacecraft is a potential source of failure. The data architecture selected for the coastal salinity measurement mission was a distributed one. In addition to the benefits described above, this allowed for a faster collection and dissemination of information in time-critical operations, such as formation control or debris avoidance.

In addition to the computationally intensive task of formation control, there are less intensive monitoring tasks that also must be performed. These include the payload task management, the management of the architecture-level safe mode and the FDIR system, and the data storage management. The hardware and software required to implement these tasks could either be co-located on a dedicated OBDH spacecraft, or they could be distributed throughout the system on other spacecraft. By distributing these functionalities throughout other spacecraft, there is one fewer major subsystem that must be launched. However, the overall flexibility of the system will be reduced because other spacecraft will be carrying more than one mission resource, resulting in more than one functionality being lost in the event of a failure or anomaly. Alternatively, if these functionalities were co-located on a dedicated OBDH subsystem, the reliability of the architecture would be reduced; if the OBDH subsystem fails then the key architecture-level tasks of system safe mode management and the FDIR system will also be lost with potentially fatal consequences for the rest of the architecture. Consequently, the proposed solution had a dedicated OBDH satellite containing the functionalities described above, but they are also distributed on other satellites to provide redundancy. In the event of the OBDH subsystem failing, the redundant

subsystems would take over until a new OBDH spacecraft can be launched to restore the full operational redundancy of the architecture.

In the event of a failure, contact with the failed subsystem will be lost. In the event of an anomaly, the affected subsystem will enter safe mode, will broadcast that information to the other satellites and will manoeuvre itself away from the formation to avoid collisions. In the event of a loss in manoeuvring capability in a spacecraft, the rest of the formation will move away from the affected satellite. The first action the system FDIR will undertake is to halt all the payload operations. If the failed subsystem is a payload element, then no further action is needed as these do not contribute to the safety critical operations of the system. Payload operations can resume once the anomaly is rectified or the failed subsystem is replaced. If any of the bus subsystems fail, then the system is at risk. In this eventuality, the other spacecraft must be placed into safe mode, whereby each spacecraft simply maintains power on its solar array (when in sunlight) and uses its relative position sensors to ensure that it does not collide with any other spacecraft subsystem. Once the failed subsystem is replaced or the anomaly is rectified, the other spacecraft can exit safe mode and normal operations can resume.

2.5.1.2.3. Fractionating the AODS

The fractionation of the AODS requires the determination of the absolute attitude and orbit of the fractionated system via a dedicated satellite subsystem. The processing of these data and its outcomes was discussed in the previous section. Now, the focus is on the sensors used on this AODS subsystem and the relative position sensors (RPS) used throughout the fractionated system. Table 4 gives the mass and power values for the AODS equipment used.

Unit	Quantity	Mass (kg)	Power (W)	Reference
SSTL SGR-10 GPS Receiver	1	0.95 kg	5.5 W	[79]
SSTL MIRAS-1 Inertial Sensor	1	1.8 kg	5 W	[80]
SSTL 2-Axis Sun Sensor	2	0.3 kg	-	[81]
SSTL Magnetometer	1	0.14 kg	0.3 W	[82]
SSTL RIGEL-L Star Tracker	2	2.4 kg	12 W	[83]

Table 4 Mass and power values for the AODS equipment

The RPS functionality was provided by a radio-frequency ranging system and supplemented by a vision-based system similar to that employed on the PRISMA mission, which demonstrated these technologies for autonomous rendezvous and docking purposes [65]. It was assumed that the pointing accuracy required for rendezvous and docking would be comparable for determining relative positions this formation flying case. A laser-based system, such as the one proposed for the X-ray Evolving Universe Spectroscopy (XEUS) formation flying mission, was not utilised here as it would require many precisely-directed laser beams.

2.5.2. Summary and conclusions

The determination of sea surface salinity in coastal waters by remote sensing requires the use of an instrument that utilises a large number of free-flying antennas given the, otherwise, large sized structure needed. The signals collected by the individual antennas were combined to generate an image. As a result, this payload was fractionated.

The fractionated system presented here consisted of a primary payload which had 72 free-flying spacecraft, each carrying one Doppler Radiometer receiver, one central hub carrying 36 receivers and the correlator unit, used to combine the signals received by the antennas, a microwave radiometer satellite and a scatterometer satellite. The three satellite bus spacecraft and their main mission assets were as follows: an AODS spacecraft, carrying an attitude and orbit determination sensor subsystems, a communications subsystems carrying a high-speed data, telemetry and telecommand link and an OBDH spacecraft carrying the main mission processor and data storage. Table 5 gives the mass budgets for each of the satellites in the architecture. To account for the hardware required to wirelessly connect the satellites an additional 10% is added to the sum of the masses of the other subsystems on-board each satellite. To account for the thermal control hardware and the structural mass an additional 10% is added to the sum of the other subsystems on-board (including the wireless connection hardware). A further 20% margin is added to the mass to give the total mass of each satellite. Table 6 gives the power budgets for each satellite in the architecture. Again 10% is added for both the wireless connectivity hardware and the thermal control hardware as well as a 20% system margin.

A monolithic approach to coastal salinity measurement would be significantly harder to achieve as a 35 m diameter instrument would have to be mounted on a single spacecraft. This drove the design towards an architecture which employed antennas on separate free-flying satellites. Without the use of fractionated concepts, each free-flying element of the system would have to carry all the required supporting subsystems on-board, along with its primary payload asset. By fractionating those support systems, it is possible to enable the payload satellites to be smaller

and to consume less power. In addition, the fractionated nature of the system allowed for the integration of two secondary payloads to support the primary mission, without having to accommodate their mass and power impacts on the primary satellites in the system.

It has been shown how the fractionated principles can also be applied to the spacecraft bus systems. The impact of the fractionation has been assessed and some preliminary architecture decisions have been discussed. In some cases, such as the ground link communications and the on-board computing, the architecture of these fractionated subsystems is directly determined by the coastal salinity payload and its operations. The decision to fractionate other subsystems such as the data architecture and the AODS system was not directly influenced by the payload and requirements, but the exploration of fractionating these architectures was a direct consequence of the fractionated payload. The architecture of the subsystems was determined by the need to have a safe, reliable and flexible system. In this chapter, it has been shown that fractionation is a mission enabling technology, allows the coastal salinity parameter to be measured from space when a conventional, monolithic system would not be capable of supporting such measurements. The responsiveness of this system to failures was not addressed in the design exercise. In the next chapter, the knowledge gained from this design process was used to feed into a study of more generic fractionated systems, examining their responsiveness to subsystem failures whilst trying to maintain continuity of operations.

Subsystem	Free flying payload satellite (x36)	Payload hub satellite	Scatterometer satellite	Radiometer satellite	Communications satellite	OBDH satellite	AODS satellite
Mission payload/resource 1	1.0 kg	36.0 kg	260.0 kg	314.0 kg	19.5 kg	2.5 kg	5.9 kg
Mission payload/resource 2	-	11.2	-	-	-	-	-
ACS	13.4 kg	13.4 kg	26.2 kg	13.4 kg	13.4 kg	13.4 kg	13.4 kg
OCS	20.0 kg	20.0 kg	20.0 kg	20.0 kg	20.0 kg	20.0 kg	20.0 kg
Power	8.3 kg	11.2 kg	24.9 kg	26.9 kg	13.5 kg	8.6 kg	8.6 kg
Fractionated connections (10%)	4.3 kg	9.2 kg	33.1 kg	37.4 kg	6.6 kg	4.5 kg	4.8 kg
Structure and thermal control (10%)	4.7 kg	10.1 kg	36.4 kg	41.2 kg	7.3 kg	4.9 kg	5.3 kg
Margin (20%)	10.3 kg	22.2 kg	80.1 kg	90.6 kg	16.1 kg	10.8 kg	11.6 kg
TOTAL	62.0 kg	133.3 kg	480.8 kg	543.5 kg	96.3 kg	64.7 kg	69.5 kg

Table 5 Mass budgets for the satellites in the coastal salinity measurement system

Subsystem	Free flying payload satellite (x36)	Payload hub satellite	Scatterometer satellite	Radiometer satellite	Communications satellite	OBDH satellite	AODS satellite
Mission payload/resource 1	2.0 W	72.0 W	215.0 W	350.0 W	131.0 W	25.0 W	22.8 W
Mission payload/resource 2	-	34.0 W	-	-	-	-	-
ACS	11.0 W	11.0 W	68.2 W	68.2 W	11.0 W	11.0 W	11.0 W
OCS	60.0 W	60.0 W	60.0 W	60.0 W	60.0 W	60.0 W	60.0 W
Fractionated connections (10%)	7.3 W	17.7 W	34.3 W	47.8 W	20.2 W	9.6 W	9.4 W
Thermal control (10%)	8.0 W	19.5 W	37.8 W	52.6 W	22.2 W	10.6 W	10.3 W
Margin (20%)	17.7 W	42.8 W	83.1 W	115.7 W	48.9 W	23.2 W	22.7 W
TOTAL	106.0 W	257.0 W	498.3 W	694.3 W	293.3 W	139.4 W	136.2 W

Table 6 Power budgets for the satellites in the coastal salinity measurement system

3. Developing design principles for fractionated space systems

Every satellite launched in the last 60 years, has been *monolithic* in design. The system engineering principles developed over this period produce optimal monolithic spacecraft. Fractionated satellites represent a completely new way of implementing space systems. As such, existing design principles do not produce optimal fractionated systems. In the following chapters, research is described which attempted to re-define the system engineering principles required for the design and implementation of fractionated systems. In addition, a much wider range of possible fractionated systems was explored to determine how best to configure such a system.

3.1. Introduction

A direct comparison was performed between fractionated spacecraft and monolithic satellites by Dubos and Saleh, using simulations of these architectures [84]. They make similar arguments to those presented here regarding the ability of a fractionated architecture to extend the mission life beyond that of a traditional system. However in research presented here, a direct comparison with monolithic spacecraft was not made. Instead, the monolithic spacecraft was included in the search space as one possible solution. Another major difference is that Dubos and Saleh considered the failure probabilities they used to be constant throughout a spacecraft's life; assuming that the age of the spacecraft does not affect this parameter. They also only consider spacecraft failures, not subsystem failures. This had the effect of not addressing external redundancy in the fractionated subsystems, which would be a major advantage of fractionation.

In this chapter, the development of a computer model to help define a new set of design principles for fractionated systems and explore the full range of potential architectures is presented. This model was developed to answer specific questions about spacecraft fractionation. These include the key question described in section 1.2.1, as well as others. The full list is as follows:

- What is the best way to design a fractionated system, such that a continuity of measurements (or other payload service) can be maintained over a period of time that is greater than a “traditional” mission lifetime (10 to 12 years) whilst minimising the mass launched over the operational lifetime?
- At what level of overall technological maturity should a fractionated architecture be implemented?

- Which subsystems should be fractionated?
- Is redundancy required across the fractionated system, not just within individual spacecraft within the system?
 - What level of redundancy is required?
- What drawbacks are there to fractionation?
- What is the best strategy to replace failed subsystems and satellites?
- What is the best strategy to upgrade the architecture with new subsystems?

These questions were addressed with respect to a LEO Earth observing mission, but can be equally applied to all types of fractionated space systems. Some fractionated systems, such as the ones considered by DARPA can be evaluated using traditional design methods. The questions posed above cannot be answered by looking at the subsystem and spacecraft level of a space system alone. Instead, the entire *architecture* of fractionated satellites must be evaluated and analysed. This architecture level does not exist for monolithic systems, and is unique to a fractionated system. It is at this level that new design procedures and engineering principles need to be developed. A computer model has been developed which allows a wide range of potential fractionated architectures to be assessed under consistent conditions. The objective is not to make direct comparisons with a monolithic spacecraft, but to include it as a possible solution in the search space.

Maintaining a continuity of measurements over an extended mission lifetime is a key feature that is truly enabled through the use of the fractionation described in this thesis. A good example which attempts this with a monolithic system is the ERS-1, ERS-2, ENVISAT, Sentinel 1-3 systems [85]. The ERS-2 and ENVISAT spacecraft provided a continuity of measurements from the ERS-1 mission, which was launched in 1991 up until early 2012, when contact with ENVISAT was lost [86]. The Sentinel 1-3 satellites, due for launch between 2014 and 2015, were intended to continue this chain of measurements, and provide the opportunity to cross-calibrate the new instruments with the outputs from ENVISAT's instruments. With the failure of ENVISAT, this opportunity has been lost. The new Sentinel system is more robust to such a failure, by distributing the payloads onto several satellites, but these are still vulnerable to a failure as with any monolithic system.

3.2. Model structure

There are two main components to the model: an evaluation (of each architecture) and a search algorithm. The evaluation provides a consistent method of assessing an architecture's performance over a 50-year period, in terms of the time that it is operational and the total mass launched into orbit. The output from this evaluation is used to inform the search algorithm, which uses a local search method to find architectures that maximise operational time and minimise mass launched. An initial architecture is evaluated and the output of this evaluation is then passed to the search algorithm, which then finds a new architecture for evaluation. This continues until a set number of architectures are evaluated, at which point the model run terminates and the final data products are produced. The following sections describe, in detail, each component of the model, the assumptions made, how the model will be used and the outputs that are generated.

3.3. The fractionation of spacecraft subsystems

In the concept presented in section 2.5, it was assumed that only three subsystems were fractionated: the communications, OBDH and AODS. This was to minimise the technology development required for a mission that would be launched by 2016. The fractionation of other subsystems would require a major breakthrough in technology. In the aspect of the research described in this and the following sections, a different approach was taken. The maturity of the required technologies for fractionation has not been considered. Instead, the fractionation of six of the main spacecraft subsystems was assumed to be possible, allowing the required level of technological maturity to be determined.

The number of subsystems has been consolidated from six to five to simplify the model. The attitude determination system and attitude control system were combined into the attitude determination and control system. This meant that the fractionated architectures evaluated were made up of the following five subsystems: communications, on-board data handling, attitude determination and control, orbital control and the power subsystem. Each spacecraft in the simulated architecture comprised between one and five of the fractionatable subsystems and possibly a payload. As with the coastal salinity system design, it was assumed that each spacecraft provides its own thermal control and a structure. In the model subsystems which are designated as fractionated are assumed to have the capability to share their resources with the rest of the architecture.

There are two universal rules that apply to all fractionated architectures regardless of operational strategy or design used, as follows:

Rule 1: *Each of the fractionated subsystems must be present at least once in the architecture*

Rule 2: *All of the non-fractionated subsystems must be present on every satellite in the architecture*

The implication of Rule 2 is that if a non-fractionated subsystem fails, its host satellite becomes inoperative, and any fractionated subsystems on-board will no-longer be able to share their resources with the rest of the architecture. These two rules were applied to all the architectures evaluated. It was assumed that for the fractionated architectures considered here, the use of “small” satellites and subsystems is the most appropriate, as these systems have developed using commercial-off-the-shelf components, to reduce the time it takes to design, build, test and launch spacecraft [3]. As such, the subsystems considered were modelled as components designed and manufactured by Surrey Satellite Technology Limited (SSTL). Each subsystem is described in the following sections. The payload was assumed to be an SSTL imaging system made up of high-and medium-resolution visible wavelength imagers [87]. The payload was always assumed to be fractionated, so that only a minimum of one was required in the architecture.

The following sections describe how these subsystems are implemented in the model, whether fractionated or not, and the associated assumptions made. Where the fractionation of a subsystem has not been discussed in section 2.5.1.2, this is discussed here as well.

3.3.1. Communications

The communications subsystem was assumed to include the transponders, any necessary power amplifiers and the antennas. If the communications subsystem is fractionated, then the dissemination of telecommands throughout the fractionated system and the collection of payload and telemetry data are handled by the OBDH network. It was assumed that inter-satellite communication takes the form of a radio-frequency system. In an architecture with a fractionated communications subsystem, a minimum of one satellite in the architecture must carry a communications subsystem to provide all the satellites with a link to the controllers on the ground. An architecture with a non-fractionated communications system requires every satellite in the architecture to host a communications subsystem and each spacecraft receives commands and sends telemetry back to the ground independently.

The communications hardware took the form of S- and X-band units, for telemetry and payload data respectively. These were assumed to be units designed and developed by SSTL. Table 7 gives the mass and power values for the communications equipment.

Unit	Quantity	Mass (kg)	Power (W)	Reference
SSTL S-Band Downlink Transmitter	2	2 kg	38 W	[66]
SSTL S-Band Transmit Filter	2	0.09 kg	-	[67]
SSTL S-Band Uplink Receiver	2	1.3 kg	1.5 W	[68]
SSTL S-Band Receive Filter	2	0.09 kg	-	[69]
SSTL S-Band Patch Antenna	2	0.08 kg	-	[70]
SSTL S-Band High Power Amplifier	1	1.25 kg	33 W	[71]
SSTL X-Band Downlink Transmitter	2	3.25 kg	55 W	[72]
SSTL X-Band Antenna Pointing Mechanism with High Gain Antenna	2	2.7 kg	3.5 W	[73]

Table 7 Mass and power properties for the communications equipment

3.3.2. On-board data handling

The OBDH subsystem was assumed to include the computers, data storage and any monitoring equipment for FDIR operations. A fractionated OBDH subsystem also accounts for the hardware required to establish and maintain the data network between the satellites in the architecture. It is assumed that each satellite in the architecture has some low-level of computational capability in order to operate, but the main processing tasks are handled by the main OBDH subsystem(s) present in the architecture. An architecture with a non-fractionated OBDH subsystem requires every satellite in the architecture to host an OBDH subsystem. These subsystems are then responsible for the computation, data storage and FDIR functions for their host satellite. The OBDH hardware consisted of a SSTL OBC750 computer (1.5 kg, 10 W) and a SSTL High Speed Data Recorder (1 kg, 15 W) [74], [75].

3.3.3. Power

The fractionation of the power system requires electrical power to be transferred wirelessly between those satellites that have a primary power source (such as solar arrays) to those that do not. The technology required to do this has been extensively researched for the purpose of providing power generated in space to the ground for terrestrial uses [55]. Demonstrations of power transfer at microwave frequencies have been demonstrated over small distances on the ground [88]. There are considerable challenges to wireless power transfer, not least that of the achievable efficiency. A study indicated that wireless power transfer between spacecraft would not be suitable to provide the primary power source for a system of free flying spacecraft [17]. In their design study, Orbital Sciences did not consider the power system for fractionation and instead included this functionality as a technology demonstration only [8].

The power system consists of solar arrays for power generation, a power conditioning and distribution unit (PCDU) and batteries. An architecture with a fractionated power system requires that at least one satellite in the architecture carry a power system. A key assumption for the power system was the location for the batteries and PCDUs in the architecture. As power is such a key resource for the architecture, relying on another spacecraft to provide power leaves a single point of failure in the transmission. Therefore it was assumed that every satellite in an architecture with a fractionated power system has a battery and PCDU to maintain the satellite's functionality in the event of a loss of transmitted power. In reality, this battery will only provide power to the host satellite for a few orbits. It is therefore assumed that the depletion of the battery is a failure that is recoverable. For architectures with a non-fractionated power system, the solar arrays, PCDU and batteries are hosted on every satellite in the architecture [60],[64], [62]. The sizing of the power systems is described below in section 3.3.6.

3.3.4. Attitude determination and control system

The design presented in section 2.5 only considered attitude determination as a fractionatable subsystem. In the model, the full attitude determination and control system (ADCS) is considered as a single subsystem for fractionation. The required technologies for the fractionation of the attitude determination subsystem were presented in section 2.5.1.2.3. Here, we discuss the required technologies for the fractionation of the attitude control subsystem followed by the implementation of the ADCS in the model.

The fractionation of the ACS requires that one (or more) spacecraft provide forces and torques to control the attitude of the other satellites in the architecture. Theoretical studies have been conducted, such as those into Coulomb formation flying (CFF), where electrostatic forces are used

to change the relative orientation and position of satellites [89]. The CFF concept only affects the *relative* position and orientation of the satellites. The system will still be affected by atmospheric drag and other orbital perturbations. To date, there have been no known demonstrations on the ground or in orbit of this attitude control technique. The fractionation of the ACS using the CFF technique will require a major investment.

The ADCS subsystem includes the attitude sensors: sun-sensors and star trackers, and attitude actuators: reaction wheels and magnetorquers. An architecture with a fractionated ADCS system requires that at least one satellite in the architecture host an ADCS subsystem. This subsystem then provides attitude determination and control capabilities for the entire architecture. An architecture with a non-fractionated ADCS system requires every satellite in that architecture to carry an ADCS subsystem, to provide the host satellite with its own attitude determination and control [57], [58], [79], [81]–[83].

The attitude control actuators consisted of four reaction wheels (assumed to be SSTL Microwheel 100SP-M wheels for the smaller satellites and SSTL Smallwheel 200SP wheels for the satellites hosting the two larger secondary payloads) and three magnetorquers for momentum dumping (assumed to be SSTL MTR30 magnetorquers) [56]–[58]. Table 8 gives the mass and power consumption for these components.

Unit	Mass (kg)	Power consumption (W)	Reference
SSTL Microwheel 100SP-M	2 kg	2 W	[56]
SSTL Smallwheel 200SP	5.2 kg	16.3 W	[57]
SSTL MTR30	1.8 kg	1 W	[58]

Table 8 Mass and power properties for the ACS components

Table 9 gives the mass and power values for the AODS equipment used.

Unit	Quantity	Mass (kg)	Power (W)	Reference
SSTL SGR-10 GPS Receiver	1	0.95 kg	5.5 W	[79]
SSTL MIRAS-1 Inertial Sensor	1	1.8 kg	5 W	[80]
SSTL 2-Axis Sun Sensor	2	0.3 kg	-	[81]
SSTL Magnetometer	1	0.14 kg	0.3 W	[82]
SSTL RIGEL-L Star Tracker	2	2.4 kg	12 W	[83]

Table 9 Mass and power values for the AODS equipment

3.3.5. Orbital control system

The fractionated system considered in section 2.5 did not consider the fractionation of the orbital control subsystem. Here, the required technologies and the implementation in the model are discussed. Of all the subsystems, the fractionation of the OCS is the most difficult as it requires one satellite to precisely change the orbit of another satellite. One possible solution is the Ion Beam Shepherd (IBS) concept, proposed for active debris removal purposes [90]. This utilises two electric propulsion (EP) systems on-board a single spacecraft. One EP system pushes against the target (spacecraft or debris) to adjust its orbit, while the second is used to maintain the orbit of the host satellite. In spite of EP being a widely used form of propulsion, the IBS concept is an immature one. There are many unknowns, including whether the ion beam would ablate the spacecraft, an obviously undesirable event! Despite this limitation, the OCS is still considered a fractionatable subsystem in order to assess if it is required before a fractionated architecture is successfully implemented. A non-fractionated OCS system takes the form of a SSTL Xenon propulsion system, which had a wet-mass of 20 kg and consumed 60 W of power [59]. The mass assumed for both a fractionated and non-fractionated OCS is the mass of the cold gas propulsion system.

3.3.6. Mass and power values used in the model

The mass and power values for each of the subsystems were modelled using information found in SSTL data sheets. Table 10 gives the mass and power values used for the payload, communications, OBDH, ADCS and OCS subsystems [56], [58], [59], [66]–[75], [79], [81]–[83], [87].

While these masses and power values may not be wholly accurate with respect to those that might be used in a future fractionated system, they were chosen as representative values for the modelling and simulation work conducted, as it allowed the development of the model to progress using real world subsystem masses.

Subsystem	Mass	Power
Payload	41.00 kg	55.0 W
Communications	20.27 kg	131.0 W
OBDH	2.50 kg	25.0 W
ADCS	45.04 kg	99.1 W
OCS	20.00 kg	60.0 W

Table 10 Mass and power values used in the model, derived from SSTL data sheets

The power subsystem was sized based upon the power requirements of the architecture. A non-fractionated power system was sized to provide power to the host satellite only. This power system consisted of a PCDU, a battery and solar arrays. The PCDU had a mass of 3.9 kg [60]. The battery which was sized to provide power to the bus subsystems in eclipse (i.e. they payload is not operated in eclipse). An 800 km circular orbit was assumed, which gave an eclipse time of 35.13 minutes [61]. The battery sizing process allowed one of five different battery sizes to be chosen. The masses of these batteries, taken from an ABSL Power Solutions Ltd. data sheet were 1.2 kg, 3.3 kg, 4.7 kg, and 15.9 kg, which corresponded to battery capacities of 130 W-h, 346 W-h, 486 W-h, 691 W-h and 1620 W-h [62]. The solar array area, A_{SA} was given by,

$$A_{SA} = \frac{P}{F_{sun} \times \eta} \quad [13],$$

where P is the power required by the spacecraft when in sunlight plus the power required to charge the battery (assuming a 30% depth of discharge). F_{sun} is the solar flux at the Earth (assumed to be 1350 W/m^2). η is the efficiency of array system, assumed to be 0.2 and encapsulates the efficiency of the solar cells, the packing efficiency of the cells and the degradation factor [63]. The mass of the solar array was given by multiplying the array area by 1.76 kg/m^2 , the specific mass of a Spectrolab panel using triple junction cells [64].

A fractionated power system was sized to firstly provide power to the host satellite, and then some surplus power to those satellites without the capability to generate power, calculated by simple division of the power requirements of these satellites. As described in section 3.3.3 each

satellite in the architecture was assumed to carry a battery and PCDU, which had the same mass as for the non-fractionated case.

The total mass and power consumption of each satellite in the architecture was given by the sum of masses and power consumption of the subsystems on-board each spacecraft. An additional 20% was added to the mass and power consumption of each satellite to account for the structure mass, the mass and power required for thermal control and the mass and power required to form the inter-satellite links for sharing the fractionated resources.

3.4. Architecture evaluation

The first main component of the model was the evaluation of fractionated architectures. This simulated a fractionated architecture's operational life by generating random failure events and evaluating how it responded to those failures. In this context, failure was defined as any malfunctions that caused the subsystem to cease functioning permanently. Temporary outages (such as a software crash) were ignored. This evaluation of fractionated architectures only addressed the responsiveness of the architectures to two sources of uncertainty: Technical uncertainty and Environmental uncertainty. The evaluation used can be adapted to understand the responsiveness to Launch, Demand and Requirements uncertainty. In order to keep the model simple and the results understandable, the decision was taken to not include these three sources of uncertainty. The responsiveness to Funding uncertainty was not included in this evaluation as it cannot be modelled using technical parameters, and would require a detailed cost model to be developed. The purpose of this evaluation was to provide an environment in which the performance of fractionated architecture can be analysed under consistent conditions. The use of random failures brought a stochastic element to the model which was entirely deterministic in all other respects. A single run of the evaluation algorithm was not sufficient to accurately represent the performance of an architecture. Instead, a Monte-Carlo (MC) method was required, where the evaluation was performed multiple times for the same architecture, and the results are then averaged. The following sub-sections will describe each component of the architecture evaluation, and how they fit in with the rest of the model.

3.4.1. Fractionated architectures

Fractionated architectures can be represented in a matrix format as shown in Table 11. Each row represents a satellite in the architecture, and each column a subsystem. The columns highlighted in blue represent subsystems that are fractionated. A cell with a “1” represents the presence of a subsystem on the satellite in the corresponding row, while a “0” represents that the subsystem is

not present on the satellite in the corresponding row. For example, the fractionated ADCS subsystem is present on satellite 1, but not on satellites 2 and 3.

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	1	1	1	1	1
2	0	1	1	0	1	0
3	0	1	1	0	1	1

Table 11 Example matrix format of a fractionated architecture with the payload, OBDH, ADCS and power subsystems fractionated

The two rules described in section 3.3 were used to derive the architecture shown in Table 11. Each fractionated subsystem must be present in at least once in the architecture, therefore each blue column must sum to greater than- or equal-to one. All the non-fractionated architectures must be present on every satellite in the architecture, therefore the sum of each white column must be equal to the number of satellites in the architecture.

Five parameters were derived from the matrix format shown in Table 11. The *number of satellites*, in the architecture, $N_{Satellites}$, was given by the number of rows in the table. The average *number of subsystems per satellite*, $N_{Subsystems}$, was given by,

$$N_{Subsystems} = \frac{\sum_{i,j=1}^{i=6, j=N_{Satellites}} M_{arch_{ij}}}{N_{Satellites}} \quad 14,$$

where, M_{arch} is the architecture matrix, i is the column number and j is the row number. For the example shown in Table 11, $N_{Subsystems} = 4.33$. Three parameters were used to describe the fractionated properties of the architecture. The *degree of fractionation*, D , described the number of subsystems which are capable of being shared within the architecture and was given by the number of blue columns in Table 11, in this case $D = 4$. The *fractionated subsystems per satellite*, F , given by,

$$F = \frac{\sum_{i,j=1}^{i=D, j=N_{Satellites}} M_{Fractionated_{ij}}}{N_{Satellites}} \quad 15,$$

where $M_{Fractionated}$ is the matrix formed by the fractionated columns only, for the above example, $F = 2.33$. This described the average number of fractionated subsystems that are housed on each satellite. The *external redundancy*, E , given by,

$$E = \frac{N_{Satellites} \times F}{D} \quad 16.$$

This describes the redundancy in the fractionated subsystems. For the example shown in Table 11 $E = 1.75$.

The number of satellites in an architecture is limited to less than or equal to six. The upper limit of six satellites in the architecture was chosen to allow a heterogeneous architecture to exist, as shown in Table 12. In this architecture $D = 6$, $N_{Subsystems} = 1$ and $E = 1$. If there were fewer than six satellites, then there would have to be more than one subsystem hosted on each satellite. If there were greater than six satellites, a $D = 6$, $N_{Subsystems} = 1$ architecture would have $E > 1$, which is not a heterogeneous architecture.

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	0	0	0	0	0
2	0	1	0	0	0	0
3	0	0	1	0	0	0
4	0	0	0	1	0	0
5	0	0	0	0	1	0
6	0	0	0	0	0	1

Table 12 Heterogeneous fractionated architecture

3.4.2. Architecture simulation

Each architecture had its operational life simulated to assess its performance over the 50-year period. The 50-year projection period was divided into discrete time-steps. At each time-step, failures can occur, and new satellites can be launched. These processes were simulated using a failure model and a replacement model, which was the process which described how the architecture responded to the failure of a subsystem.

3.4.2.1. Failure model

The failure model was one of the key components of the overall model, as it was the only stochastic element in the simulation. At each time-step in the projection period, every subsystem present on each satellite was assessed to determine if it has failed. Failures were determined by the use of failure curves which give the probability of failure as a function of time.

The initial data exploration used a single failure curve applied to all the subsystems equally. This curve has been designed such that the probability of a subsystem failing over 12 years of time is 100%. It was an approximation, and not based on satellite failure data. The assumption was made that the design life of the subsystems was seven years and, as such, the curve was designed so

that the average age of a subsystem at failure was at seven years. This was approximately the design lifetime of many LEO earth observation spacecraft [61].

There were three components of the failure rate function: early wear-out, random failure and end of life wear-out. These three components were combined to give the final failure rate function. The values used have been selected to be indicative of a failure curve made up from the three components and dominated by wear out [91]. The early wear-out component represented failures at the start of the subsystem's life, such as the failure of a mechanism, or a power switch to turn on the subsystem. The early wear-out failure probability rate is,

$$R_e = \frac{1}{(a+1.45)^{10}} \quad 17,$$

where a is the age of the subsystem in years. The random failures component represented events such as a micro-meteoroid, debris impact or a solar flare event, and was described by a constant failure rate. The random failure probability rate is,

$$R_r = 0.0026625 \quad 18.$$

The third component was the end of life wear-out. This represented the depletion of on-board consumables and the natural mechanical or electrical wear-out of all components. The end of life wear-out rate is,

$$R_{EOL} = \frac{a^{11.5}}{a_{max}^{11.5}} \quad 19.$$

The value of a_{max} has been chosen to be 12 years such that probability of failure is 1 after this period of time. The three components were combined to give the overall failure probability rate function over the 12-year lifetime of the subsystems,

$$R = R_{EOL} + R_e + R_r \quad 20.$$

The integral of this function, R gives a probability of failure is 1 after 12 years. The curve described by this function is shown in Figure 27.

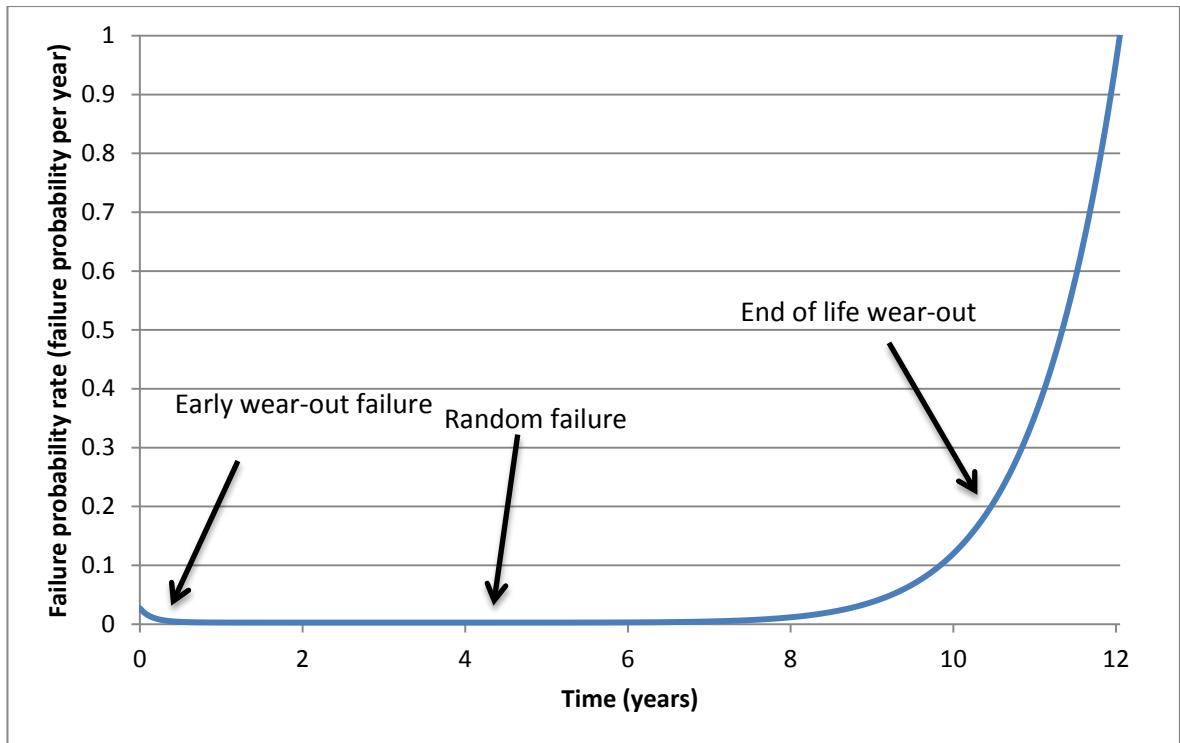


Figure 27 Failure rate curve for scenario 1. On the x-axis is the age of the subsystem and on the y-axis is the failure rate in failures per year.

The probability of failure for each subsystem during each time-step is calculated by a rectangle to approximate the area under the failure rate curve during time-step, T . This is calculated by the value on the lifetime curve corresponding to the age of the subsystem at time t ($R(a)$) multiplied by the time-step size, T ,

$$P_{failure}(t) = R(a) \times T \quad 21.$$

For the subsequent analyses, a new type of failure curve was implemented, based upon the work of Jean-Francois Castet and Joseph Saleh of Georgia Institute of Technology who studied the failures of over 1500 spacecraft to produce reliability curves [92]. Using a Maximum Likelihood Estimation (MLE) method, they modelled the spacecraft failures as 2-mixed Weibull functions. This treats the reliability of a satellite as a probability density function. The 2-mixed Weibull function comprises two components, the infant mortality and the end of life wear out. The probability density at time, t , is

$$P_D(t) = \alpha_1 \left(\frac{\beta_1 t^{\beta_1-1}}{\theta_1^{\beta_1}} \right) \exp \left[-\left(\frac{t}{\theta_1} \right)^{\beta_1} \right] + \alpha_2 \left(\frac{\beta_2 t^{\beta_2-1}}{\theta_2^{\beta_2}} \right) \exp \left[-\left(\frac{t}{\theta_2} \right)^{\beta_2} \right] \quad 22,$$

where β is the shape parameter, θ is the scale parameter and α is the mixture parameter where, $\alpha_1 + \alpha_2 = 1$. The subscript 1 refers to the infant mortality component and the subscript 2 refers to

the end of life wear out component. Table 13 lists the parameters used for the reliability curve in this scenario [92].

Infant Mortality	α_1	0.9484
	β_1	0.2575
	θ_1	982100
End of Life Wear Out	α_2	0.0516
	β_2	1.997
	θ_2	10.2

Table 13 2-mixed Weibull function parameters

These parameters were used along with Equation 23 to produce the graph shown in Figure 28, which describes the probability density function for the spacecraft reliability. Using this method, there was no predetermined age at which the spacecraft is guaranteed to fail, unlike scenario 1 where the failure curve was designed to guarantee a failure after 12 years of operations. Therefore, this was a major assumption for using this failure curve as it allowed the spacecraft to be operational beyond the 15-year nominal lifetime implied by the data.

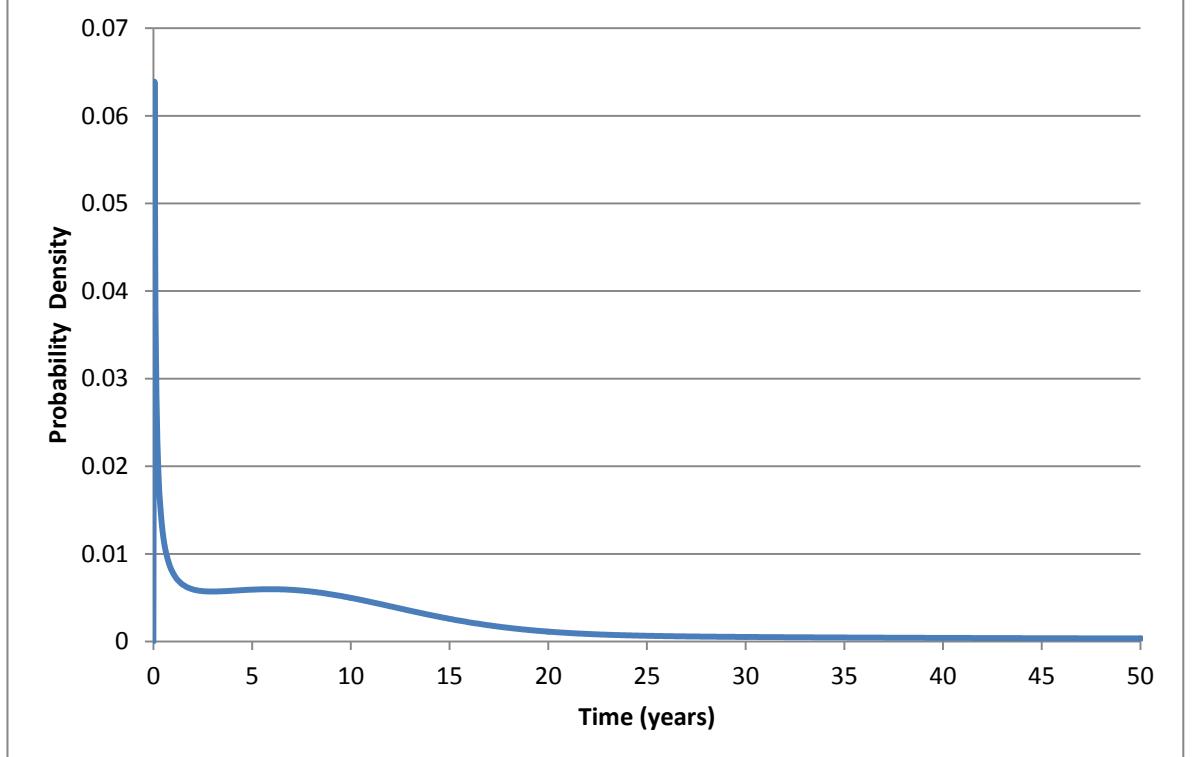


Figure 28 2-Mixed Weibull probability density function

A trapezium rule was used to calculate the probability of a subsystem failing in the period time t to $t + T$,

$$P_{failure}(t) = (a(t + T) - a(t)) \left[\frac{P_D(t) + P_D(t+T)}{2} \right] \quad 23.$$

Where $a(t)$ is the age of the subsystem at time t , and T is the time-step size. A random number, N_{rand} , was then generated, and if $N_{rand} < P_{failure}(t)$ then the subsystem was determined to have failed.

Regardless of the failure curve used, when a failure occurred, the effect on the architecture depended on whether the failed subsystem was fractionated or not. If the failed subsystem was fractionated, then that shared resource would no-longer available to the rest of the architecture and the host satellite was otherwise unaffected. Rule 2, given in section 3.3 states that all the non-fractionated subsystems must be present on each satellite in the architecture in order for them to be functioning. Therefore, if a non-fractionated subsystem failed, it would result in the failure of its host satellite and any fractionated subsystems on-board would no longer be available to the rest of the architecture. Therefore, a failure in a non-fractionated subsystem could have a more severe impact on the architecture than a failure in a fractionated subsystem. It had the potential to remove more shared resources than just the subsystem that has failed.

3.4.2.2. Replacement of failed subsystems

The other key process in the model was the response of the architecture to the failures it experiences. After a failure occurred, a strict set of deterministic rules were followed to establish whether a new satellite was required and, if so, which subsystems should be present on it. A replacement period of 0.7 years (8 months 12 days) was used to represent a modular assembly and integration of the replacement satellites, using COTS components [3]. Figure 29 shows the flow diagram that describes the replacement process. This process followed a “Fail-then-replace” strategy, whereby new satellites were only launched in response to a failure. Two other strategies were considered, the first was an on-orbit spares strategy similar to the one used for the navigation systems, such as the Navstar GPS [93]. The second replacement strategy required the prediction of an up-coming failure and the launch of replacement satellite before the failure can occur. These strategies were not used as they were more complex to implement than the fail-then-replace strategy. It was felt that the use of a simpler strategy would make the outputs of the model easier to understand.

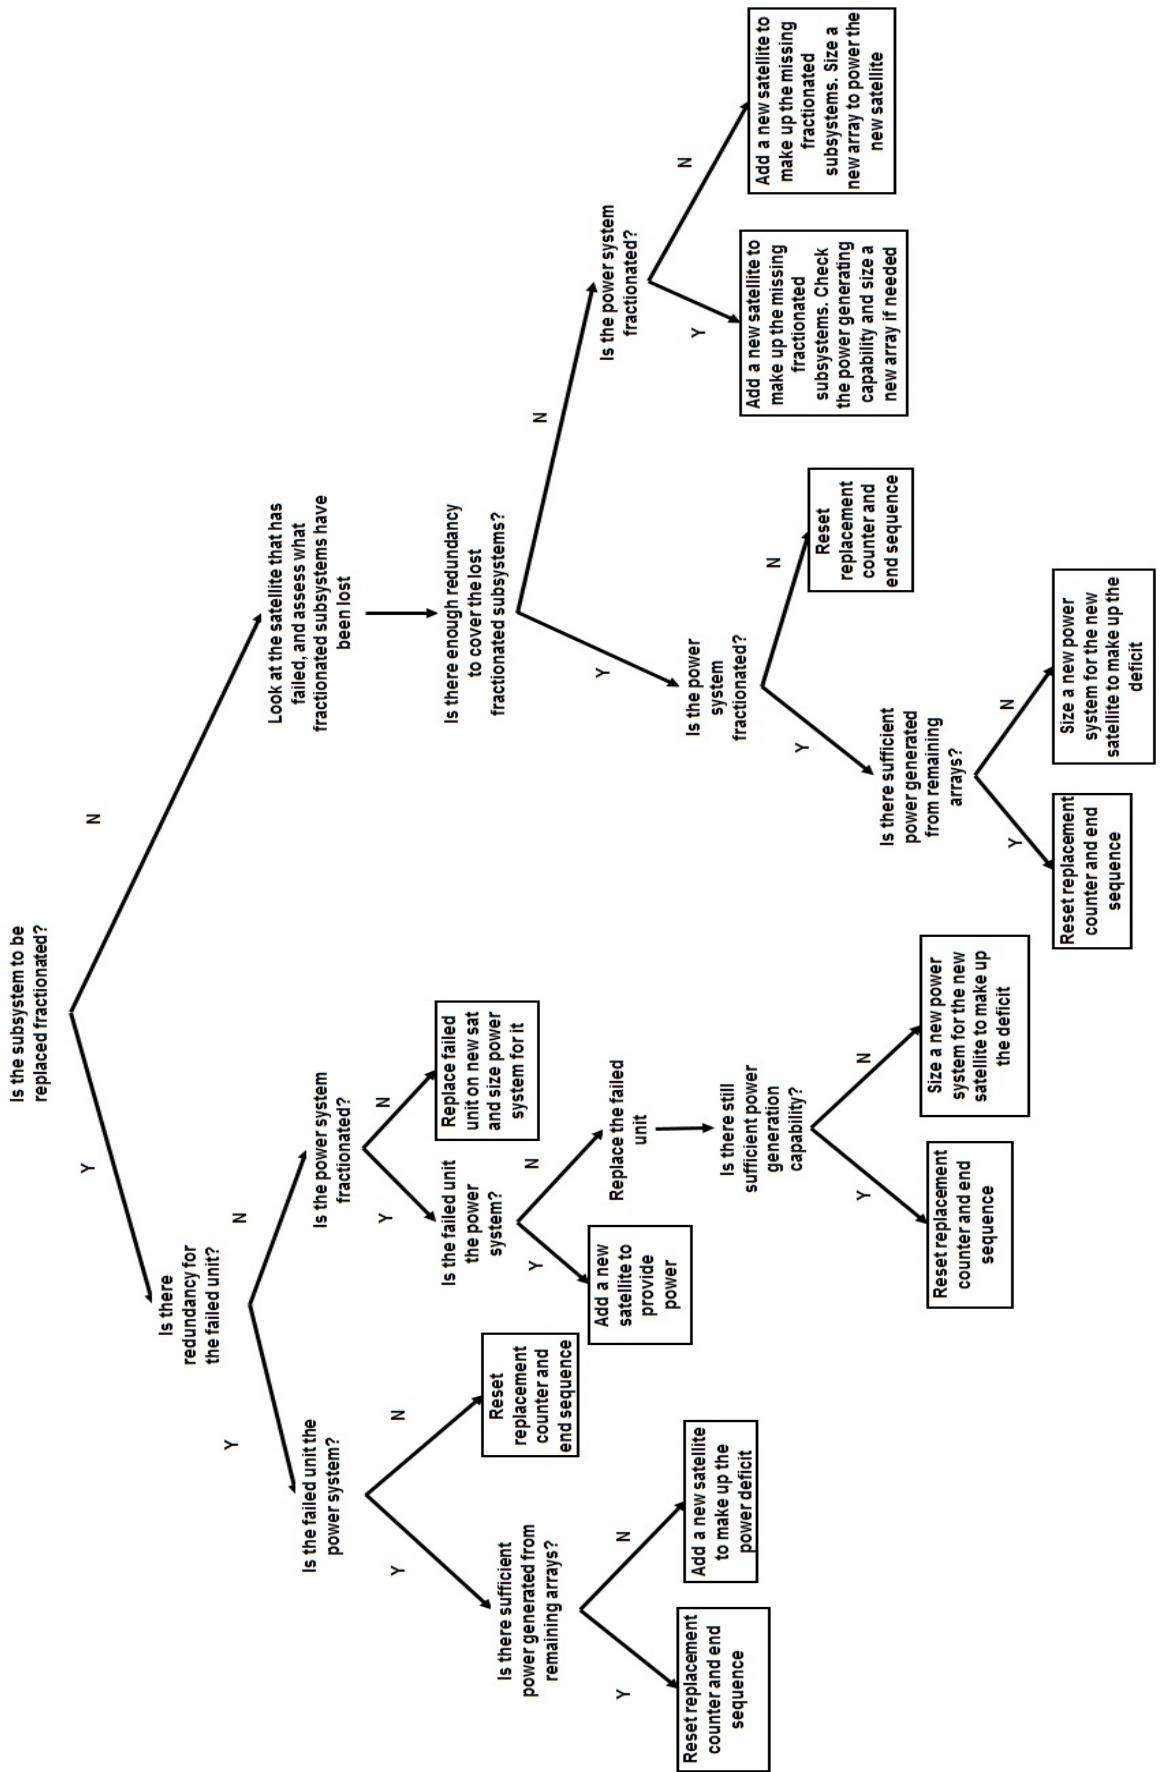


Figure 29 Replacement flow diagram

The process shown in Figure 29 ensured that when a new satellite was launched it restored the architecture back to full operational capacity. Each new satellite launched must contain all the fractionated subsystems which are not present in the architecture, to ensure that the architecture was fully functional. This included the payload, which was treated as fractionated in this model. In addition, each new satellite required all non-fractionated subsystem to be present. When a new fractionated power system was launched it was sized to provide enough power to make up the deficit left by the failure using the equations given in 3.3.6.

3.4.3. Outputs from the architecture evaluation

Failures were characterised using four measures: the *number of failures* (n_f) that have occurred during the architecture's 50-year life, the average age of the subsystem when it fails (*average age at failure*, a_f), the *number of satellites launched* over the satellite's lifetime (n_{SL}), and the *number of failures which require a new satellite to be launched* (n_{fs}). This final measure was an important one, because it was only the failures which cause the system to lose operational capability that required a new satellite to be launched. Therefore, the best architectures will have the greatest difference between the n_f and the n_{fs} . In practical terms, it means that the architecture must be robust to failures. This results in operational down-time being minimised, and helps to minimise the mass launched. The architecture properties that lead to maximum robustness were determined from the results of the model. Two outputs were produced by the model for each architecture. The first was the proportion of the architecture's 50-year lifetime that it was operational, (*operational time*, t_{op}). This was evaluated at each time-step throughout the projection period. The architecture was considered to be operational if each of the fractionated subsystems were present on at least one satellite, all the non-fractionated subsystems were present on every satellite and the payload was present on at least one satellite. The second output was the total mass that has been launched over the architecture's lifetime (*total launched mass*, M). The operational time and the total launched mass are combined into a single fitness parameter,

$$Fitness = \frac{t_{op}}{M} \quad 24.$$

The operational time represented the parameter to maximise, whereas the total launched mass was the parameter that was minimised, representing a constraint on the architecture. Both were normalised, using a zero-mean unit-variance approach,

$$N_{F_k} = \frac{k - \bar{k}}{Var(k)} \quad 25.$$

Where N_{Fk} is the normalised value of parameter k , \bar{k} is the mean of k and $Var(k)$ is the variance of k . This approach re-scales the parameters to values approximately between +3 and -3 with a normal distribution. The mean of this distribution is centred around zero, therefore when the two parameters are divided the range of possible fitness values is very large. To rectify that issue, the normalised values have 3 added to them to make the mean approximately +3,

$$\text{Fitness (normalised)} = \frac{N_{F_{top}} + 3}{N_{F_M} + 3} \quad 26.$$

This has the result of making the fitness approximately +1. Figure 30 shows the fitness values observed in a typical run of the model and the progression of the search algorithm, with its tendency to produce higher scoring architectures as it progresses.

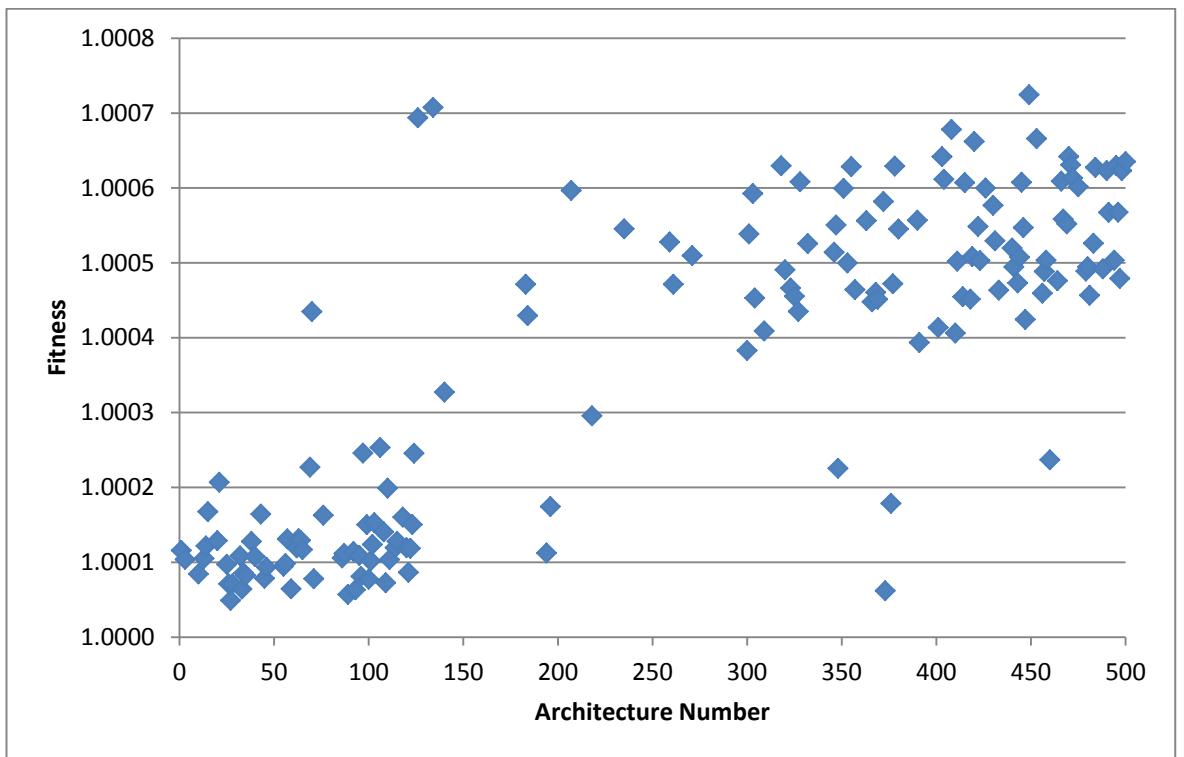


Figure 30 Graph showing the fitness values for a typical run of the model

This fitness parameter is not a standard spacecraft design parameter. However, it aims to encapsulate two important aspects of any spacecraft: the operational time, a measure of how much utility is achieved by the system and the mass launched, a proxy for the cost to the operator of implementing the system.

3.4.4. Analysis on the number of Monte-Carlo runs required

The number of Monte-Carlo simulations performed on each architecture determined how accurately the probability distributions in the failure curves were reconstructed. Increasing the number of Monte-Carlo simulations had the effect of increasing the run-time. The model was run for different numbers of Monte-Carlo runs and the fitness, standard deviation of the fitness and run-time were recorded to find the optimal accuracy and run time. Figure 31 to Figure 34 show the fitness values for the viable architectures found in each model run plotted against the architecture number on the x-axis, with one standard deviation error bars. Table 14 shows the run-time and mean standard deviation for each run. There was a 42% reduction in the mean standard deviation between the 50 Monte-Carlo simulations run and the 300 Monte-Carlo simulations run. Diminishing returns were achieved for the 1000 Monte-Carlo simulations case. As a result, the decision was taken to run the model with 300 Monte Carlo runs on each architecture, as no real benefit could be seen by performing more, at a cost of a significant increase in run time.

Number of Monte-Carlo simulations per architecture	Run-time	Mean standard deviation
50	1 hour 31 minutes	0.015952107
100	4 hours 22 minutes	0.012182525
300	8 hours 33 minutes	0.009290697
1000	26 hours 31 minutes	0.009216548

Table 14 Number of Monte-Carlo simulations, run-time and mean standard deviation

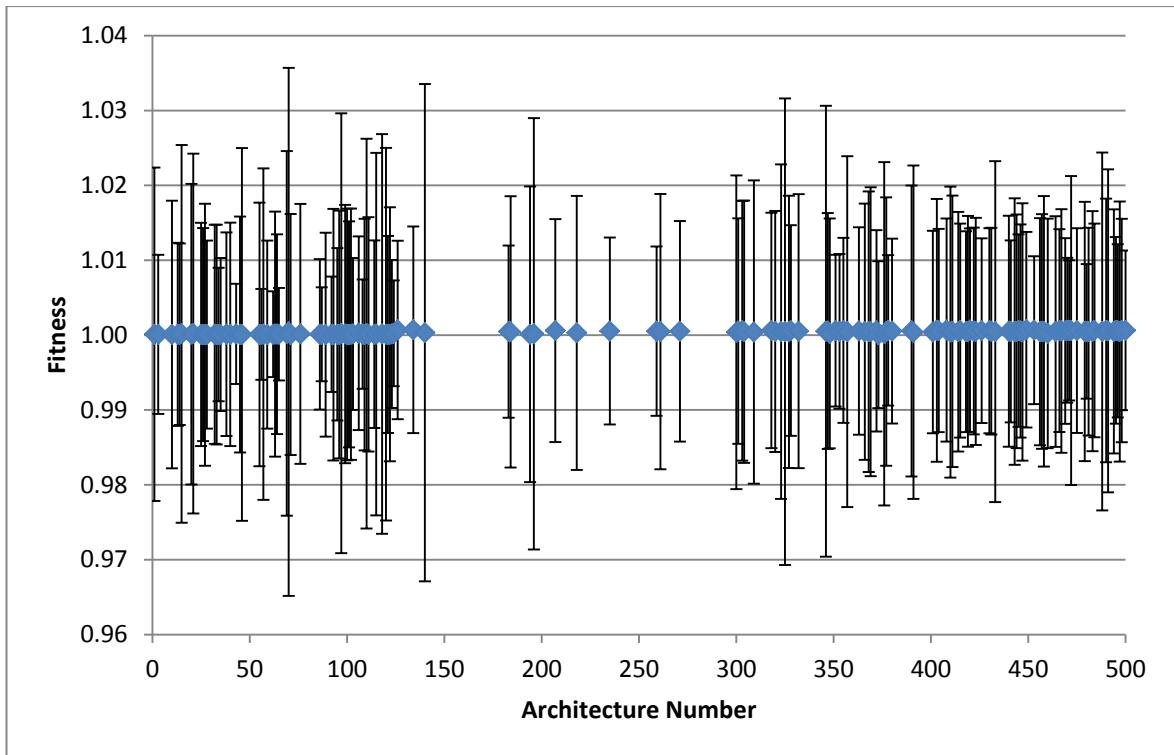


Figure 31 Model run with 50 Monte-Carlo simulations per architecture with 1 standard deviation error
bars

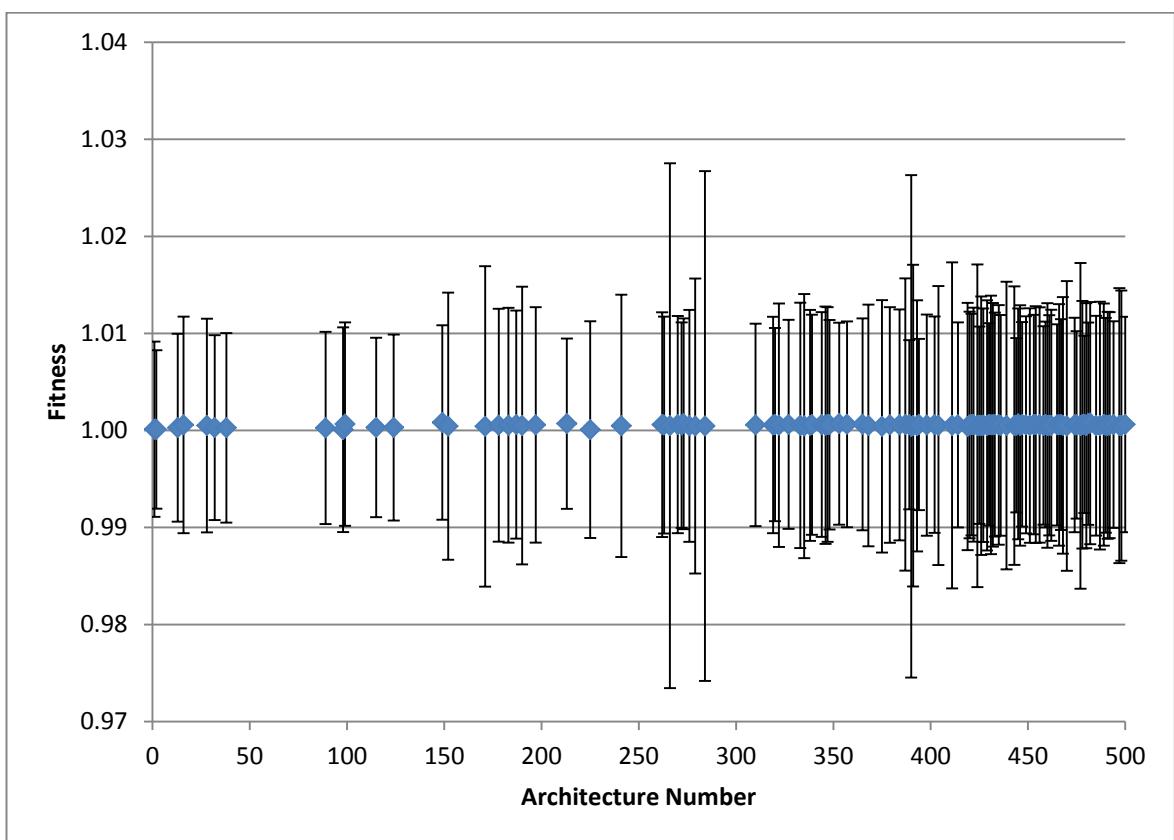


Figure 32 Model run with 100 Monte-Carlo simulations per architecture with 1 standard deviation error
bars

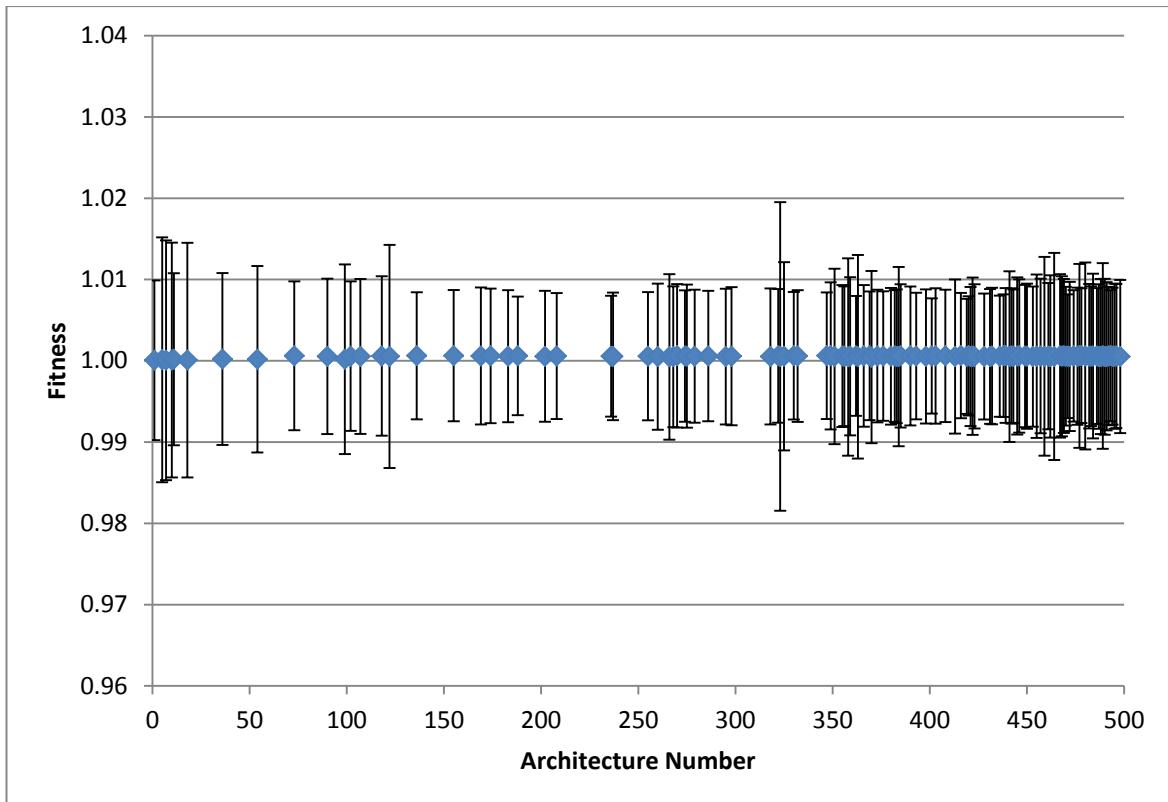


Figure 33 Model run with 300 Monte-Carlo simulations per architecture with 1 standard deviation error bars

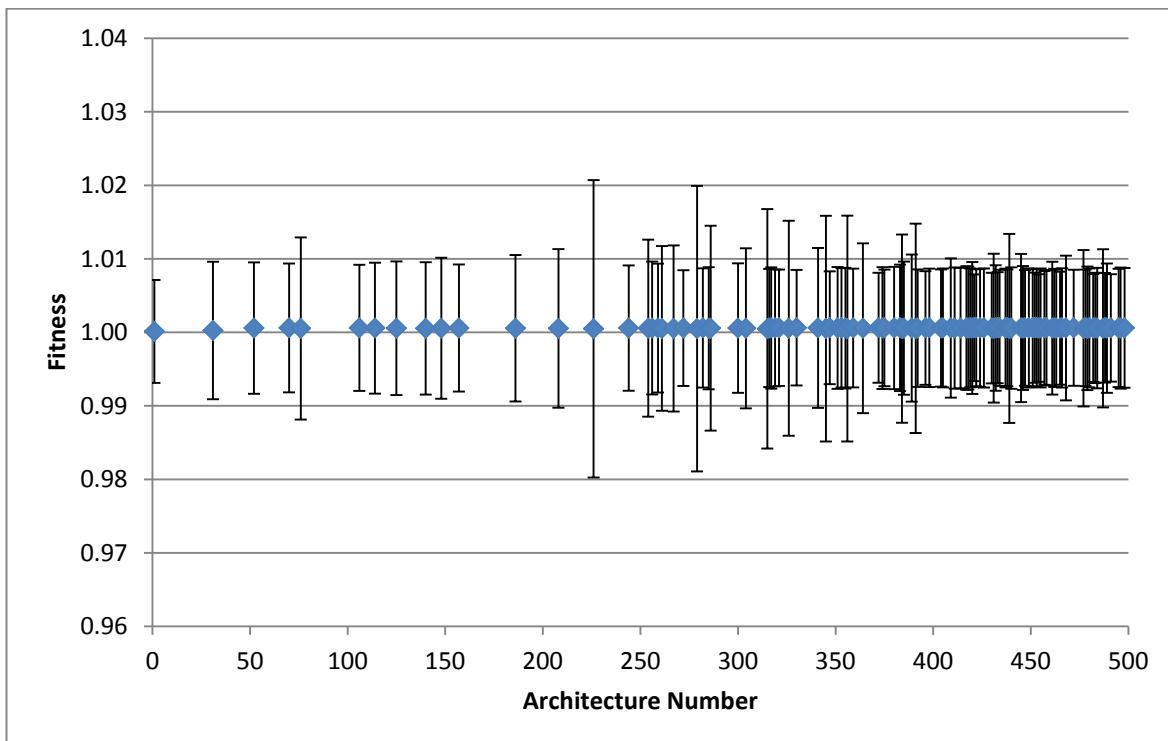


Figure 34 Model run with 1000 Monte-Carlo simulations per architecture with 1 standard deviation error bars

3.5. Search algorithm

With six subsystems, the number of possible combinations of subsystems is 4.95×10^{13} (given by 2^{36} combinations of satellites and subsystems multiplied by 720 possible permutations of fractionated subsystems). While not all of these will be fully functional fractionated architectures, to evaluate every one and perform a Monte-Carlo simulation on every viable architecture would take decades of computing time. Therefore, a search algorithm was required to find the optimal architectures. Two options were explored for the search algorithm: a simple, local search algorithm and a genetic algorithm.

Genetic algorithms (GAs) are inspired by the biological process of natural selection and genetics. It is one example of a set of optimisation techniques known as evolutionary algorithms. GAs differ from many other optimisation techniques in that they operate on populations of potential solutions simultaneously rather than on a single solution at a time [94], [95]. With a large population of potential fractionated architectures to explore, a GA was good candidate search algorithm for use in the model because it acts on populations of potential solutions. The initial population would be formed of a randomly determined subset of the possible fractionated architectures. This population would be assessed using the evaluation algorithm described in section 3.4. The result of this assessment would be a population of fractionated architectures with a fitness score associated with each one. The next step in the GA process is to mix the population of evaluated solutions to produce a new population set [94], [95]. In this study, the architectures would be ranked in order of fitness and the top half would be passed on to this mixing process. The mixing process consists of two components: crossover and mutation. Crossover involves pairs of evaluated solutions exchanging information that describes their composition [94], [95]. This requires the solutions to be encoded in a form that allows the crossover to be performed by the computer model. In this research this encoding takes the form of the matrix notation for a fractionated architecture as shown in Table 15. Crossover would take the form of two architectures exchanging a sub-matrix of these arrays.

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	1	1	1	1	1
2	0	1	1	0	1	0
3	0	1	1	0	1	1

Table 15 Example matrix format of a fractionated architecture with the Payload, OBDH, ADCS and Power subsystems fractionated

The second component of the mixing process is mutation. This involves random changes being made to individual solutions in the population. In this research, this would take the form of increasing or decreasing the number of satellites (number of rows in Table 15), degree of fractionation (number of blue shaded columns in Table 15) and the locations of the subsystems on each satellite (changing values from one to zero or *vice-versa* in individual cells in Table 15). The new population of potential solutions would then be subjected to the evaluation algorithm and the process would be repeated for a set number of cycles (generations) or until a fitness threshold was reached. The GA technique was rejected for use in this research because it was found that the matrix notation encoding was not suitable for the crossover technique used, as it did not produce, on average, better fractionated architectures than the previous generation. Secondly, each fractionated architecture in each generation required 300 Monte-Carlo simulations to produce a fitness score. Therefore any GA that used a large initial population would take a very long time to run.

The local search algorithm is a more simple optimisation technique than GAs. It acts on a single architecture at a time, rather than a population of architectures. It used the fitness parameter described in section 3.4.3 as the value to be maximised. To produce a new generation of fractionated architecture for evaluation, only the mutation algorithm was used. The number of mutations made to each architecture was called the *degree of mutation*. The degree of mutation used determined how much of the *search space* is explored. The search space was the 4.95×10^{13} possible architectures. A high degree of mutation meant that the search algorithm moved around the search space widely and successive fractionated architectures looked very different from one another. A low degree of mutation resulted in the search algorithm moving around a small area of the search space and successive fractionated architectures looking very similar to one another. To allow the search algorithm to explore as much of the search space as possible and still search around maxima that were found, a large degree of mutation was used at the beginning of the model run, and decreased as the search algorithm progressed. The cycle of architecture evaluations followed by mutations was carried out for 500 generations of fractionated architectures. Six levels of degree of mutation were used, starting with six mutations on each of the parameter. The degree of mutation for each parameter was decreased by a value of one for each $1/6^{\text{th}}$ of the generations explored i.e. after 84 generations the degree of mutation was five, four after 168 generations and so on down to a degree of mutation of one.

The sensitivity of the outputs to the number of generations explored is presented in section 3.7.1.1. The mutation process was always performed on the highest fitness architecture

found thus far in the model run. The final output of the model was the overall highest fitness architecture found. This process is illustrated in Figure 35.

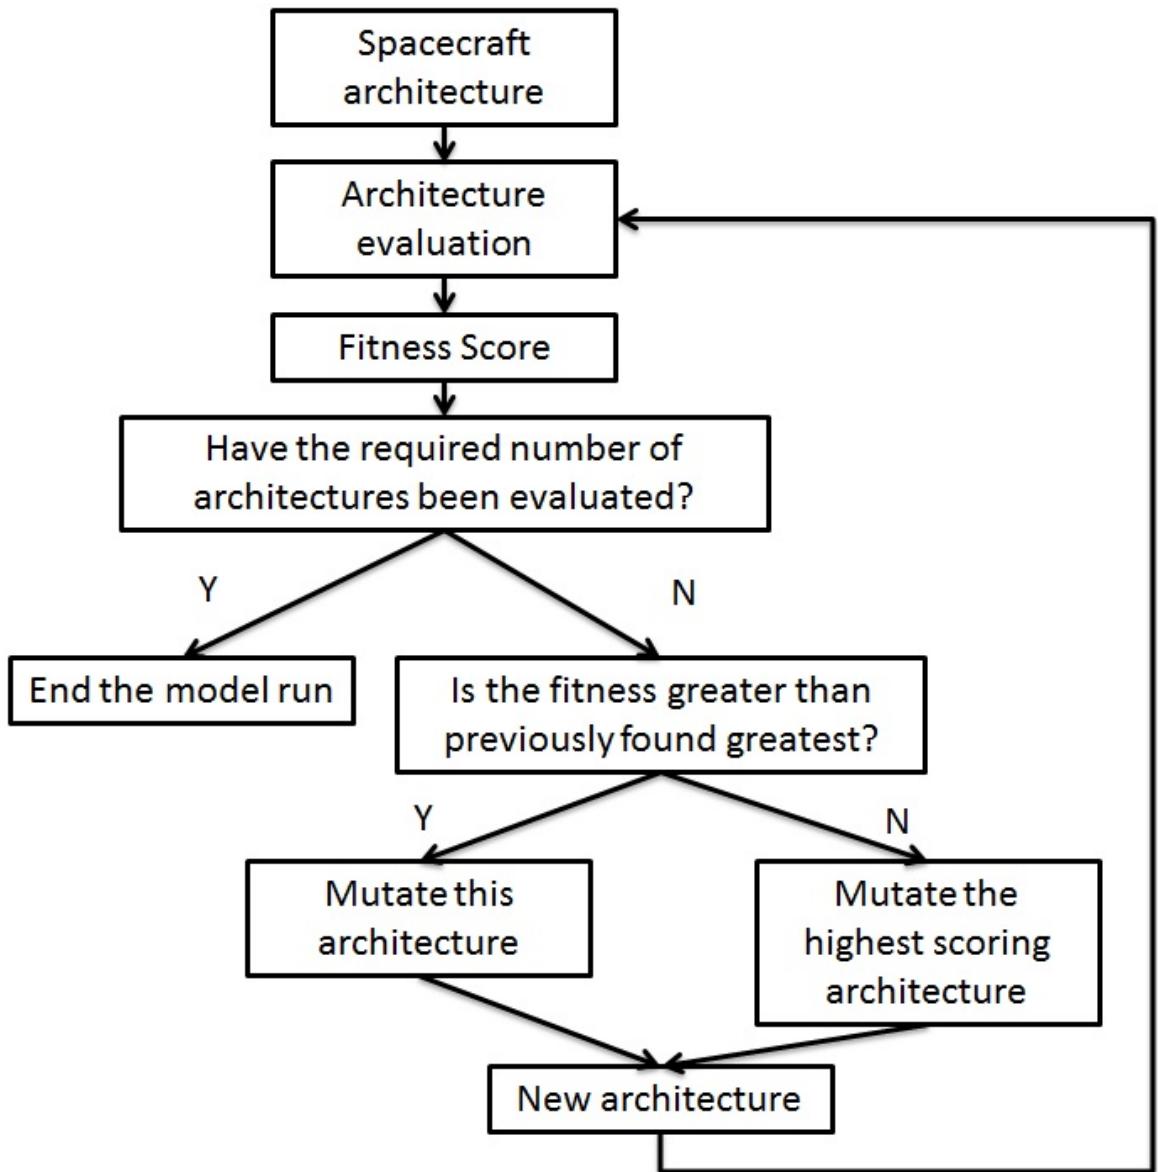


Figure 35 Block diagram showing search algorithm process

The locations of the mutations made to the architecture were random, so the search algorithm may produce new architectures that do not meet the rules for viable fractionated architectures described in section 3.3. To prevent computational time being wasted by evaluating architectures that are not functional, a check is performed on each new architecture and any which do not meet all the rules are discarded and a new mutation is performed to generate an alternative architecture. This process may take several attempts, depending on how sparsely populated the area of the search space is with viable architectures and on the degree of mutation used.

3.6. Testing, debugging and validation

A large proportion of the time spent on the development of the model was devoted to testing, debugging and validation. As the model was a new development and not based on any existing software, this was an important process to go through to ensure that it operated in a way that was consistent with expectations and produced results that were representative of the architectures that were simulated. Text outputs giving detailed information on the failure and replacement of subsystems as the projection progressed as well as the search algorithm's progress were used to validate those processes through manually checking the information presented. Key graphs, such as the relationship between operational time and the number of failures were also compared against expected relationships.

The first version of the model (including architecture evaluation and optimising for operational time only) was implemented for architectures which consisted of nine subsystems (payload, telemetry, telecommunications and control, payload data downlink, on-board computer, data storage, attitude and orbit determination, attitude control, orbital control and power). This gave 8.77×10^{29} possible architectures (given by 2^{81} combinations of satellites and subsystems multiplied by 362880 permutations of fractionated subsystems). This large number of subsystems made the model very complicated to validate. Therefore the decision was taken to reduce the number of subsystems modelled to three for some of the testing and validation activities. The three subsystems consisted of the payload, the power subsystem and the “other” subsystems.

To test the replacement algorithms, the failure curve was adjusted such that all subsystems failed after 7 years of operation. This removed the random nature of the failures and, in theory, every architecture should produce the same operational time of 91.6%, given six periods of loss of operation in the 50 year projection period (one every 7 years) which each result in 0.7 years of operational downtime (4.2 years in total). However some architectures only achieved an operational time of 13.9%. Upon investigation of the failure process it was discovered that this was caused by architectures with either the power subsystem or the “other” subsystem not fractionated. After the initial failure of one of these subsystems the algorithm which checked the architecture for its operational status included the satellites which had experienced these failures. The result was that after this initial failure, the architecture was evaluated as always being non-operational for the remainder of the projection. The solution was to adjust the code so that once satellites have experienced a failure which removes them from the architecture (one in a non-fractionated subsystem or in the final fractionated subsystem on-board) they are no longer included in the operational status check.

Following this testing, the number of subsystems modelled was increased to the six subsystems as described in section 3.3. This represented a compromise between the high fidelity of the nine subsystem case and the simplicity of the three subsystem case.

When running the model and optimising for maximal operational time and minimal mass launched it was found that the zero-mean, unit-variance normalisation was breaking down for architectures which consisted of more than one satellite and were non-fractionated (monolithic constellation). These architectures had an operational time of 100% in every Monte-Carlo run, which resulted in a divide by zero error in the normalisation process. If such an architecture was the first explored by the model the result was that the run would crash. If the search algorithm found such an architecture mid-way through a run of the model then the fitness score was given as “not a number”, and therefore was effectively treated as a non-viable architecture by the search algorithm. The solution to this problem was to assume that any architecture with an average operational time of greater than 99.995% and a variance of 0 had a normalised operational time score of 3.00000000000006, which was greater than the values for architectures with lower operational times.

While investigating the above issue, a bias was discovered which unfairly affected the total launched mass recorded for low degree of fractionation architectures (with the payload as non-fractionated). In the event of a failure in one of the non-fractionated subsystems the replacement algorithm always replaced the satellite. This was because the replacement process was designed assuming that there could only be one payload in the architecture, a rule not obeyed by the search algorithm code which generated new architectures (i.e. new architectures could have more than one payload). The solution was to treat the payload as if it were always “fractionated”. This does not mean that any assumptions were made about sharing of payload data, it was merely a simple method to ensure that the replacement process functioned correctly.

3.7. How the model was used and the data analysed

3.7.1. Sensitivity analysis

In addition to the number of Monte Carlo runs, there were two other parameters that had an effect on the results: the number of generations of fractionated architectures explored by the search algorithm and the time-step used. To understand the influence of these parameters, a sensitivity analysis was performed. The baseline conditions for this analysis were: 50 year projection period, 300 Monte-Carlo simulations performed on each fractionated architecture, 500 generations of fractionated architectures explored and a time-step of 0.05 years. The last two

parameters were varied as a part of their respective sensitivity analyses with the other remaining constant.

3.7.1.1. Number of architectures explored by the search algorithm

To understand how the number of different fractionated architectures explored by the search algorithm affected the outputs, the model was run using the baseline conditions given above, changing only the number of architectures explored. Different values for this parameter were used between 100 and 3000 architectures and the fitness and standard deviation of the average fitness of each architecture were recorded. The results are shown in Figure 36 to Figure 39 where only the viable architectures in each run are plotted.

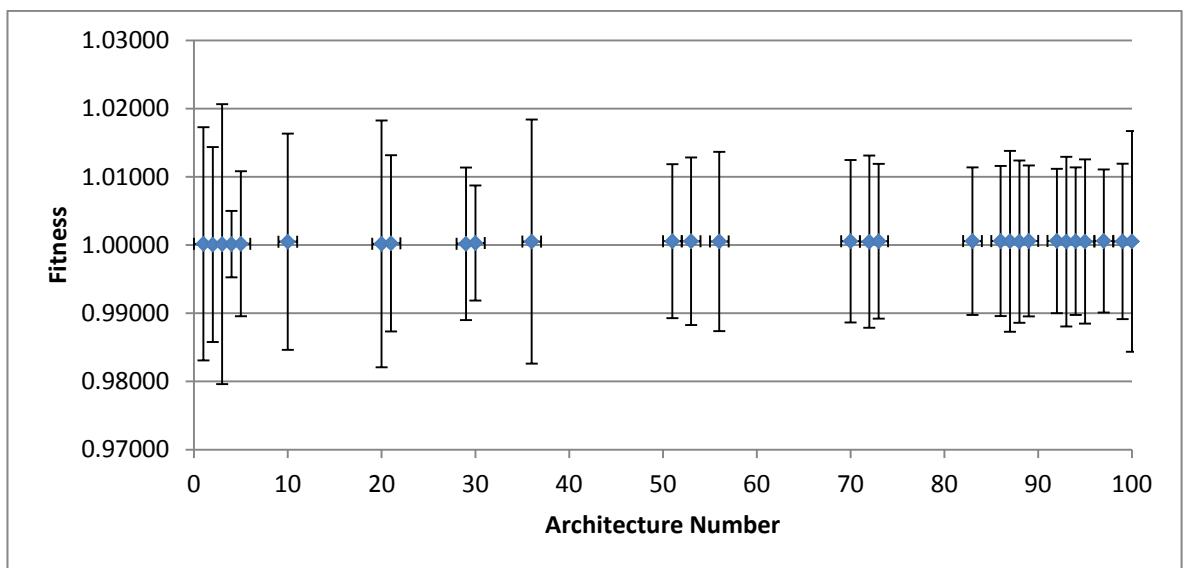


Figure 36 Model run with 100 architectures explored and 1 standard deviation error bars

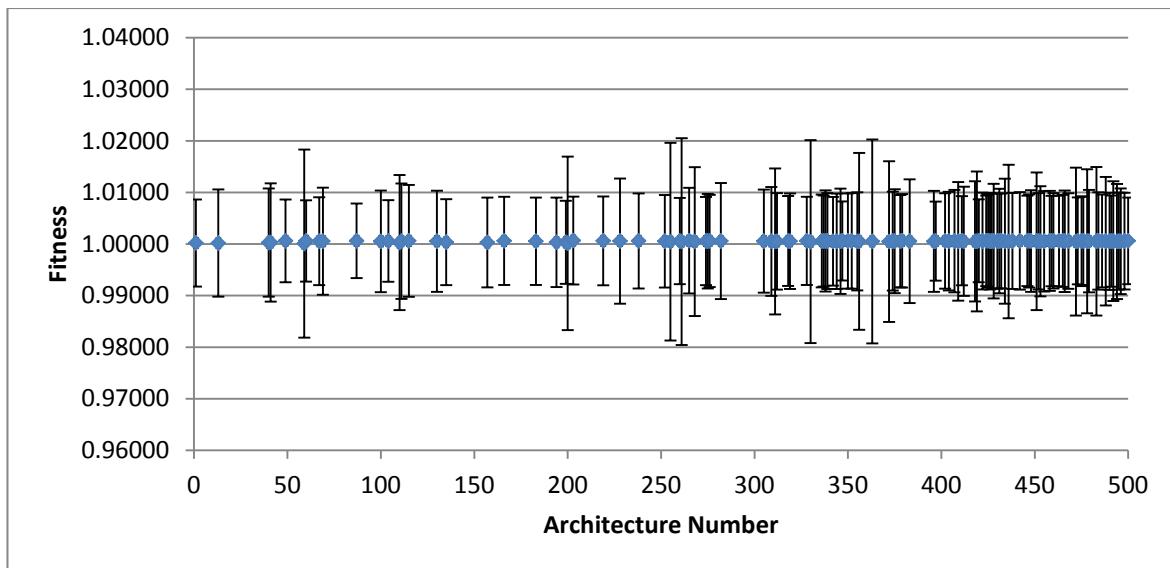


Figure 37 Model run with 500 architectures explored and 1 standard deviation error bars

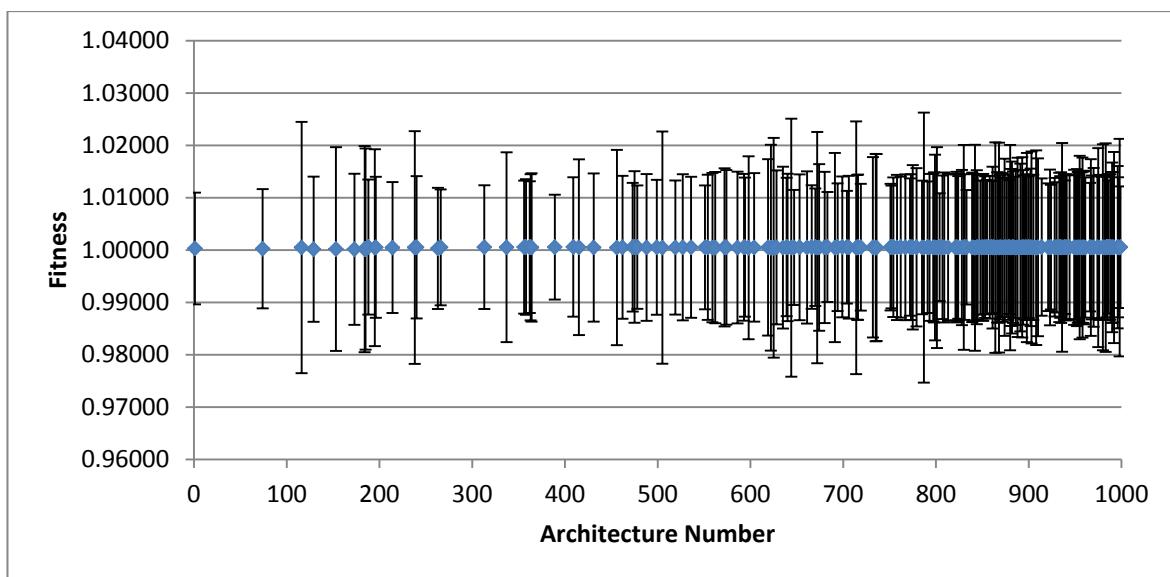


Figure 38 Model run with 1000 architectures explored and 1 standard deviation error bars

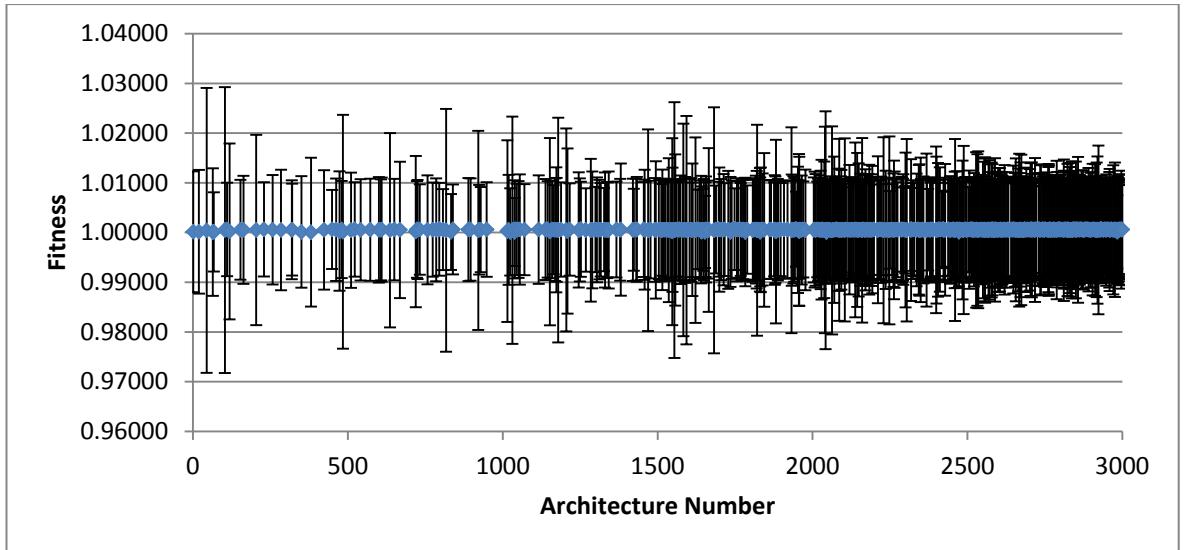


Figure 39 Model run with 3000 architectures explored and 1 standard deviation error bars

The number of architectures explored by the algorithm dictates how long the search algorithm explores the search space at each degree of mutation level. This is shown in Table 16.

Total number of architectures evaluated	Run time	Number of architectures evaluated at each degree of mutation level	Percentage of architectures that were viable	Maximum fitness found
100	1 hour 55 minutes	17	29%	1.0006042307
500	9 hours 6 minutes	84	24%	1.0006623954
1000	17 hours 32 minutes	167	20%	1.0005954398
3000	45 hours 1 minute	500	22%	1.0006839159

Table 16 Total number of generations evaluated, run time, number of generations evaluated at each mutation level, percentage of generations with viable architectures and maximum fitness found for the number of generations sensitivity analysis

The large proportion of viable architectures being evaluated towards the end of the model runs shown in Figure 36 to Figure 39 are a result of the decreasing degree of mutation used in the search algorithm. At this stage of the run, the degree of mutation is at its smallest, and the algorithm is more likely to find a new viable architecture in proximity to the previous architecture than at earlier stages. A balance is needed between the time in which the search algorithm uses each degree of mutation and the time to run the model. For all four of the model runs, only 20%

to 30% of the architectures produced by the mutation algorithm were viable. The 500 architectures run represents a good compromise with a run taking around 9 hours and the search algorithm searches at each mutation level for 84 generations.

3.7.1.2. Time-step

The time-step is a key parameter, as it defines the accuracy of the model. A shorter time-step gives a more accurate reconstruction of the failure curve(s). Figure 40 to Figure 43 give the same failure curve plotted using time-steps of 0.05, 0.5, 1, and 5 years. These figures show the decreasing accuracy with which the failure curve is generated. Table 17 gives the total probability of failure for each of the failure curves generated, given by the area under the curves. This shows the total probability of failure decreases as the time-step increases.

Time-step (years)	Total Probability of failure
0.05	0.068108736
0.5	0.063406332
1	0.062525857
5	0.057650267

Table 17 Time-step and total probability of failure for the probability density functions generated using the time-steps given

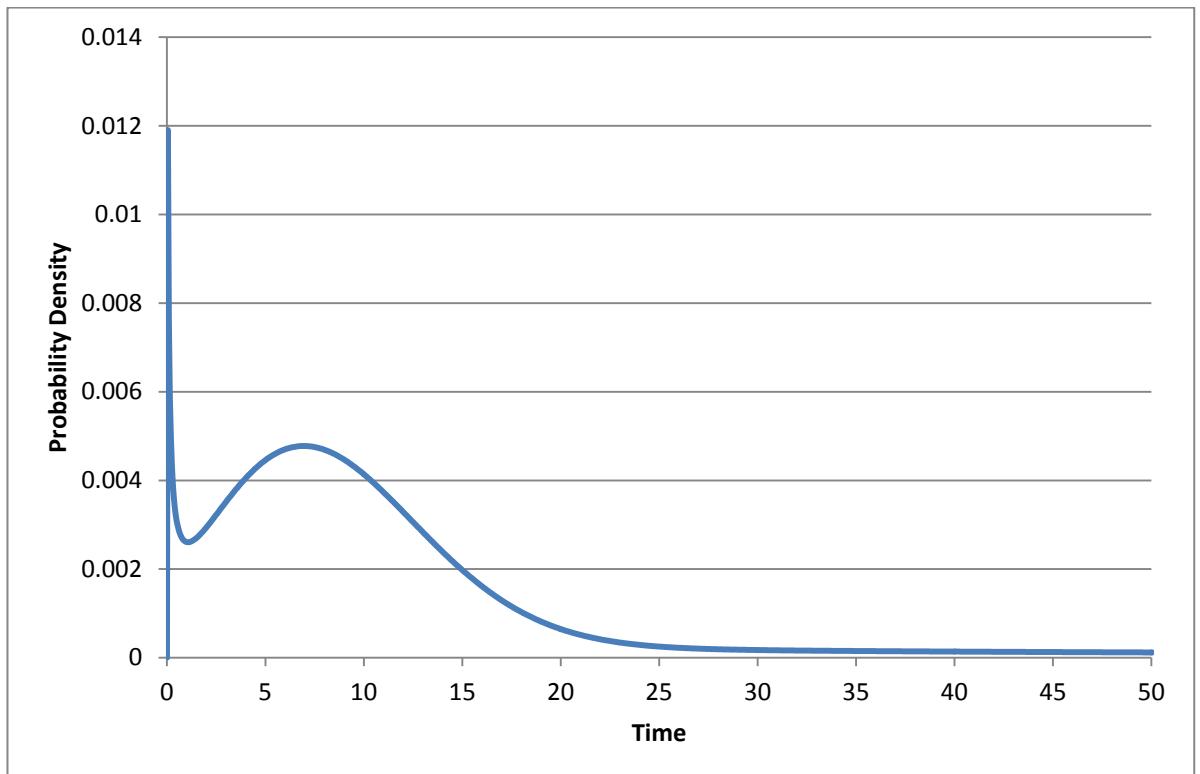


Figure 40 2-Mixed Weibull probability density function generated with a time-step of 0.05 years

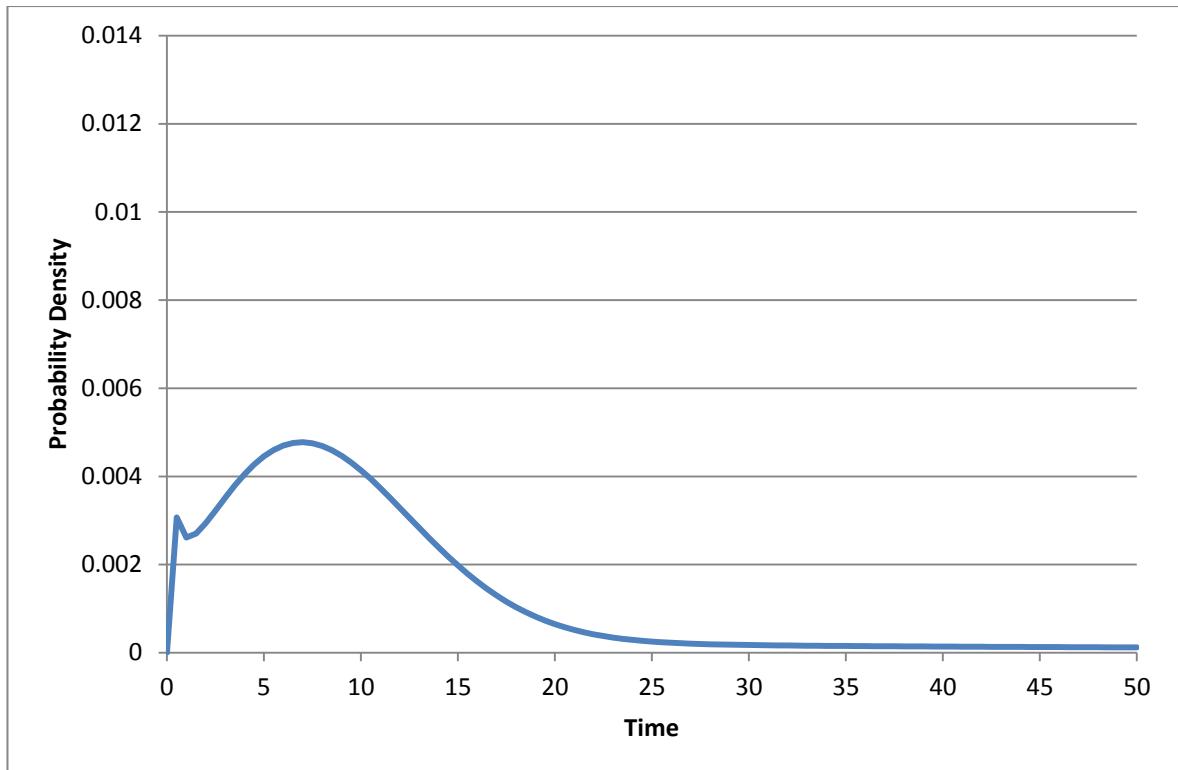


Figure 41 2-Mixed Weibull probability density function generated with a time-step of 0.5 years

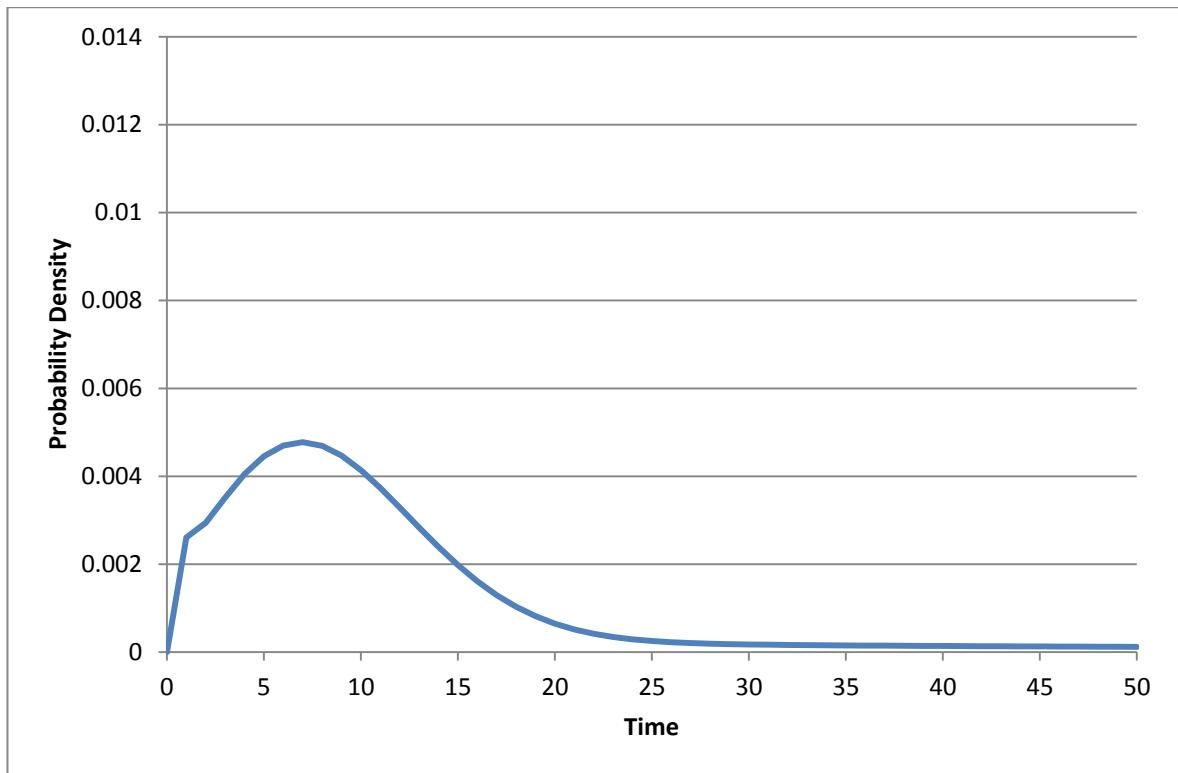


Figure 42 2-Mixed Weibull probability density function generated with a time-step of 1 year

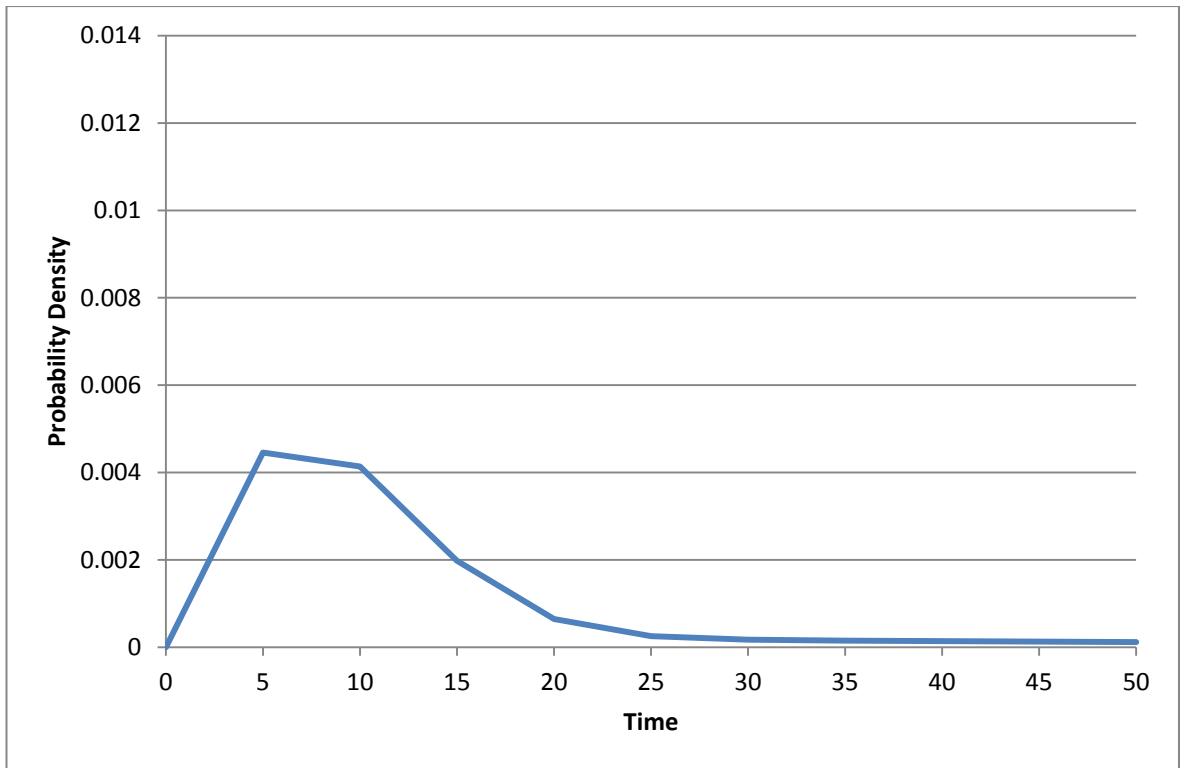


Figure 43 2-Mixed Weibull probability density function generated with a time-step of 5 years

A shorter time-step resulted in longer simulation time. Figure 44 to Figure 48 show model runs consisting of 500 architectures explored with each one subjected to 500 Monte-Carlo simulations for time-steps of 0.05 years, 0.1 years, 0.5 years, 1 year and 5 years. The fitness value for each viable architecture is plotted with the error bars showing one standard deviation. Table 18 gives the model run time and average standard deviation for each of the model runs.

Time-step (years)	Model run time	Mean standard deviation
0.05	10 hours 6 minutes	0.0090375525
0.1	8 hours 1 minute	0.0158612677
0.5	4 hours 20 minutes	0.0201131190
1	34 minutes	0.0133996337
5	11 minutes	0.0313134400

Table 18 Time-step, model run time and average standard deviation for the time-step sensitivity analysis

Increasing the time-step had the effect of increasing the standard deviation, which is the manifestation of the accuracy with which the failure curve was sampled. Although the mean standard deviation was comparable for the final architectures in the 0.05, 0.1 and 0.5 years graphs (0.009038, 0.015861 and 0.020113 respectively), the maximum standard deviation is much higher for the longer time-steps (0.02979 for the 0.05 years time-step compared to 0.04742 and 0.08799 for the 0.1 and 0.5 years time-steps respectively). It was therefore decided to use the

0.05 years time step for maximum fidelity in the model, and the longer run time was deemed acceptable to achieve this level of accuracy.

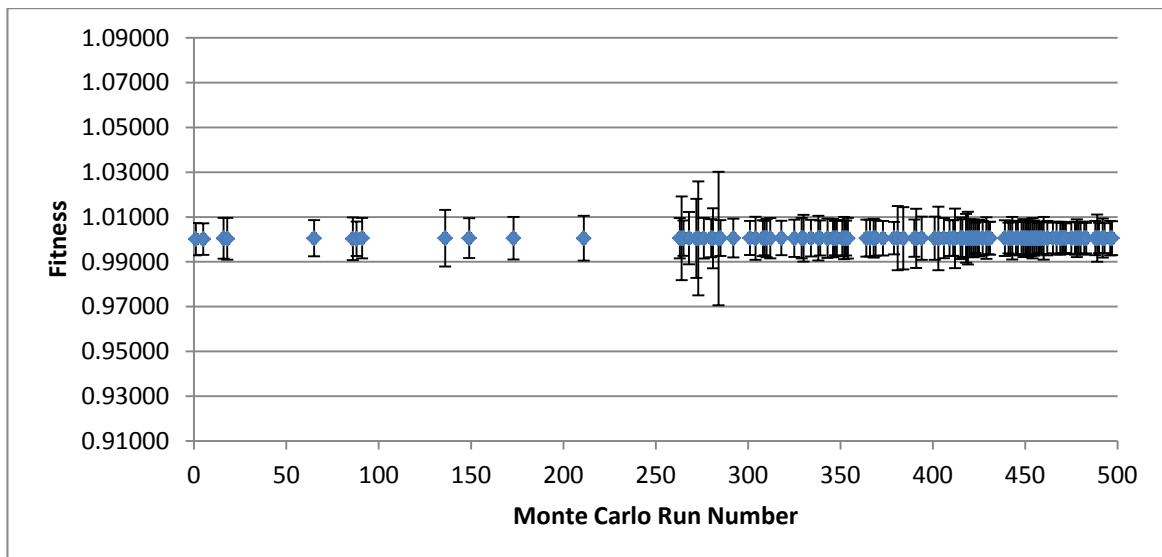


Figure 44 Model run with a time step of 0.05 years and 1 standard deviation error bars

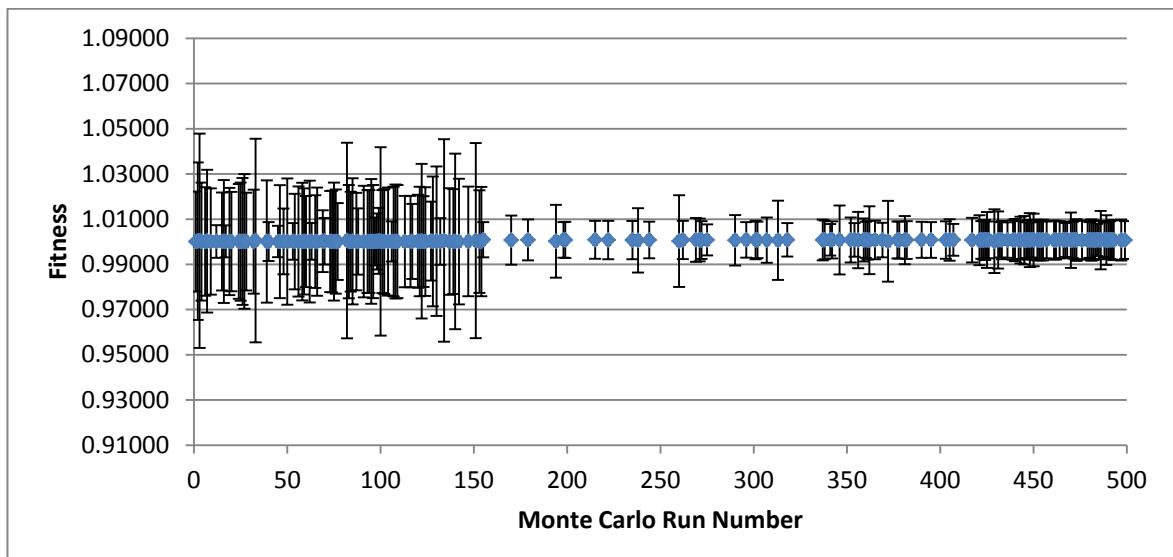


Figure 45 Model run with a time step of 0.1 years and 1 standard deviation error bars

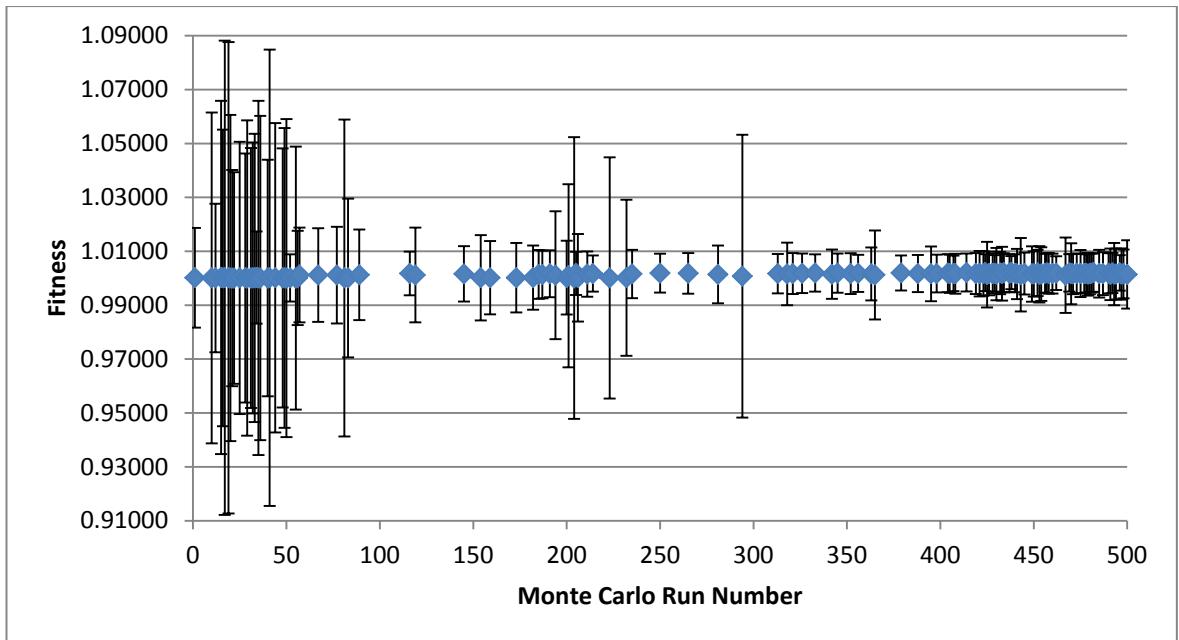


Figure 46 Model run with a time step of 0.5 years and 1 standard deviation error bars

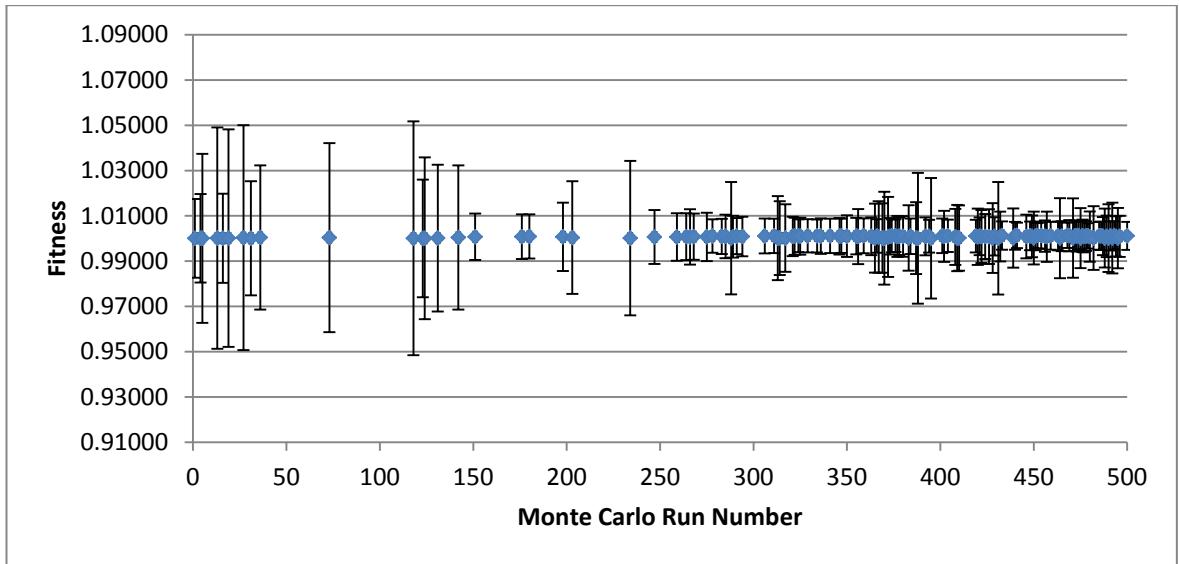


Figure 47 Model run with a time step of 1 year and 1 standard deviation error bars

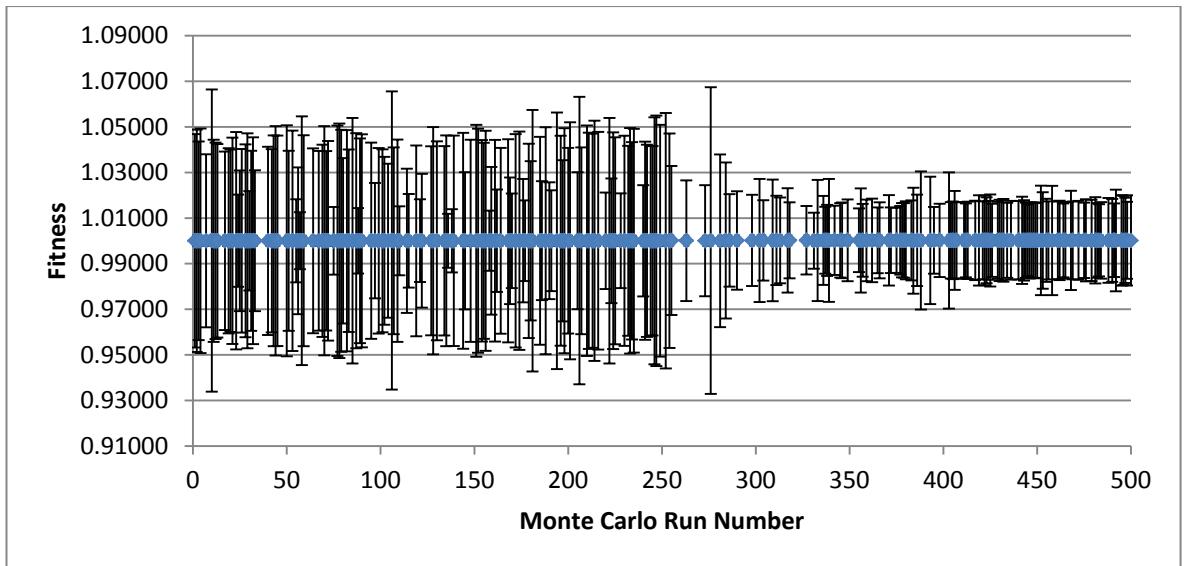


Figure 48 Model run with a time step of 5 years and 1 standard deviation error bars

3.7.2. How the model was used

The local search algorithm described in section 3.5 is not efficient in fully traversing the search space in a single run, as it can get stuck around local maxima. In addition, as shown in Table 16, only around 20% of all the architectures explored in each model run are viable. Therefore, multiple runs of the model were required, starting in different parts of the search space. For each scenario, a full data set comprised of fifteen separate model runs. This set of model runs was made up of three runs with randomly determined starting architectures and twelve with pre-determined architectures consisting of two runs each of the following starting architectures: a single monolithic satellite (Table 19).

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	1	1	1	1	1

Table 19 Matrix format for a single monolithic satellite

A six-satellite monolithic architecture (Table 20).

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	1	1	1	1	1
2	1	1	1	1	1	1
3	1	1	1	1	1	1
4	1	1	1	1	1	1
5	1	1	1	1	1	1
6	1	1	1	1	1	1

Table 20 Matrix format for a six satellite monolithic architecture

A six-satellite architecture, which is fully fractionated and a single subsystem on each satellite (Table 21). The fractionated subsystems are shown shaded in blue.

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	0	0	0	0	0
2	0	1	0	0	0	0
3	0	0	1	0	0	0
4	0	0	0	1	0	0
5	0	0	0	0	1	0
6	0	0	0	0	0	1

Table 21 Matrix format for a six satellite architecture which is fully fractionated with a single subsystem on each satellite

A single satellite, which is fully fractionated (Table 22). The fractionated subsystems are shown shaded in blue.

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	1	1	1	1	1

Table 22 Matrix format for a single satellite which is fully fractionated

A six-satellite architecture, which is fully fractionated and each satellite has a full complement of subsystems (Table 23). The fractionated subsystems are shown shaded in blue.

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	1	1	1	1	1
2	1	1	1	1	1	1
3	1	1	1	1	1	1
4	1	1	1	1	1	1
5	1	1	1	1	1	1
6	1	1	1	1	1	1

Table 23 Matrix format for a fully fractionated architecture of six satellites with a full complement of subsystems on each

A three-satellite architecture, which is fully fractionated and each satellite has a full complement of subsystems (Table 24). The fractionated subsystems are shown shaded in blue.

Satellite Number	Payload	Communications	OBDH	ADCS	OCS	Power
1	1	1	1	1	1	1
2	1	1	1	1	1	1
3	1	1	1	1	1	1

Table 24 Matrix format for a fully fractionated architecture of three satellites with a full complement of subsystems on each

Architectures which are fully fractionated and contain satellites which carry all the subsystems operate using the hybrid approach to the replacement of failed subsystems outlined in section 1.1.2. This meant that as replacement subsystems were launched, they consisted only of the new subsystem and would connect wirelessly to the rest of the architecture to share their resource.

For each of these scenarios, a model run was performed comprising 300 Monte-Carlo simulations per architecture and 500 generations of architectures explored. In addition to the fitness parameter and the associated standard deviation on the average, an additional eleven parameters are produced for each architecture. These fall into three categories: the architecture properties, the failure properties and the outputs from the evaluations. The architecture properties were: $N_{Satellites}$, D , $N_{Subsystems}$, F and E . The failure properties were: n_f , a_f , n_{SL} and n_{fs} . The evaluation outputs were the M and the percentage of the 50-year lifetime that the architecture

was operational (t_{op}). The failure properties and evaluation outputs that were used in the analyses were the averaged values from the Monte-Carlo runs.

Table 25 describes all the scenarios explored using the model, giving the failure curve used, the replacement time, the optimisation parameter, the purpose and the section where each scenario is presented.

Scenario number	Failure curve	Replacement time	Optimisation parameter	Purpose	Section
1	Simple designed curve	0.7 years	Operational time	Data exploration	4.1.1
2	Spacecraft level 2-mixed Weibull curve	0.7 years	Normalised fitness parameter	Data exploration	4.1.2
3a	Subsystem level 2-mixed Weibull curves	0.7 years	Normalised fitness parameter	Results	4.2.1.1
3b		0.05 years	Normalised fitness parameter	Results	4.2.1.2
3c		1.5 years	Normalised fitness parameter	Results	4.2.1.3

Table 25 Table describing all scenarios explored using the model

3.7.3. Analysis method

A probabilistic approach was taken to interrogating the model and analysing the outputs. This allowed the following question to be addressed: for a given (high) fitness value observed, what are the most likely architecture properties and failure characteristics to have produced it? The answer to this question requires the estimation of the conditional probability, $P(b_i | S)$ where b_i is one of the following six model parameters: $N_{satellites}$, $N_{subsystems}$, D , F , n_f and a_f and S is the optimisation parameter used (operational time for the first scenario followed by the fitness parameter outlined in section 3.4.3). The first four model parameters are architecture properties and the last two are failure properties. $P(S)$ gives the probability of obtaining a particular fitness or operational time. Given a particular value of S is obtained using the model, (e.g. fitness is “high”), then $P(b_i | Fitness)$ is the probability that a particular input value was used in obtaining that fitness result. This is the conditional probability, where

$$P(b_i | S) = \frac{P(b_i, S)}{P(S)} \quad 27,$$

and $P(b_i, S)$ is the probability that input b_i and fitness output occur together in the model (i.e. the joint probability). When calculated over the range of fitness values, the conditional probabilities, $P(b_i | S)$, can be described using a histogram.

4. Results and data analysis

In this chapter the key findings from the use of the model are presented. Two model runs are presented in section 4.1. This section explores the data produced by the model runs using simpler failure curves and optimisation parameters. The outputs from this data exploration were used to inform the final version of the model, presented in section 4.2. For each of the scenarios presented, the justification for each set of conditions and the assumptions made are given. This is followed by the key findings for each scenario, and the graphs that illustrate these findings. Each dataset produced for each condition contains over 100 graphs, and it would be impractical to present every graph here, therefore, only a selection of graphs is presented.

4.1. Data exploration

In this section two data sets are presented. In the first, the search algorithm identified architectures that maximised the operational time only. This allowed the contribution of the operational time to the outputs to be understood before the more complex optimisation parameter described in section 3.4.3 is introduced into the model in the second data set.

4.1.1. Scenario 1

This scenario used the designed failure curve described in 3.4.2.1. This curve (Figure 49) was applied to all the subsystems equally, using the method described in section 3.4.2.1.

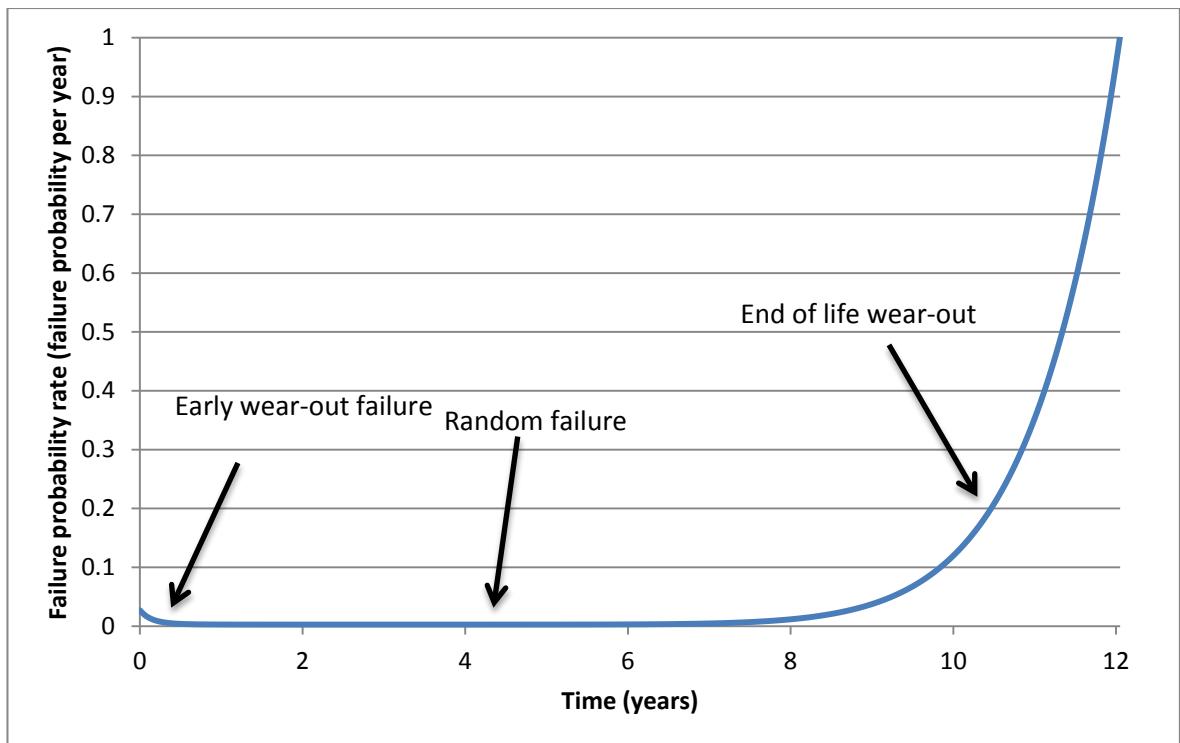


Figure 49 Failure rate curve for scenario 1. On the x-axis is the age of the subsystem and on the y-axis is the failure rate in failures per year.

This scenario represented the first opportunity to learn empirically about the relationships between the architecture parameters, failure properties and outputs. Although the failure curve used for this data set was only an approximation and not based on real satellite data, it was suitable for use in this case as a test of the model where the primary purpose is to understand the way the model functions. The behaviour of the model can therefore be understood before more complex features are used. More accurate failure curves, incorporating different failure properties for different subsystems, were added in later scenarios.

4.1.1.1. Results and data analysis

In this section, the analysis of the data produced by the first scenario is presented. Figure 50 gives the probability distribution, $P(t_{op})$ for this scenario, showing that high operational times ($t_{op} \geq 90\%$) are at least four times more likely to occur in the dataset than low operational times ($t_{op} \leq 80\%$). This graph also shows that some operational times do not exist in the dataset, for example 92% or 81%. This is due to the random path through the search space that the search algorithm takes and the random nature of the subsystem failures.

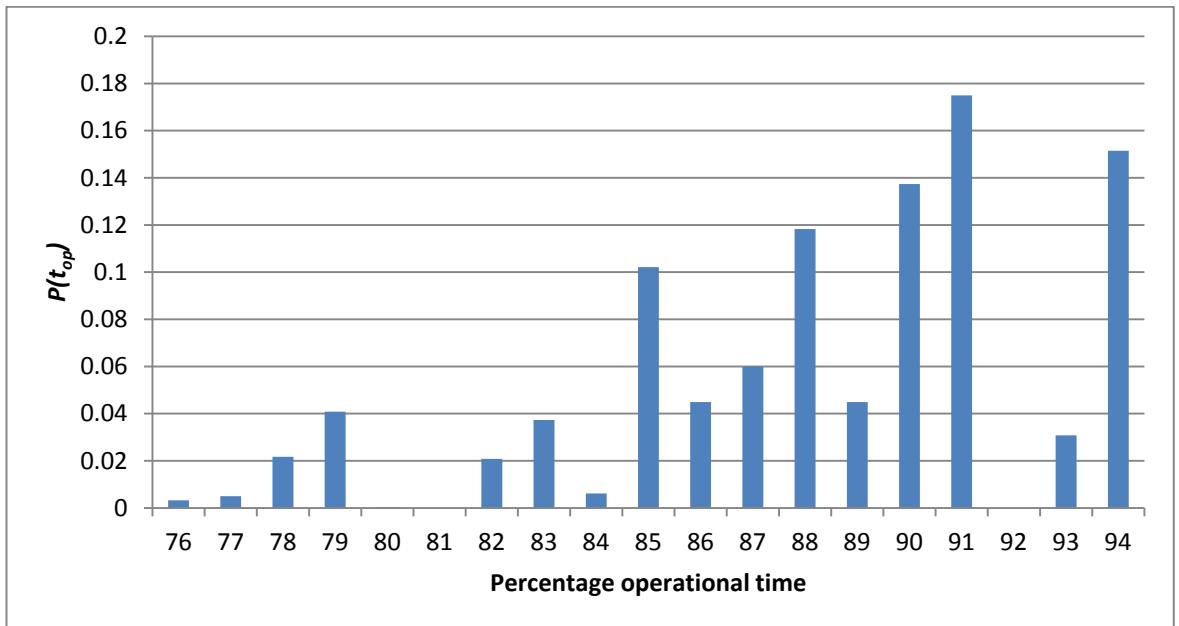
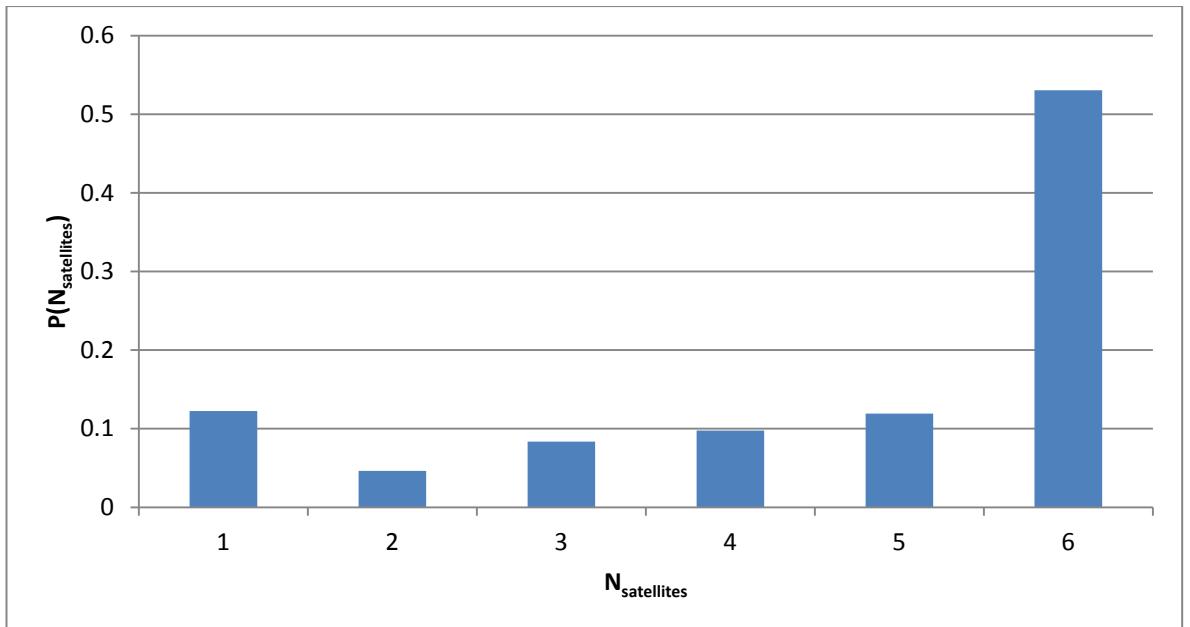
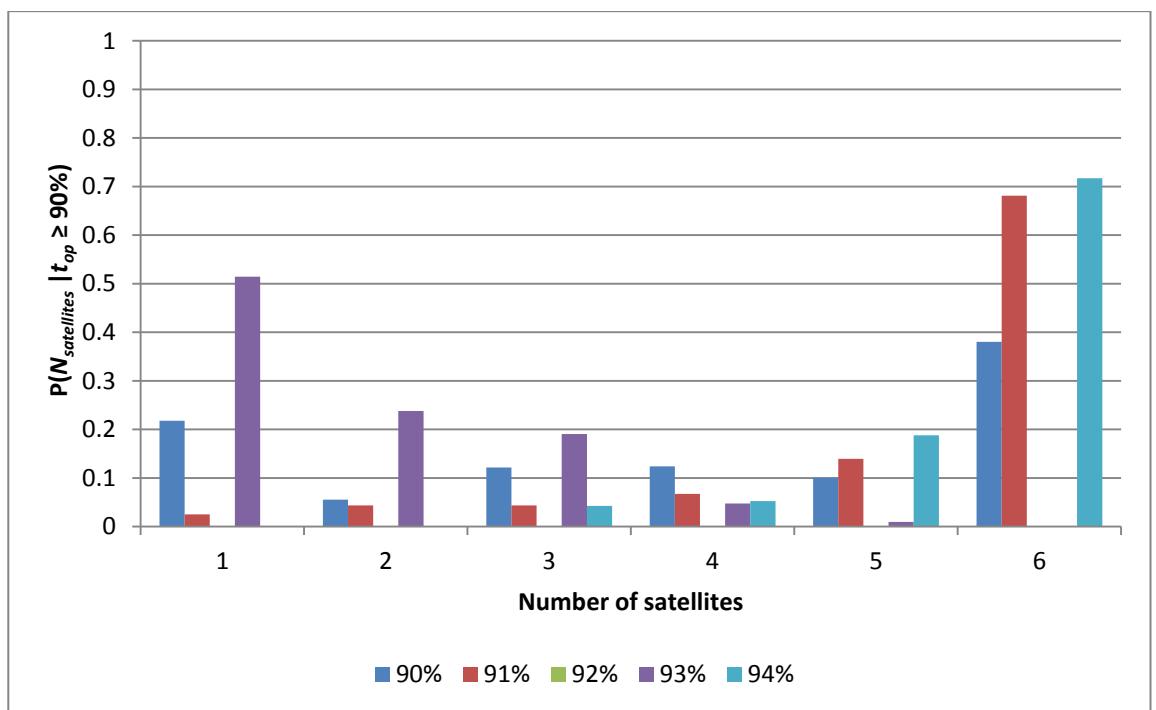


Figure 50 Probability distribution for $P(t_{op})$ for scenario 1

Figure 51 gives the probability distribution for $P(N_{satellites})$. This graph shows that six satellite architectures are at least five times more likely to occur in the dataset. Figure 52 gives the conditional probability distribution $P(N_{satellites}/t_{op})$ for high operational time architectures ($t_{op} \geq 90\%$) while Figure 53 gives the distribution for low operational time architectures ($t_{op} \leq 80\%$). Both graphs show that $N_{satellites} = 6$ are the most likely to produce both high and low operational times. However in the context of probability distribution, given in Figure 51 this is not a surprising result.

Figure 51 Probability distribution for $P(N_{\text{satellites}})$ for scenario 1Figure 52 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} \mid t_{\text{op}} \geq 90\%)$

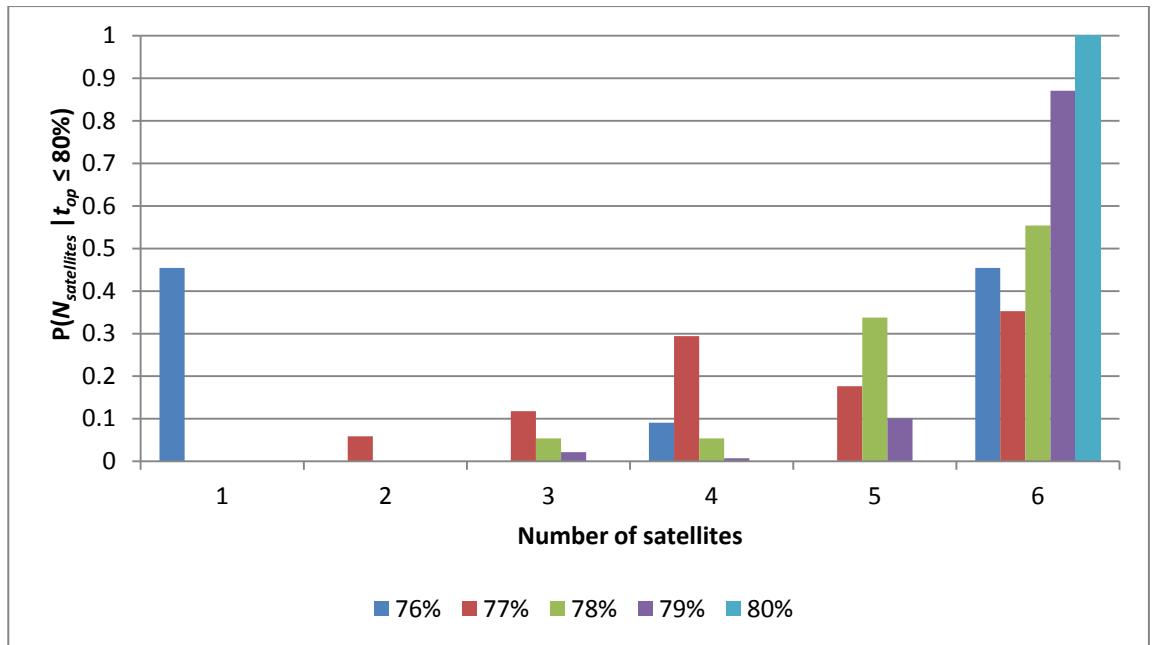
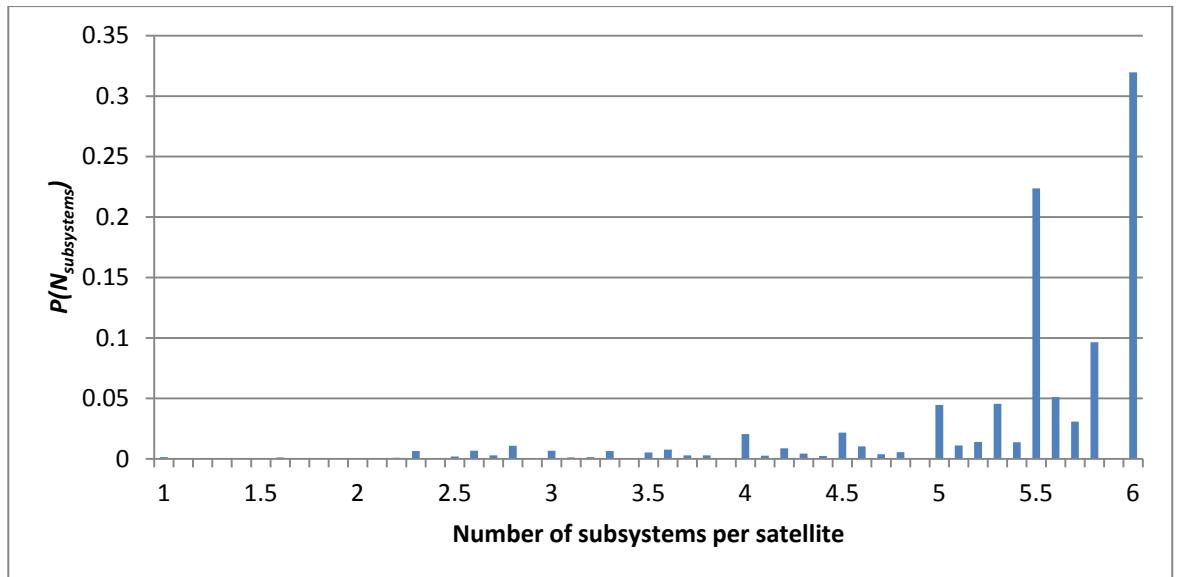
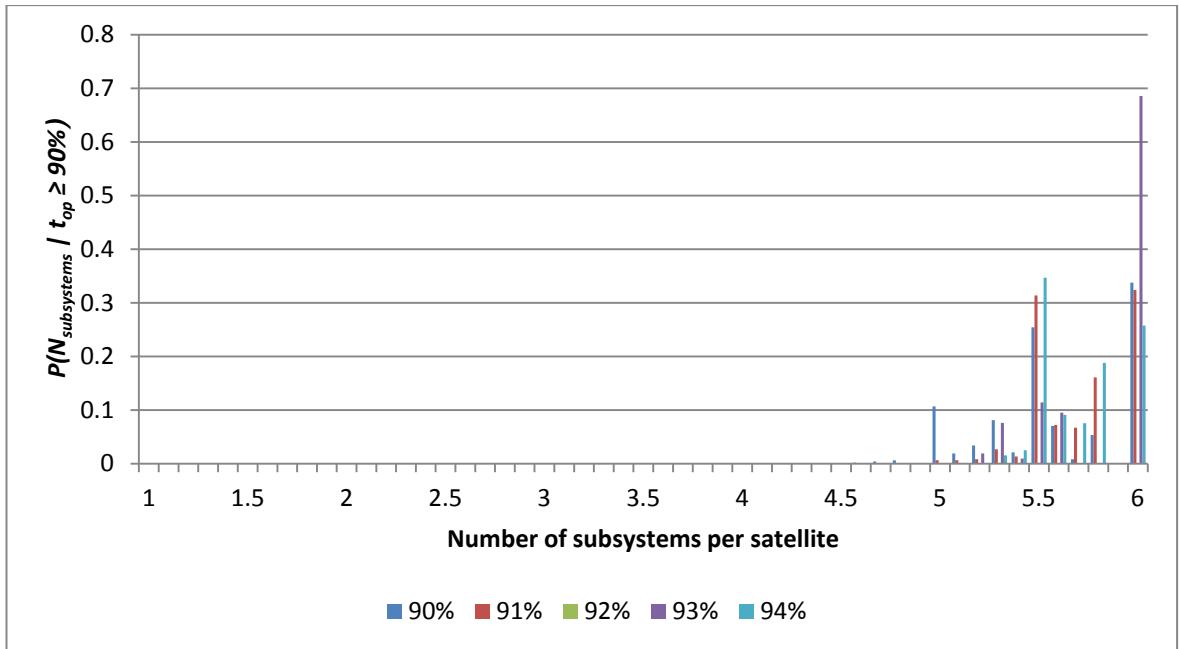


Figure 53 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} | t_{\text{op}} \leq 80\%)$

Figure 54 gives the probability distribution, $P(N_{\text{subsystems}})$ for this scenario, and shows that 85% of the architectures evaluated in this data set had $N_{\text{subsystems}} \geq 5$. Figure 55 gives the conditional probabilities for $P(N_{\text{subsystems}} | t_{\text{op}} \geq 90\%)$ and shows that operational times of greater than 90% are only produced by architectures with $N_{\text{subsystems}} \geq 5$. Figure 56 gives the conditional probabilities for $P(N_{\text{subsystems}} | t_{\text{op}} \leq 80\%)$, showing that low operational times are most likely to be the result of $N_{\text{subsystems}} \leq 4$. There are anomalous data points at $N_{\text{subsystems}} = 6$ caused by a small sample sizes for the $t_{\text{op}} = 80\%$ and 76% cases, illustrated in the probability distribution shown in Figure 50.

**Figure 54 Probability distribution for $P(N_{\text{subsystems}})$ for scenario 1****Figure 55 Histogram showing conditional probability distributions for $P(N_{\text{subsystems}} | t_{\text{op}} \geq 90\%)$**

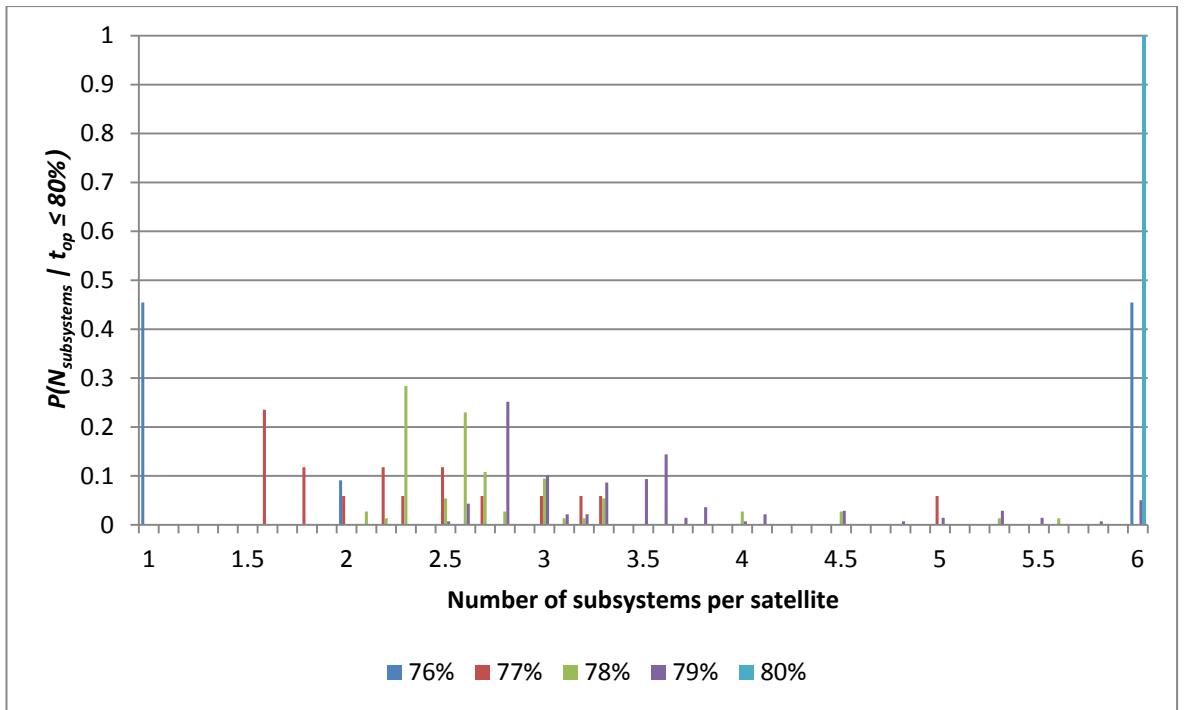
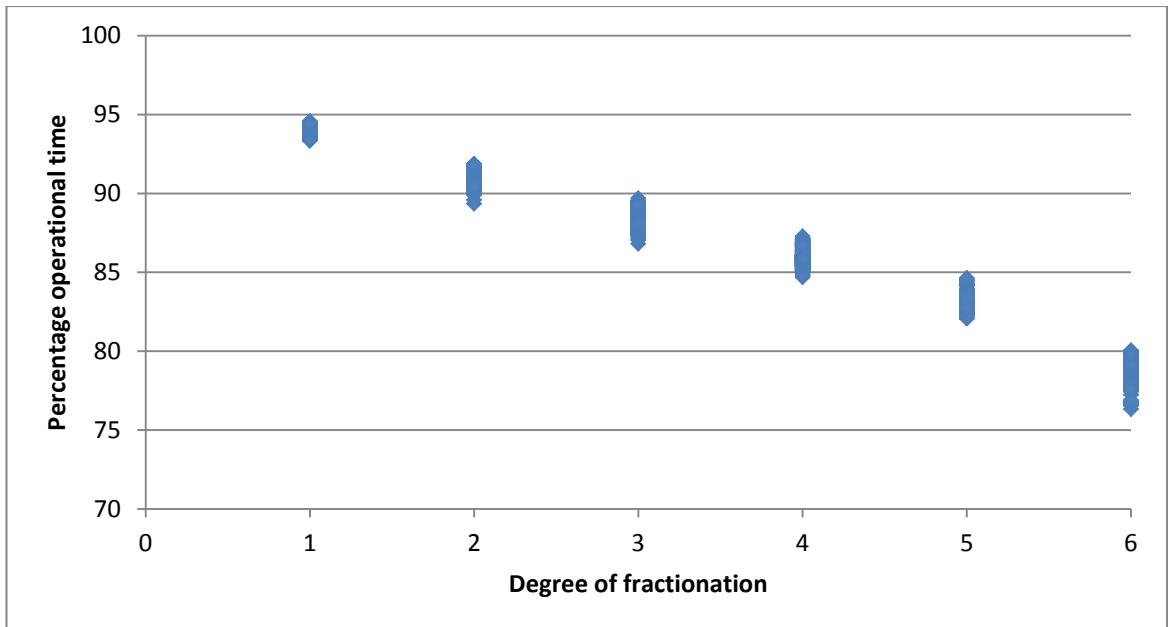
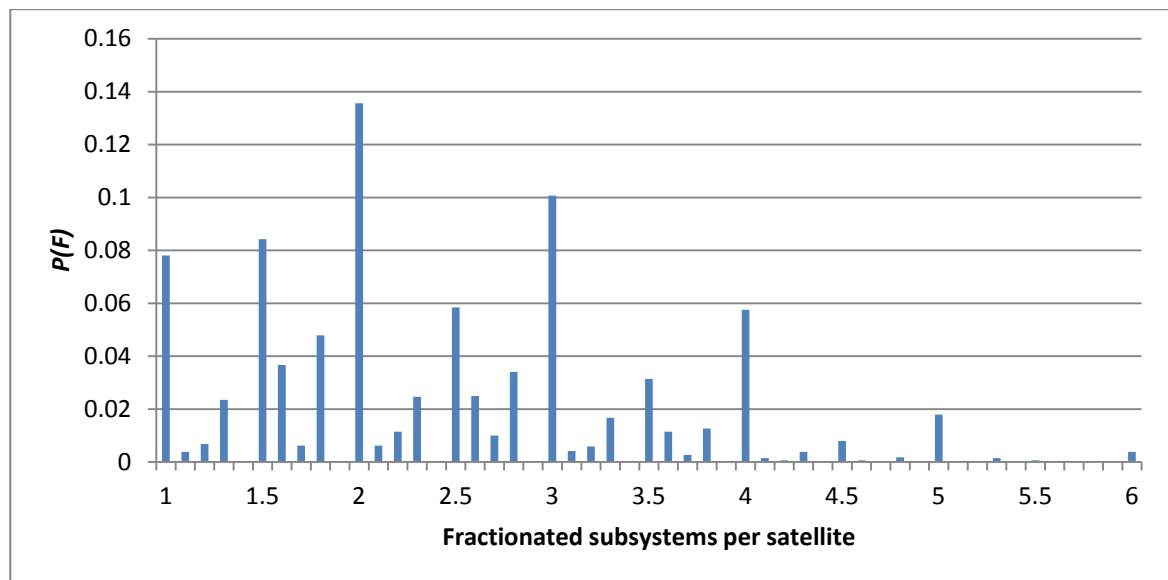


Figure 56 Histogram showing conditional probability distributions for $P(N_{\text{subsystems}} | t_{\text{op}} \leq 80\%)$

Figure 57 shows that operational time was inversely proportional to D . Only architectures with $D = 1$ or 2 exclusively resulted in operational times greater than 90%. This result means that, when optimising for operational time only, more monolithic architectures are preferable. Architectures with low values for D must also have had low values for F (i.e. operational time is inversely proportional to F). Figure 58 shows that 83% of architectures evaluated in this dataset had $F \leq 4$. Figure 59 shows that operational times of greater than 90% were only observed from architectures with $F \leq 2$, supporting the finding illustrated in Figure 57 described above. Figure 60 shows that operational times of less than 80% were experienced by architectures with the full range of values of F .

Figure 57 D versus t_{op} for scenario 1Figure 58 Probability distribution for $P(F)$ for scenario 1

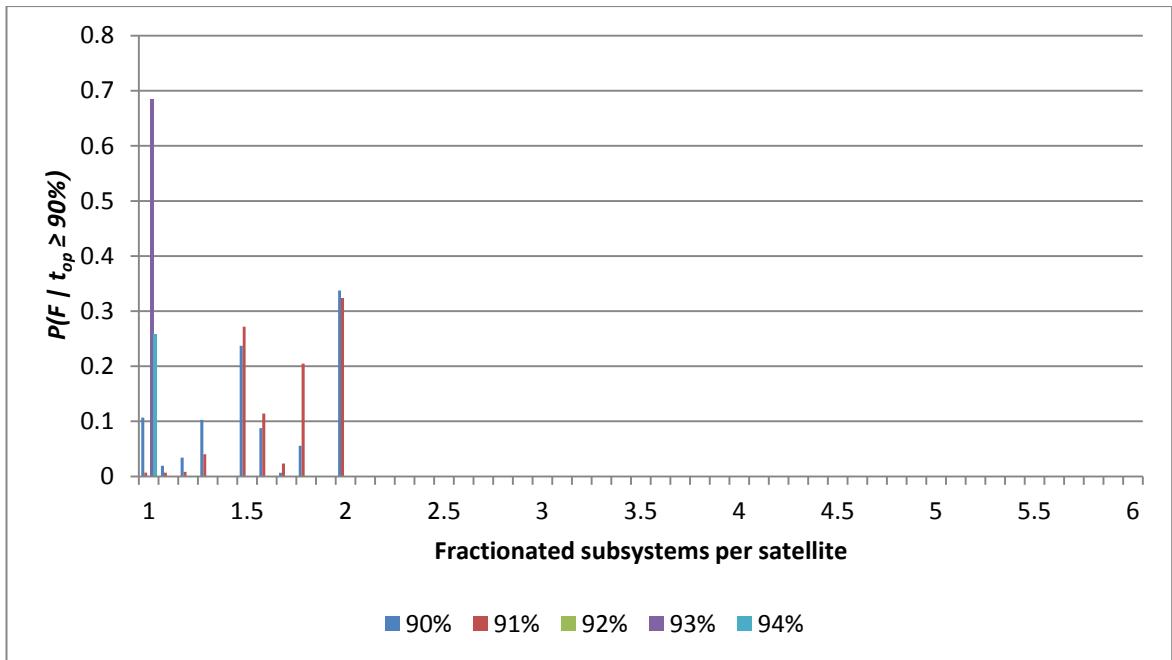


Figure 59 Histogram showing conditional probability distributions for $P(F / \text{top} \geq 90\%)$

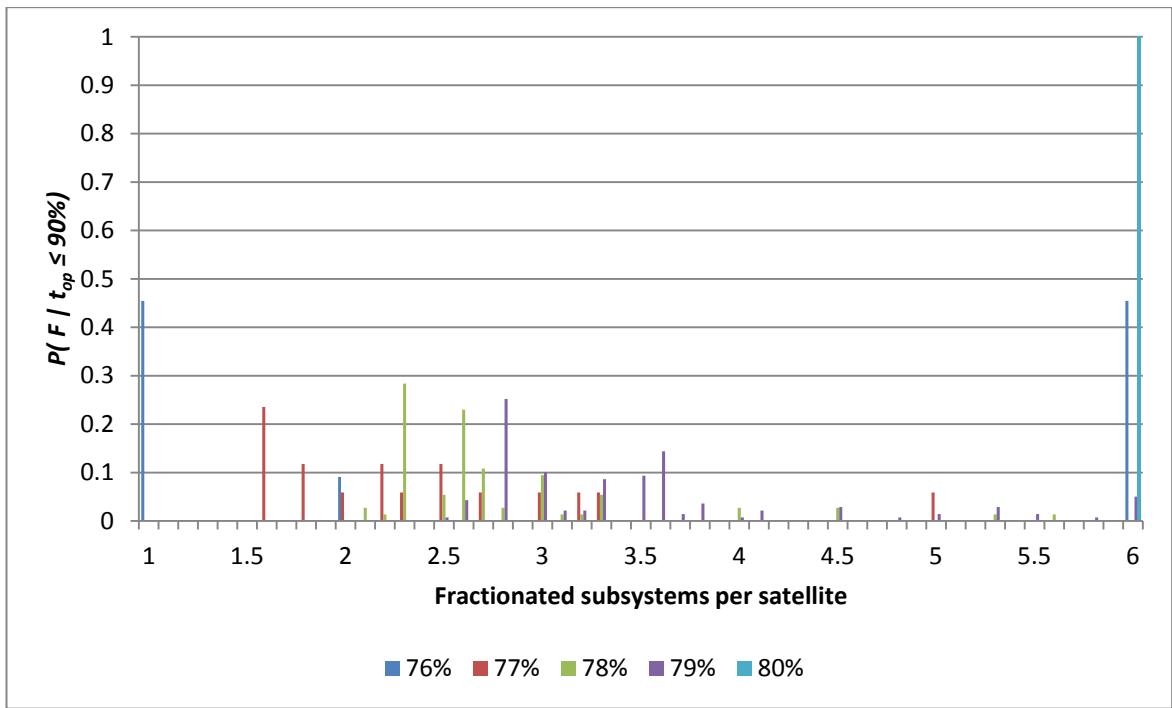


Figure 60 Histogram showing conditional probability distributions for $P(F / \text{top} \leq 80\%)$

Figure 50 to Figure 60 showed that architectures with, $N_{\text{satellites}} = 6$, $N_{\text{subsystems}} \geq 5$, $D \leq 2$, $F \leq 2$ (essentially a constellation of monolithic spacecraft), produced the highest operational times. Figure 61 gives the probability distribution for $P(n_f)$ and shows that 83% of architectures

experienced less than 25 failures. Figure 62 shows that all architectures which achieved operational times of greater than 90% experienced 16 or fewer failures while the lowest operational times were experienced by architectures with 22 or more failures (Figure 63). Figure 64 shows that a_f was inversely proportional to the operational time.

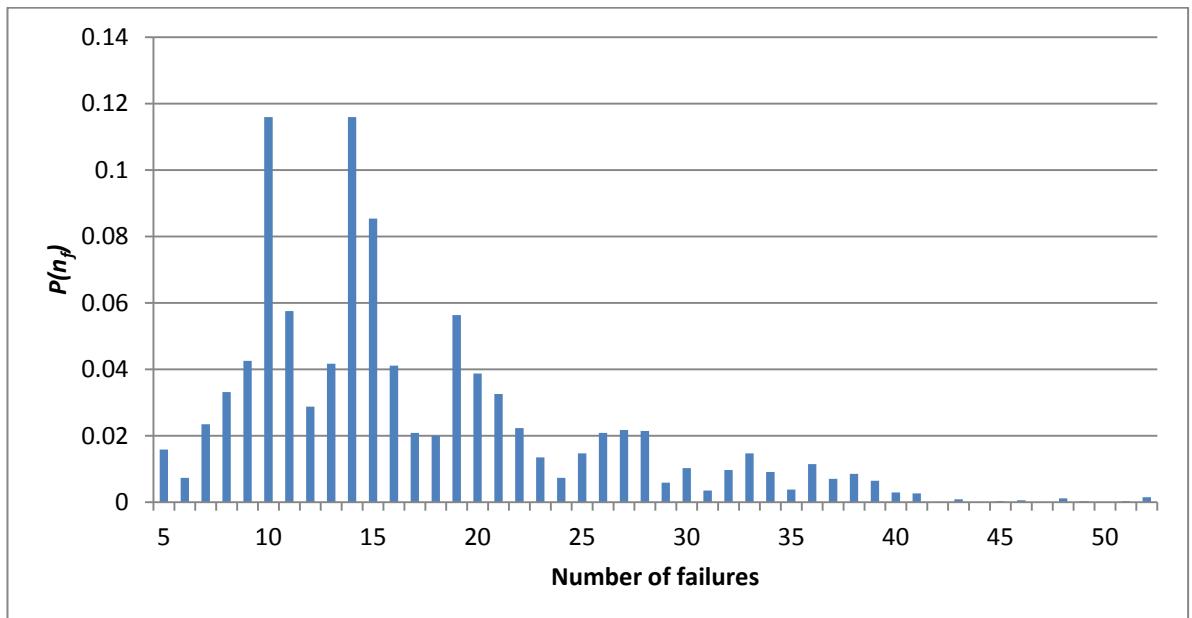


Figure 61 Probability distribution for $P(n_f)$ for scenario 1

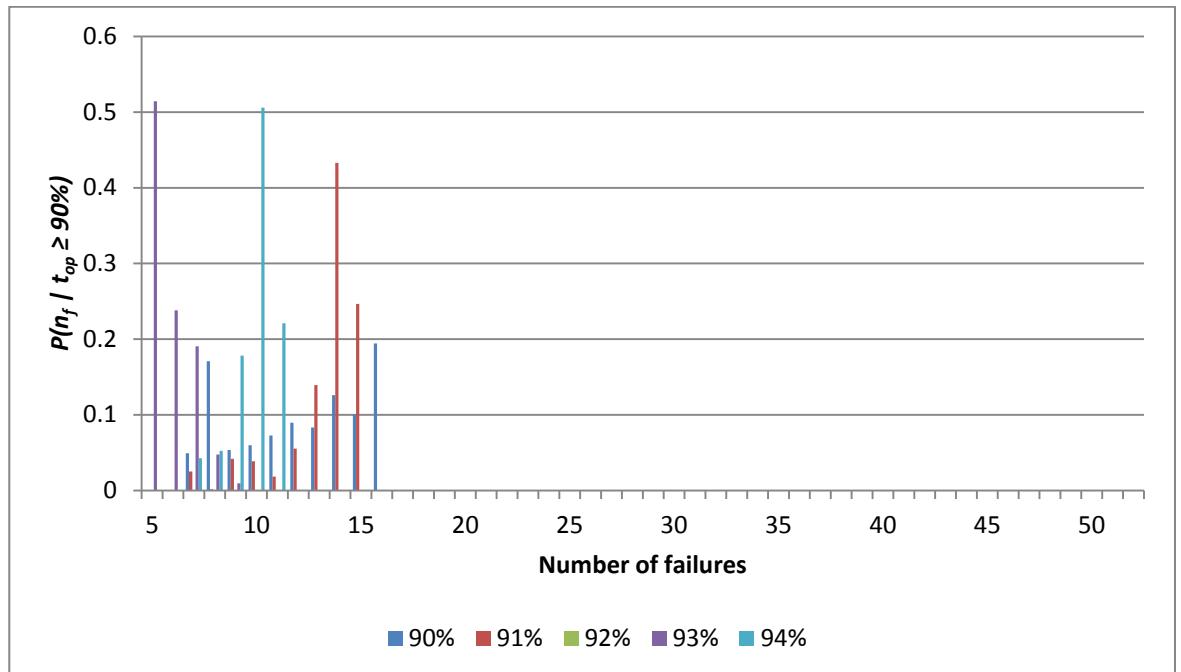


Figure 62 Histogram showing conditional probability distributions for $P(n_f / \text{top} \geq 90\%)$

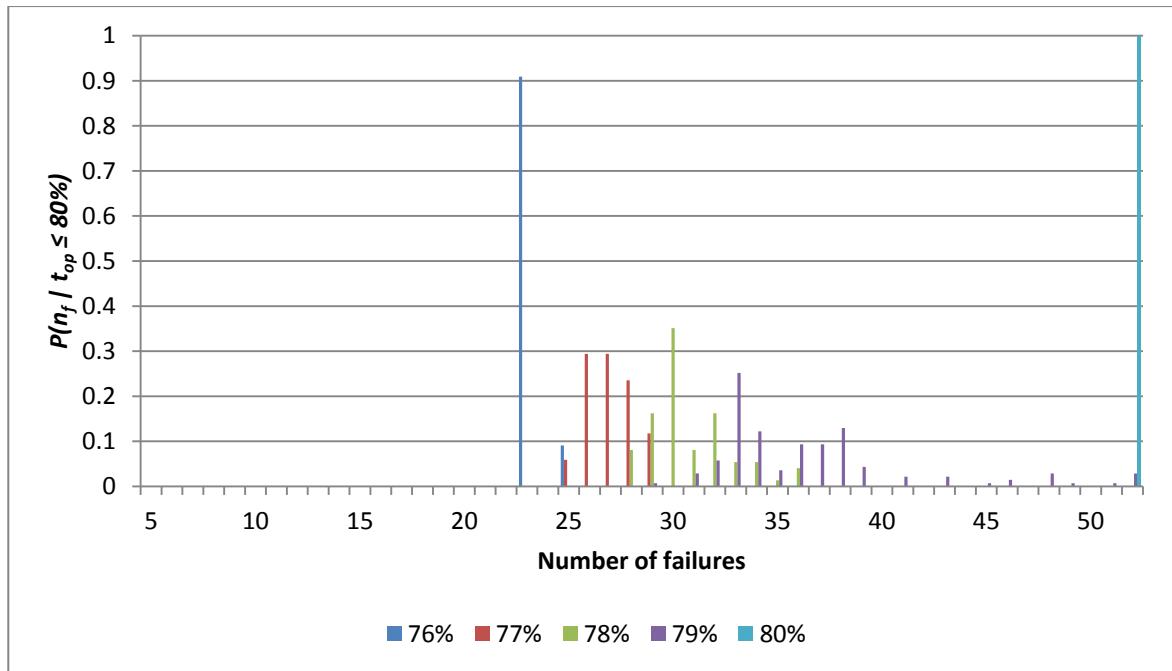


Figure 63 Histogram showing conditional probability distributions for $P(n_f / t_{op} \leq 80\%)$

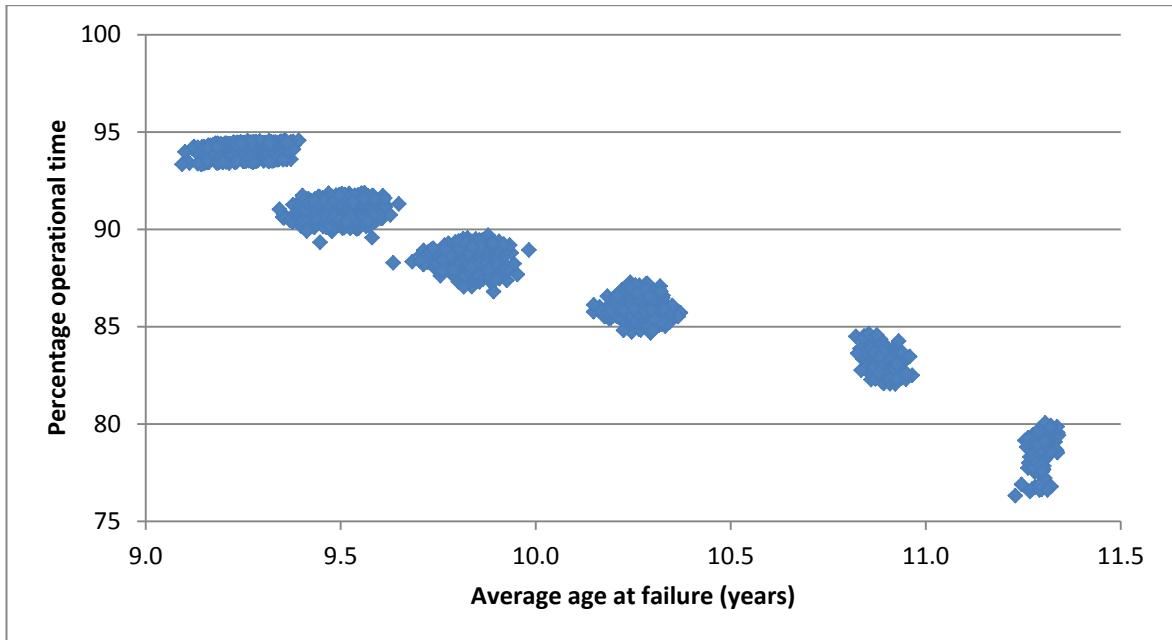


Figure 64 a_f versus t_{op} for scenario 1

As described above, a traditional constellation of monolithic spacecraft can reliably provide measurements for very long duration missions when optimising for operational time only. Given that this type of architecture is currently in use (for example the Disaster Monitoring Constellation [9]), it does not make sense to move to a fractionated type system if the designer is concerned

with operational time alone. This high operational time comes at a cost of increased mass launched. Figure 65 shows that 70% of architectures evaluated experienced a launched mass of 1500kg over their lifetime. Figure 66 shows that architectures with operational times greater than 90% were at least 3 times more likely to experience launched masses greater than 1500kg than less than this value. The 93% case is an anomaly caused by the relatively few architectures that experienced this operational time (Figure 50). Figure 67 shows that low operational time architectures exclusively experienced launched masses of less than 1600kg.

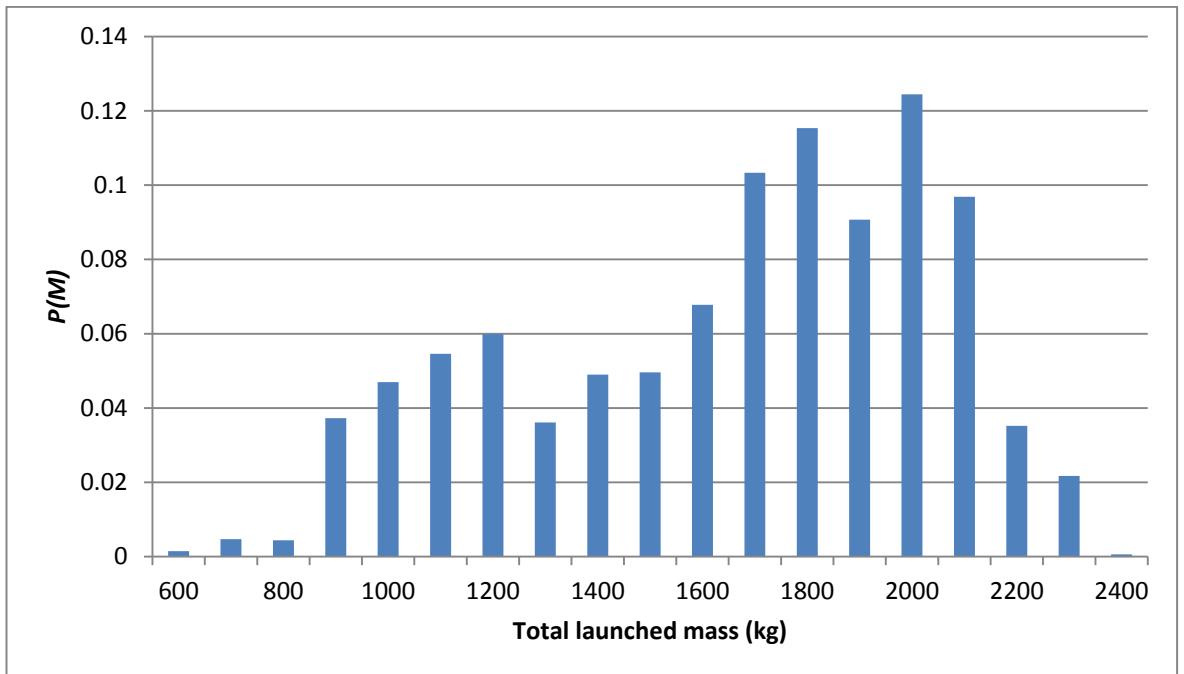
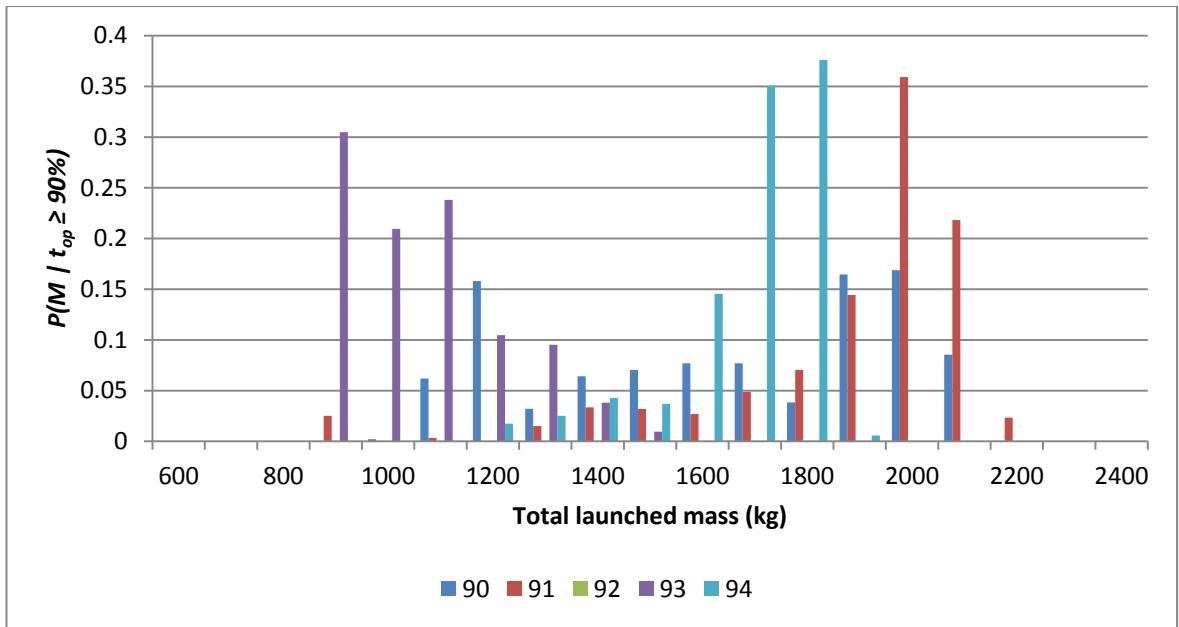
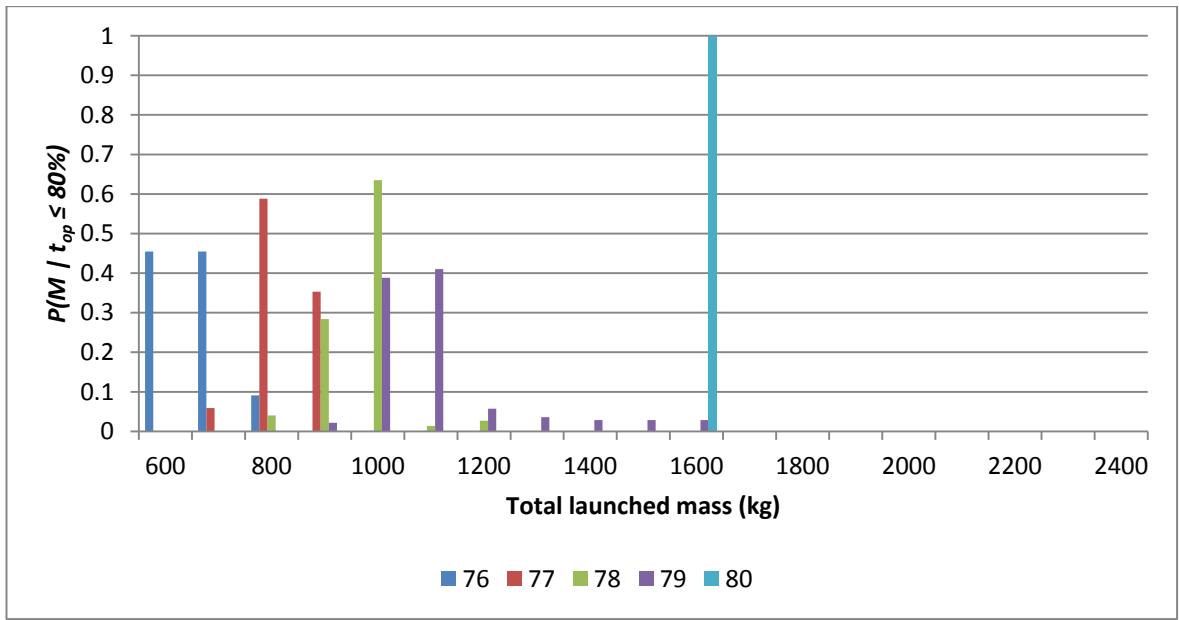


Figure 65 Probability distribution for $P(M)$ for scenario 1

Figure 66 Histogram showing conditional probability distributions for $P(M | \text{top} \geq 90\%)$ Figure 67 Histogram showing conditional probability distributions for $P(M | \text{top} \geq 90\%)$

This scenario optimised fractionated architectures for operational time only. In the next scenario, the complexity was increased by adding the mass into the optimisation parameter.

4.1.2. Scenario 2

The results for scenario 1 showed that achieving a high operational time was possible with a system that could be implemented with technology available today, using a constellation type system. However, this came at the cost of an increased total launched mass. To account for this trend the total launched mass and operational time were combined into a single fitness parameter as described in 3.4.3.

As well as modifying the fitness parameter, the failure model was also modified. The second failure curve, a 2-mixed Weibull function (Figure 68) which was described in 3.4.2.1 was implemented in this scenario.

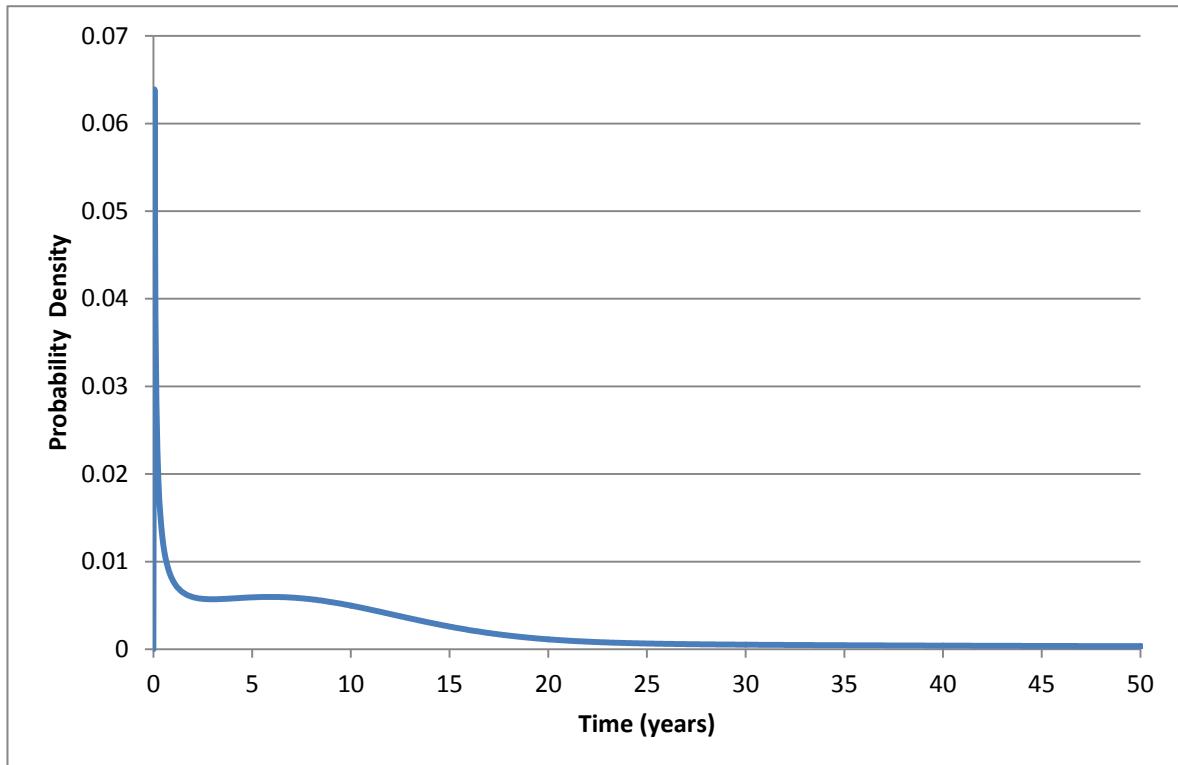


Figure 68 2-Mixed Weibull probability density function

4.1.2.1. Results and data analysis

Figure 69 gives the probability distribution for scenario 2 and shows that 57% of the architectures had fitness values of less than 1.001, while only 5.7% had fitness values of greater than 1.002. This is a shift in where the majority of architectures lie with respect to the optimisation parameter indicates that while many architectures result in high operational times, very few do so with low launched masses.

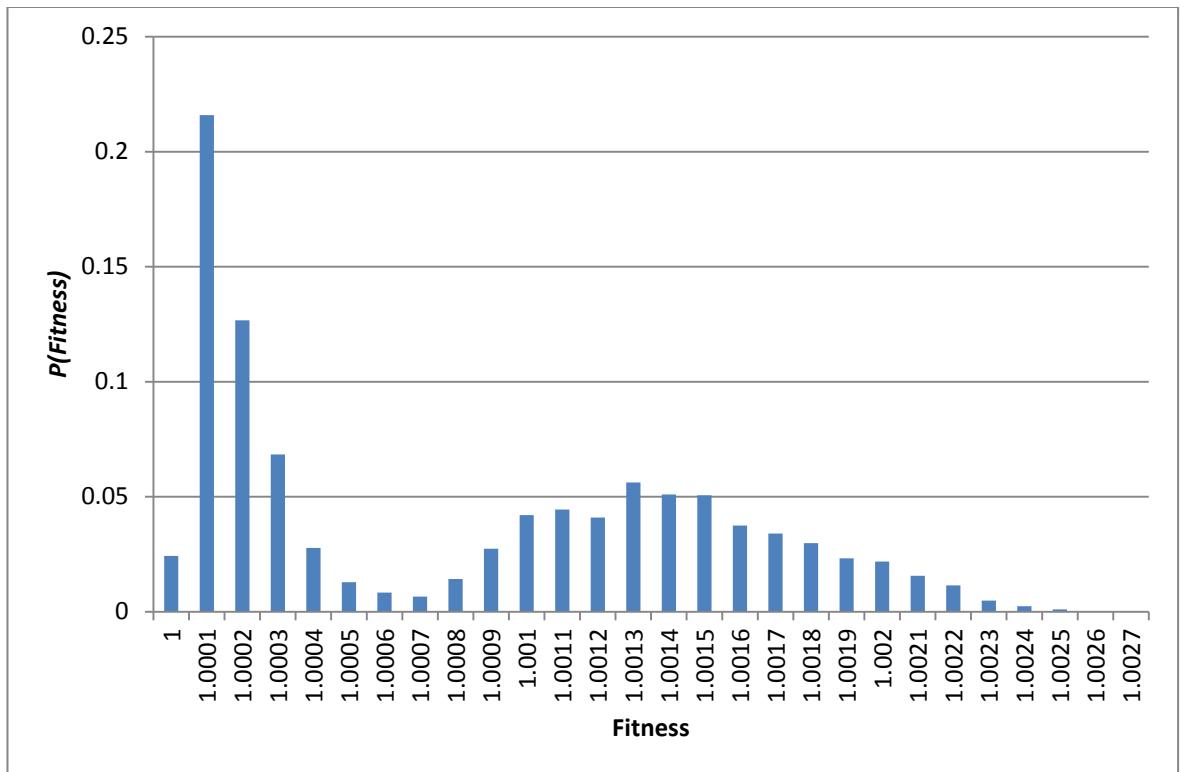
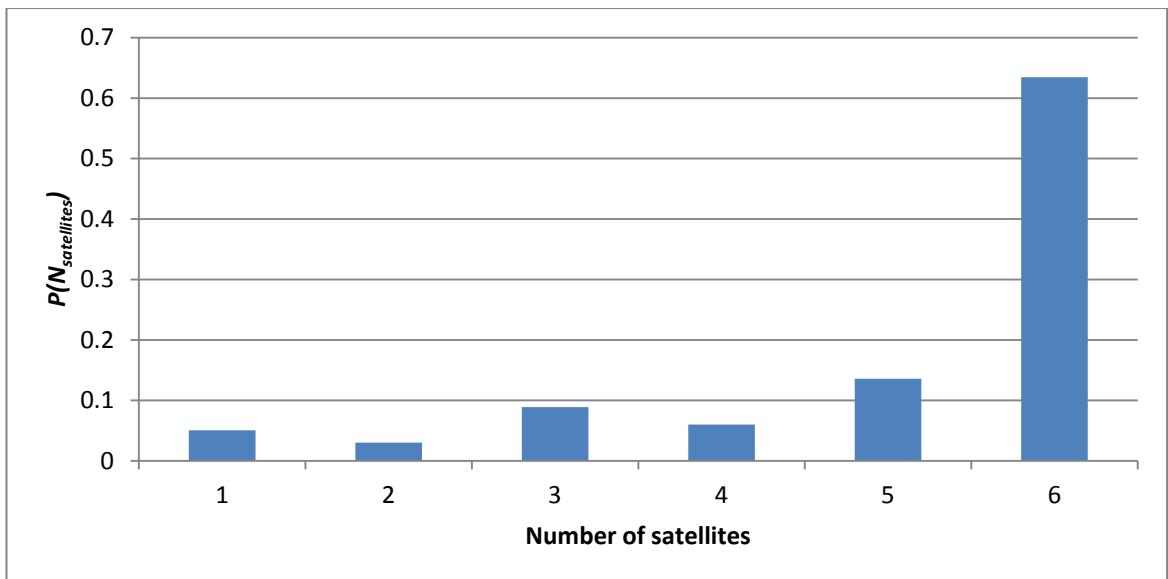
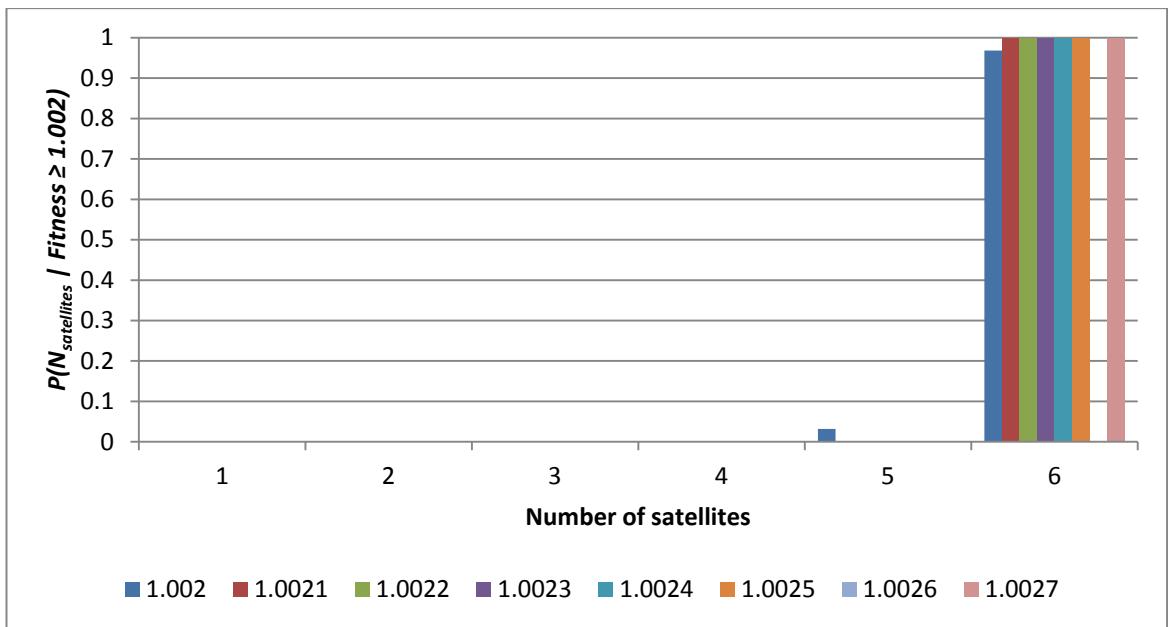


Figure 69 Probability distribution for $P(\text{Fitness})$ for scenario 2

Figure 70 gives the probability distribution for $P(N_{\text{satellites}})$ and shows that architectures containing six satellites were at least six times more likely to occur in the dataset than other numbers of satellites. Figure 71 shows that the highest fitness architectures ($\text{fitness} \geq 1.002$) were almost always produced by architectures with six satellites in them. Figure 72 shows that the lowest fitness architectures could have contained any number of satellites, although the most likely was six satellites in accordance with Figure 70.

Figure 70 Probability distribution for $P(N_{\text{satellites}})$ for scenario 2Figure 71 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} | \text{Fitness} \geq 1.002)$

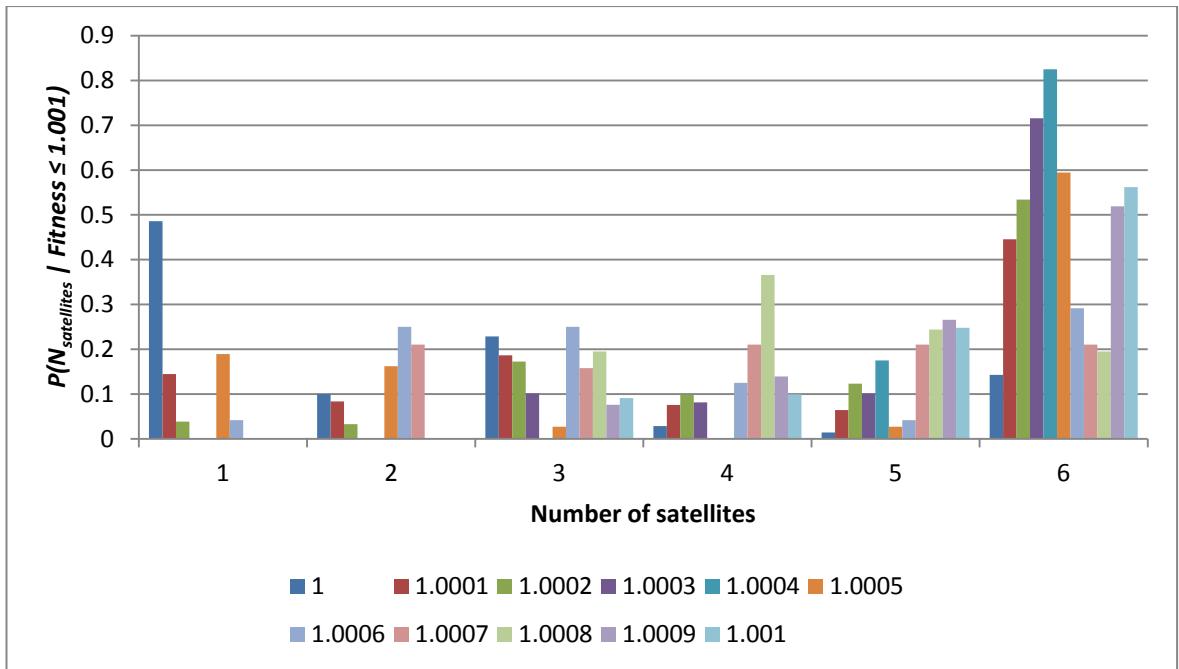


Figure 72 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} | \text{Fitness} \leq 1.001)$

Figure 73 shows that 52% of architectures evaluated in this scenario had $N_{\text{subsystems}} \geq 5$. Figure 74 shows that the highest fitness architectures were most likely to have $N_{\text{subsystems}} \geq 5.5$, while Figure 75 shows that the lowest fitness architectures experienced the full range of values for $N_{\text{subsystems}}$ but still with $N_{\text{subsystems}} \geq 5$ the most likely value.

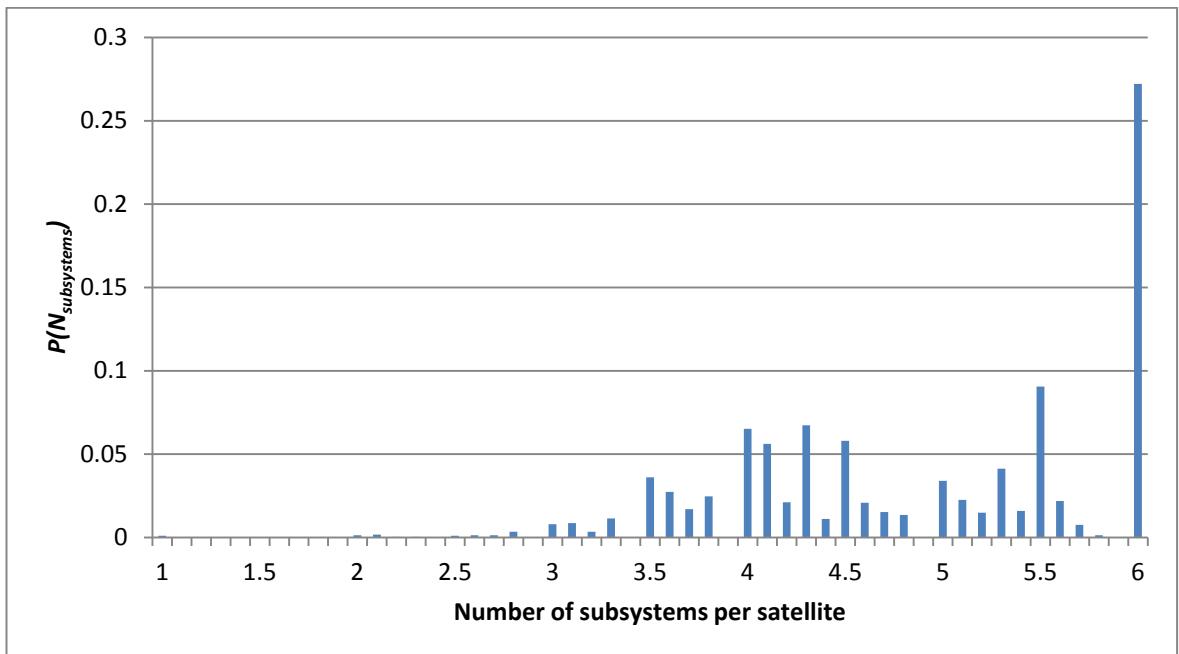


Figure 73 Probability distribution for $P(N_{\text{subsystems}})$ for scenario 2

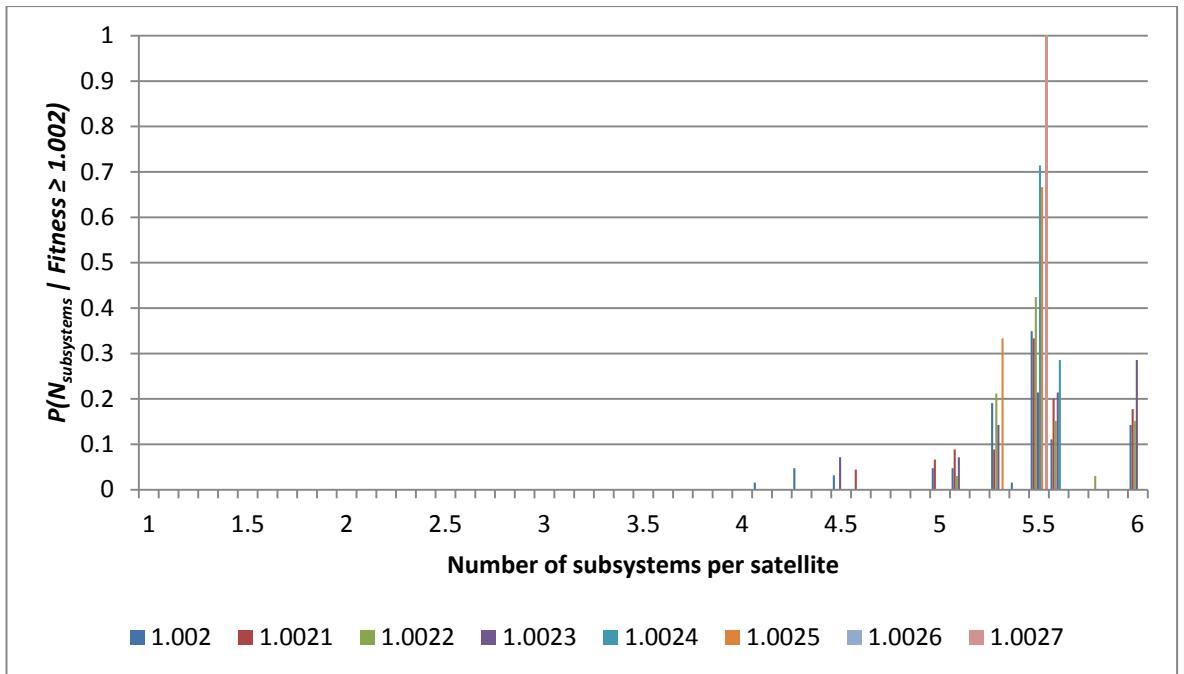


Figure 74 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} | \text{Fitness} \geq 1.002)$

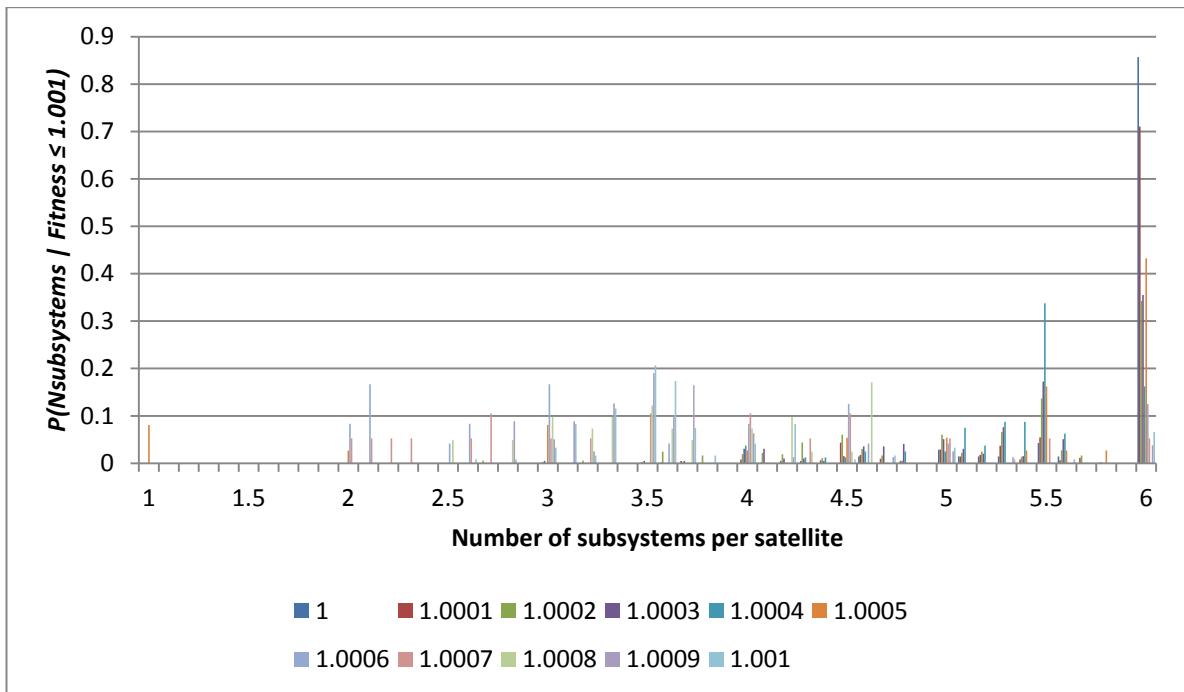


Figure 75 Histogram showing conditional probability distributions for $P(N_{\text{satellites}} | \text{Fitness} \leq 1.001)$

Figure 76 shows that fitness values of greater than 1.0005 were only observed in architectures with $D = 6$. 53% of architectures explored had $D = 6$.

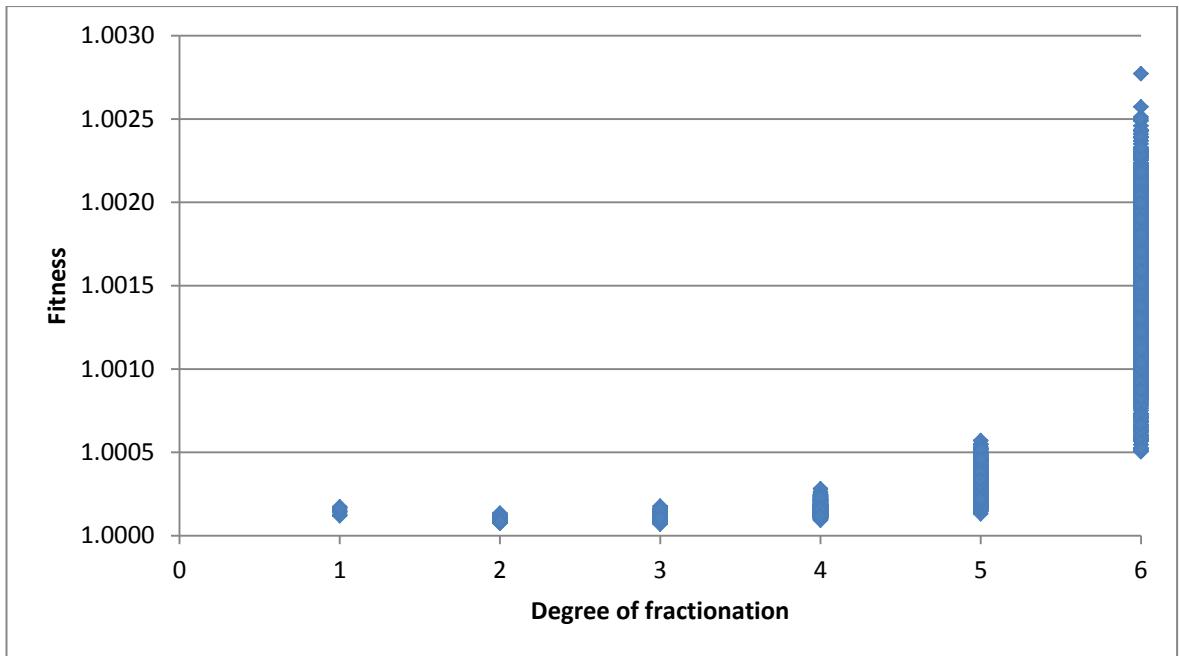
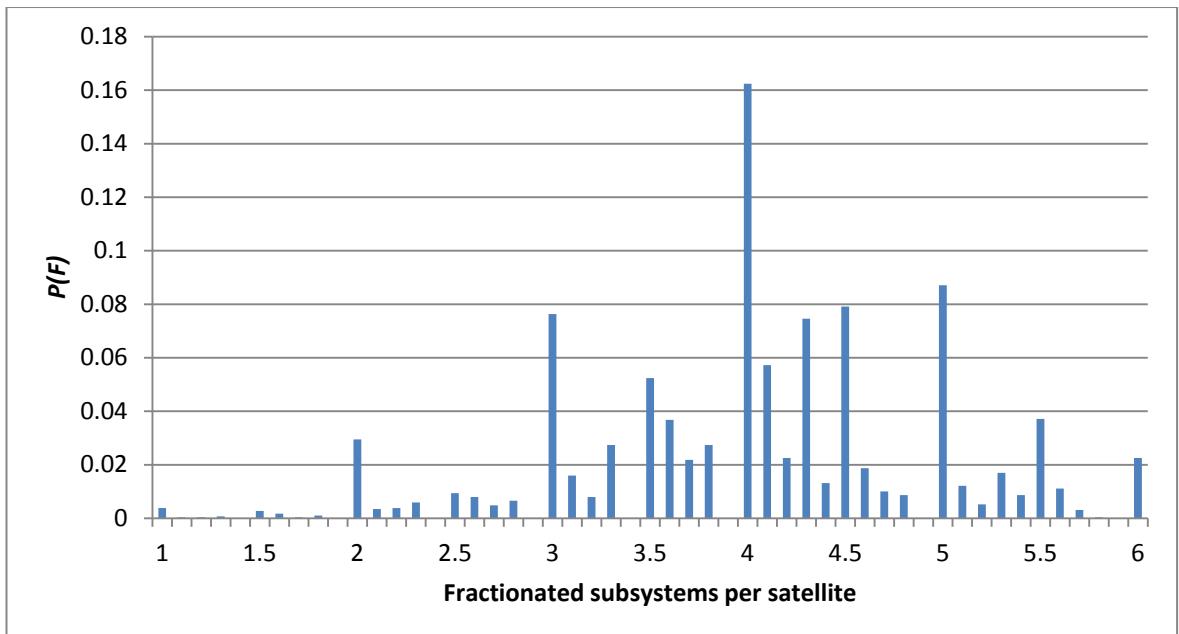


Figure 76 D versus Fitness for scenario 2

Figure 77 shows that 65% of architectures had $F \geq 4$. Figure 78 shows that architectures with high fitness values had a likelihood of having a value of $F \geq 5$ of up to 90%, whereas architectures with low fitnesses could have almost any possible value of F (Figure 79).

Figure 77 Probability distribution for $P(F)$ for scenario 2

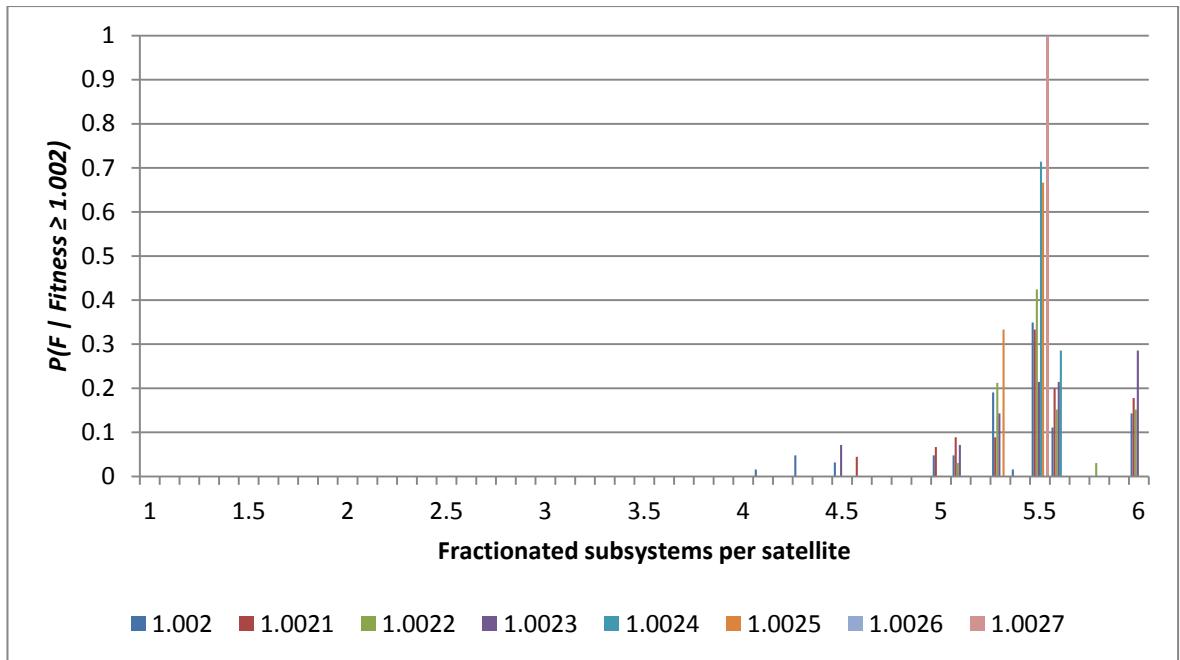


Figure 78 Histogram showing conditional probability distributions for $P(F | \text{Fitness} \geq 1.002)$

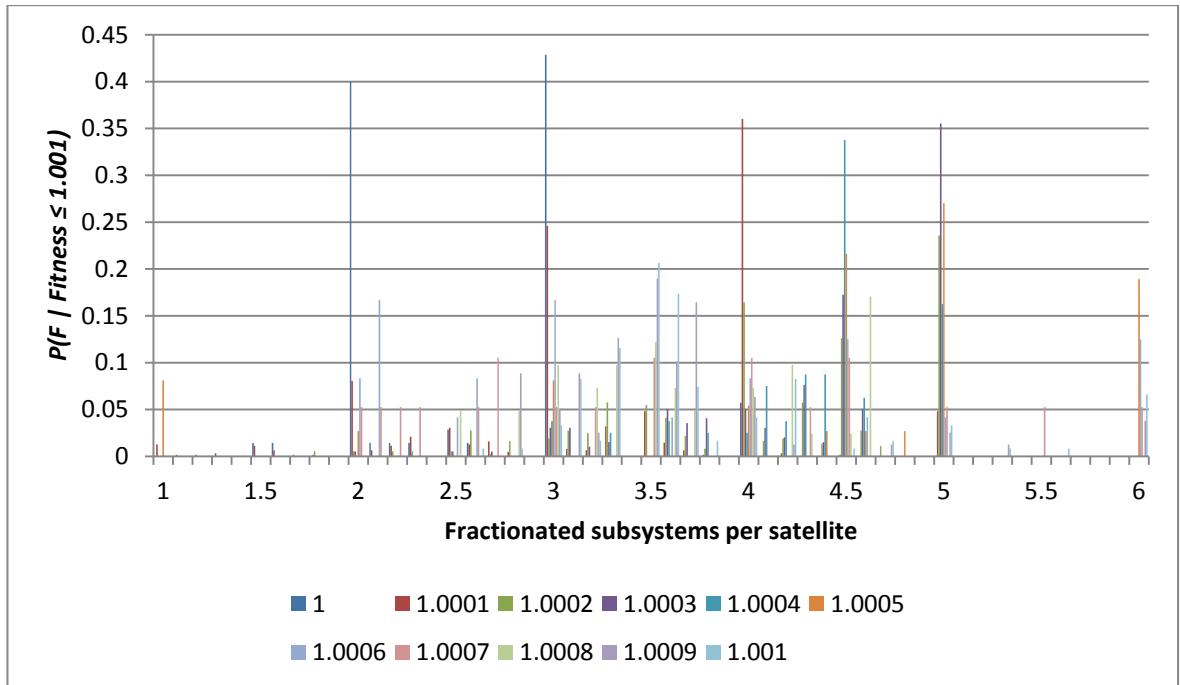


Figure 79 Histogram showing conditional probability distributions for $P(F | \text{Fitness} \leq 1.001)$

To support the probability analysis and to understand what conditions produced the highest fitnesses, an analysis was performed which looked at the top architecture from each of the fifteen separate model runs that made up the data set for this scenario. Fourteen of these top

architectures had six satellites, with the other one having five. All the architectures had $D = 6$ and Figure 80 and Figure 81 show that fitness is directly proportional to both F and E . These three results show that highly fractionated, and highly externally redundant architectures resulted in the highest fitness values. Figure 82 shows that E is directly proportional to the difference between n_f and n_{fs} , meaning that the robustness of an architecture to failures increases with external redundancy.

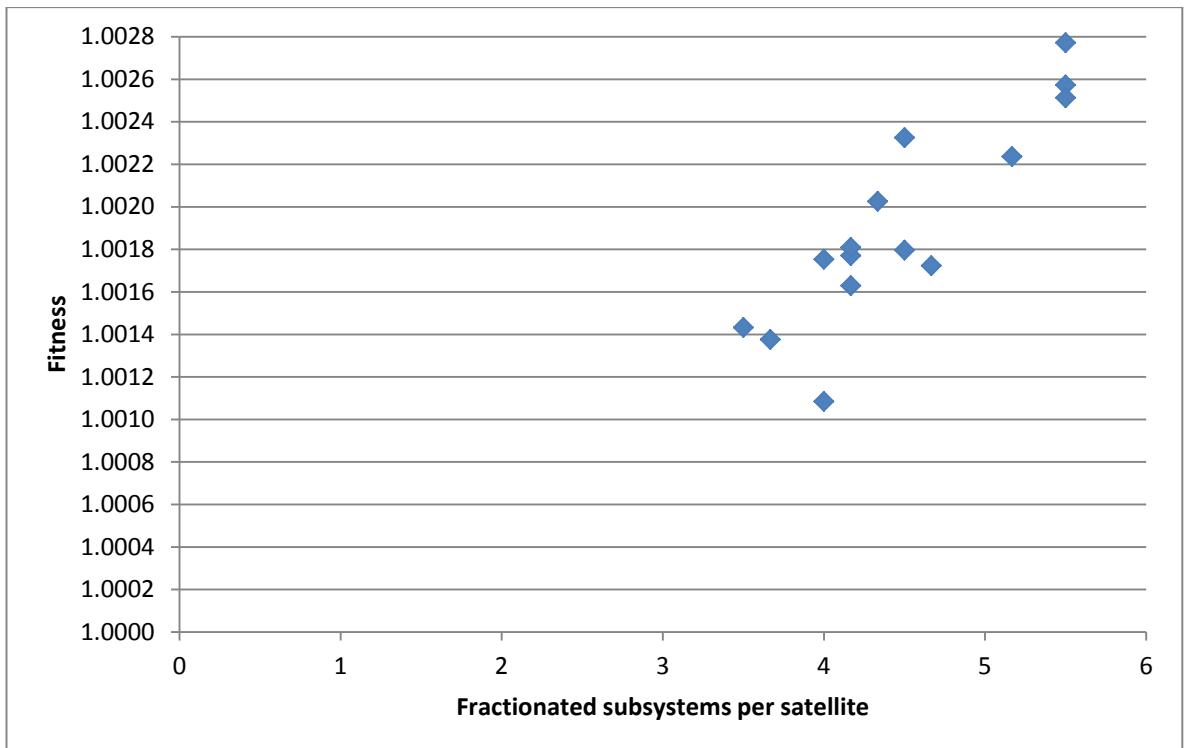


Figure 80 F versus fitness for the top architecture found by each model run in scenario 2

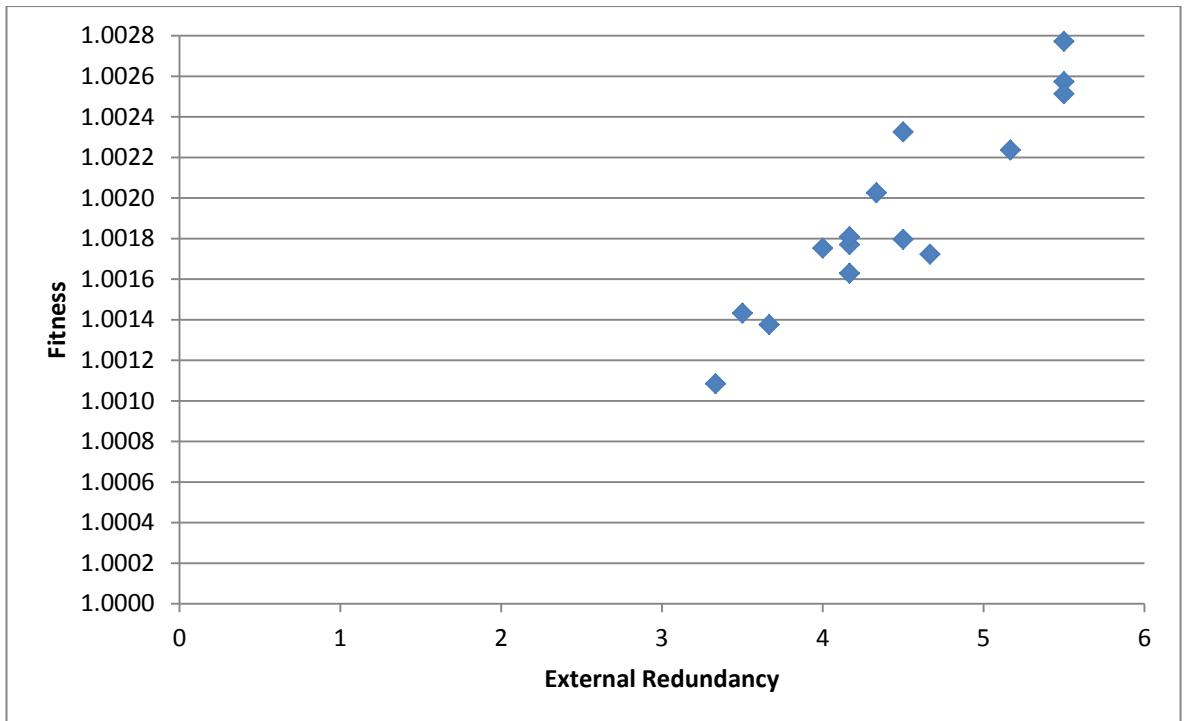


Figure 81 E versus fitness for the top architecture found by each model run in scenario 2

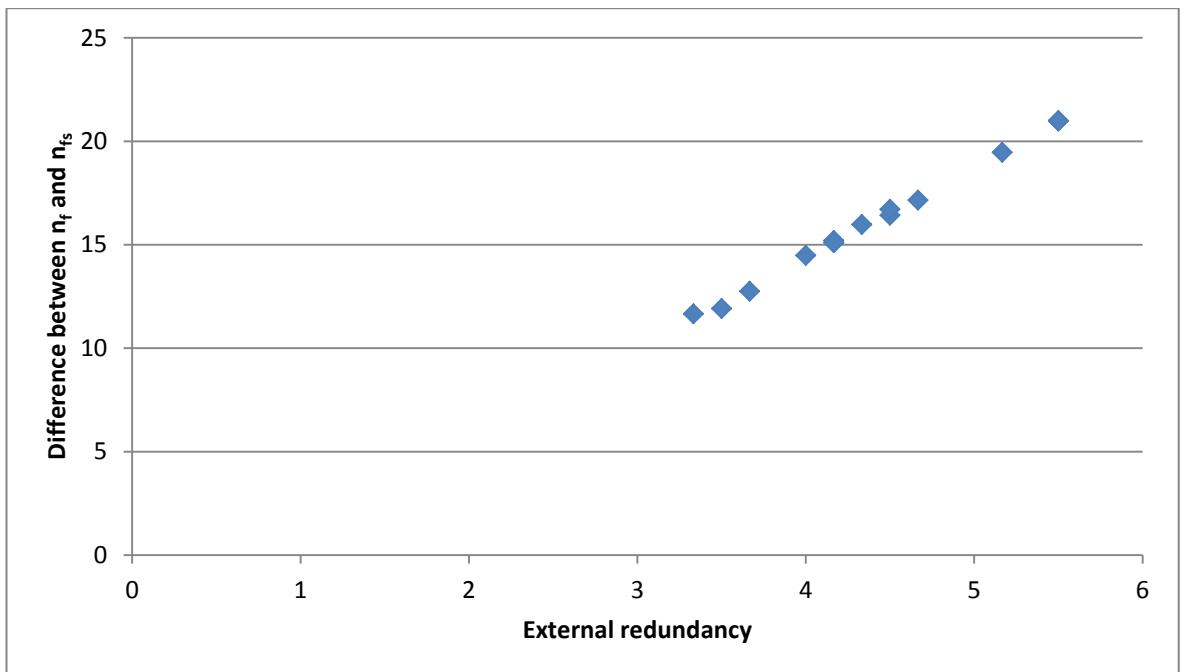


Figure 82 E versus the difference between n_f and n_{fs} for the top architecture from each model run in scenario 2

Figure 83 shows that approximately 13% of architectures evaluated in this dataset experienced 20 or more failures. Figure 84 and Figure 85 show that the architectures with highest fitness values were likely to experience 20 or more failures, while the lowest fitness architectures were likely to

experiences fewer failures. This means that, as shown above in Figure 82, that these high fitness architectures must be the most robust to failure. Figure 86 shows that architectures with fitness greater than 1.005 experienced a value of $a_f > 20.05$ years, and accounted for 53% of the total number of architectures explored.

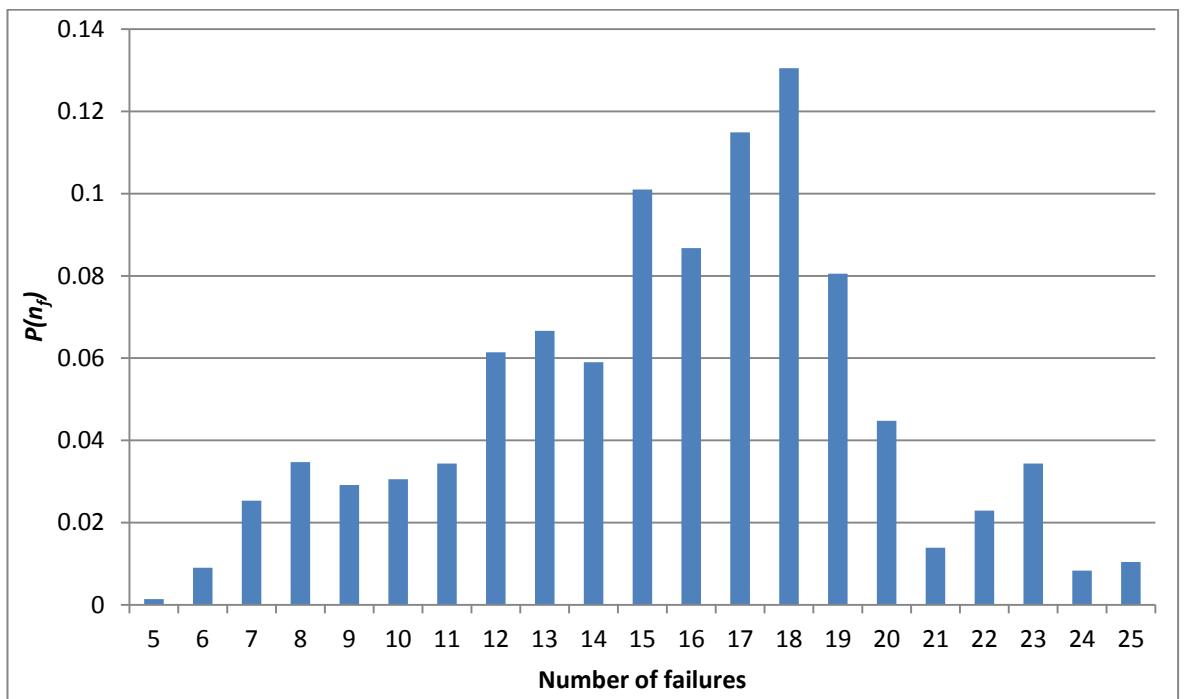


Figure 83 Probability distribution for $P(n_f)$ for scenario 2

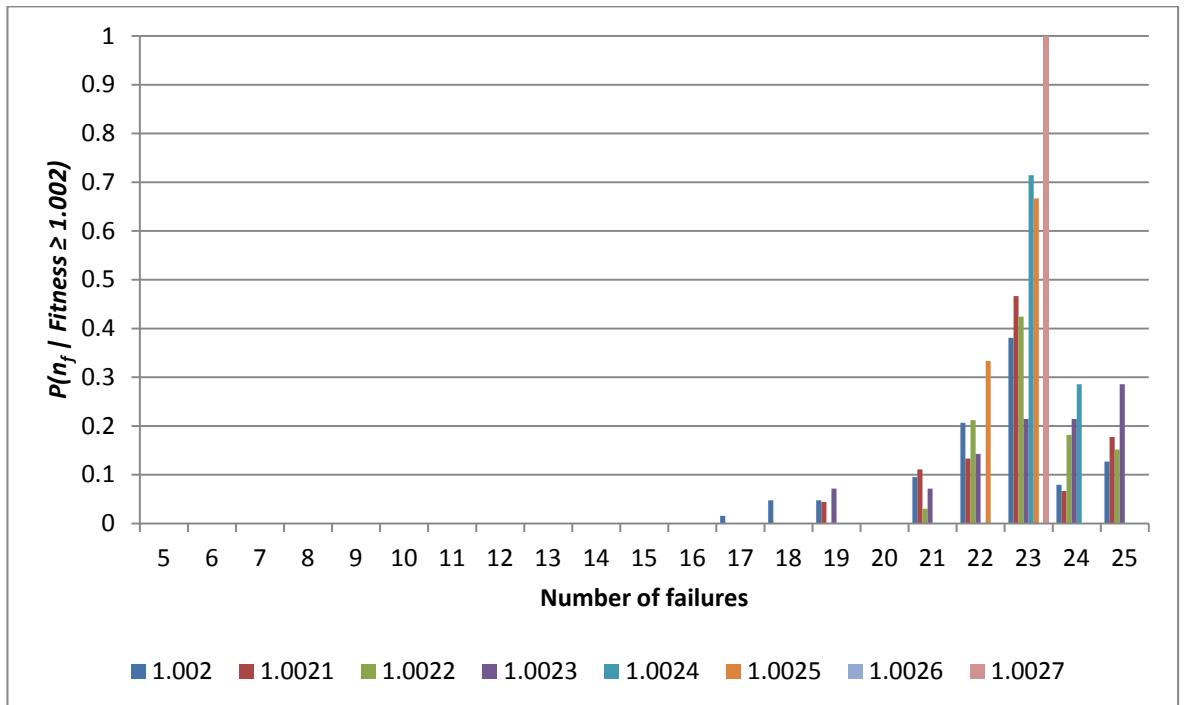


Figure 84 Histogram showing conditional probability distributions for $P(n_f / \text{Fitness} \geq 1.002)$

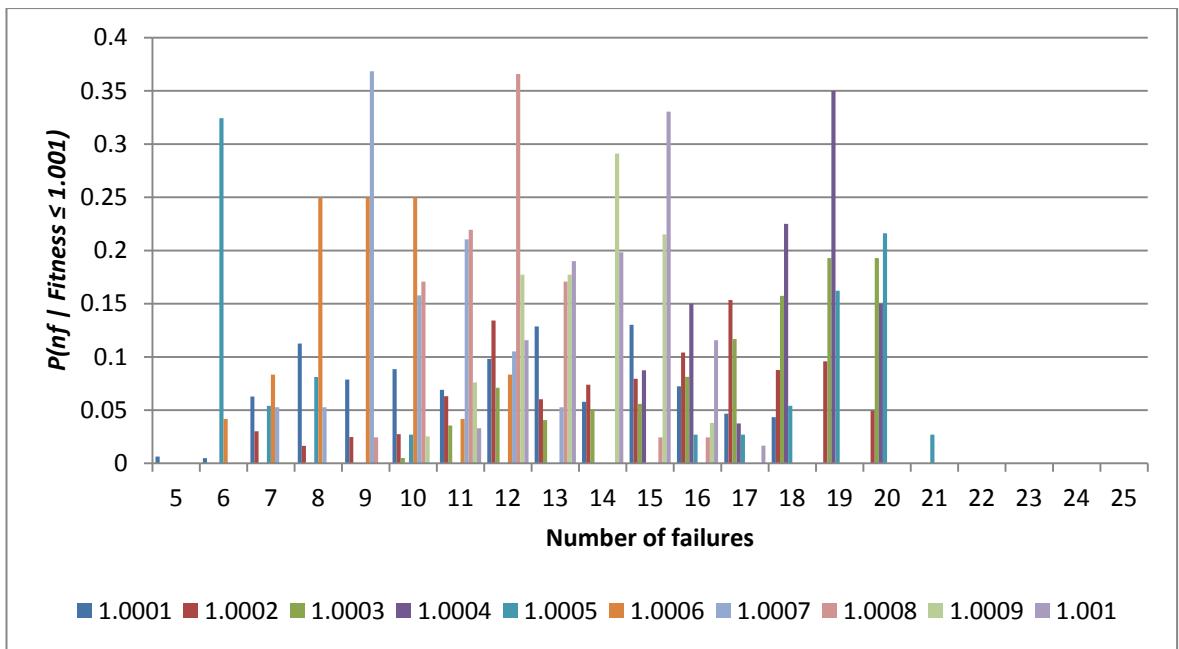


Figure 85 Histogram showing conditional probability distributions for $P(n_f / \text{Fitness} \leq 1.001)$

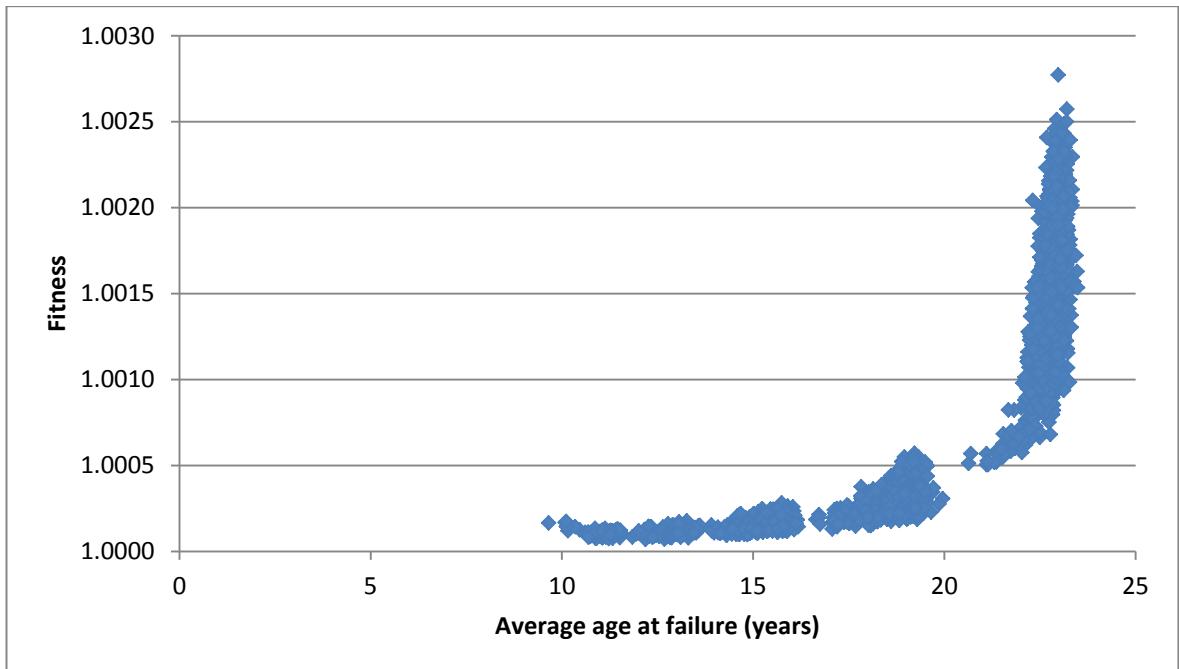


Figure 86 a_f versus fitness for scenario 1

4.1.3. Summary

The results presented in section 4.1.1 showed that 50% of architectures explored maximise operational time. The results presented in section 4.1.2 showed that only about 6% of architectures explored can successfully maximise operational time while minimising the mass

launched, resulting in fitness values greater than 1.002. The distributions presented in section 4.1.2 provided some baseline expectations for the main scenarios explored in the next section for which the model was run using multiple mixed Weibull failure curves.

4.2. Results

In the data analysis it was assumed that all the subsystems used the same failure curves. Research by Castet and Saleh into satellite failures produced reliability curves for individual subsystems, similar to that used for scenario 2 [96]. However, unlike the curve in scenario 2, they only produced a single Weibull distribution, which only accounts for infant mortality. In order to account for the end of life failure as well, a 2-mixed Weibull distribution was used, including the shape and scale parameters for the infant mortality for each individual subsystem. The end of life parameters would be taken from the 2-mixed Weibull distribution used in scenario 2.

Castet and Saleh considered eleven spacecraft subsystem categories in their analysis, as follows [96]:

1. Gyro/Attitude Sensors/Reaction Wheels
2. Thruster/Fuel system
3. Beam/Antenna Operation/Deployment
4. Control Processor
5. Mechanisms/Structures/Thermal Control
6. Payload Instrument/Amplifier/On-board Data/Computer/Payload Transponder
7. Batteries
8. Electrical Distribution
9. Solar Array Deployment
10. Solar Array Operation Mechanism
11. Telemetry, Tracking and Command

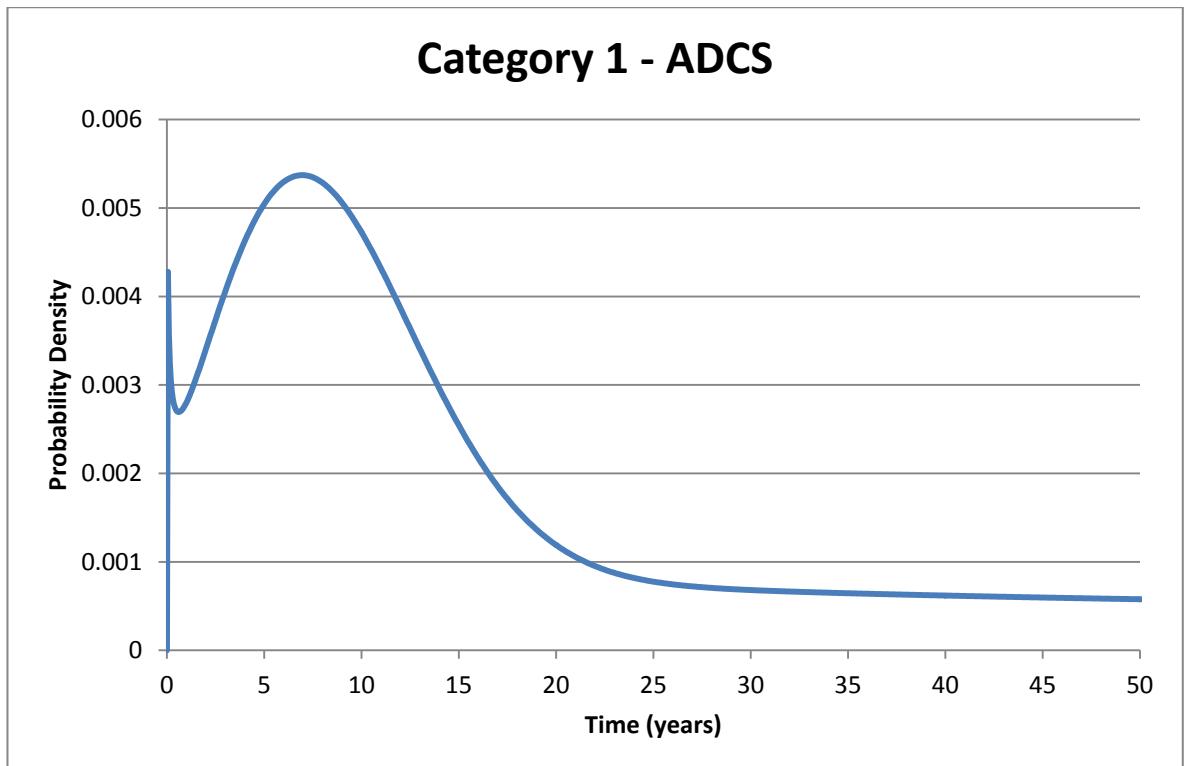
These categories were assigned to the six spacecraft subsystems modelled in this research, as given in Table 26. Category 3 was not assigned to the communications subsystem in this research, as it was assumed that this category is relevant to geostationary communications spacecraft only. Category 9 is also not assigned to the power subsystem because this is a single-use subsystem and does not have a reliability curve associated with it. For the subsystems with more than one assigned reliability curve, a single category is randomly selected at each time step to be used to check for a failure. If a subsystem were to be tested for failure more than once per time-step, it would increase the probability of this subsystem failing and lead to a bias in the failure model.

Subsystem	Payload	Communications	OBDH	ADCS	OCS	Power
Reliability curve category	6	6, 11	4, 6	1	2	7, 8, 10

Table 26 Spacecraft subsystems and associated reliability curve categories

Table 27 gives the shape and scale parameters for each of the reliability curves [96] used in this research and Figure 87 to Figure 94 show these curves.

Category	β	θ
1	0.7182	3831
2	0.3375	6206945
4	1.4560	408
6	0.8874	7983
7	0.7460	7733
8	0.5021	169272
10	0.4035	1965868
11	0.3939	400982

Table 27 Shape and scale parameters for the reliability curves used [96]**Figure 87 2-Mixed Weibull probability density function, Category 1**

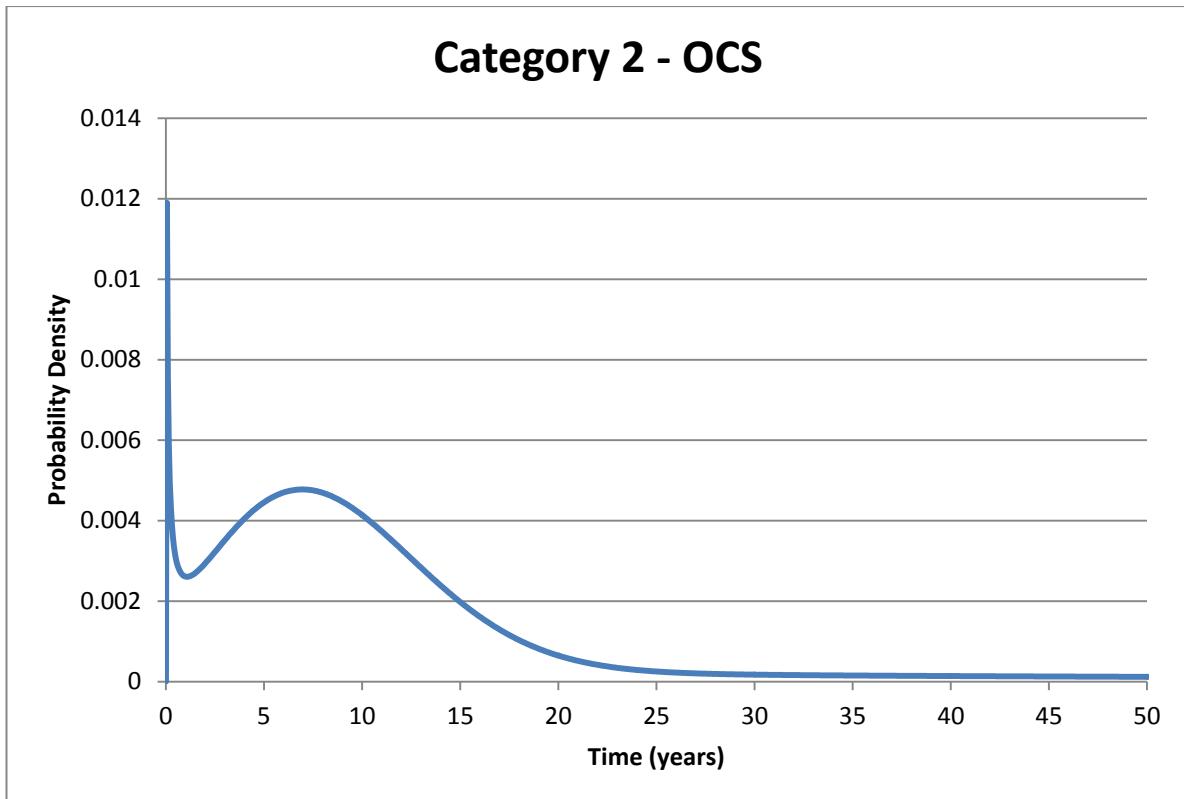


Figure 88 2-Mixed Weibull probability density function, Category 2

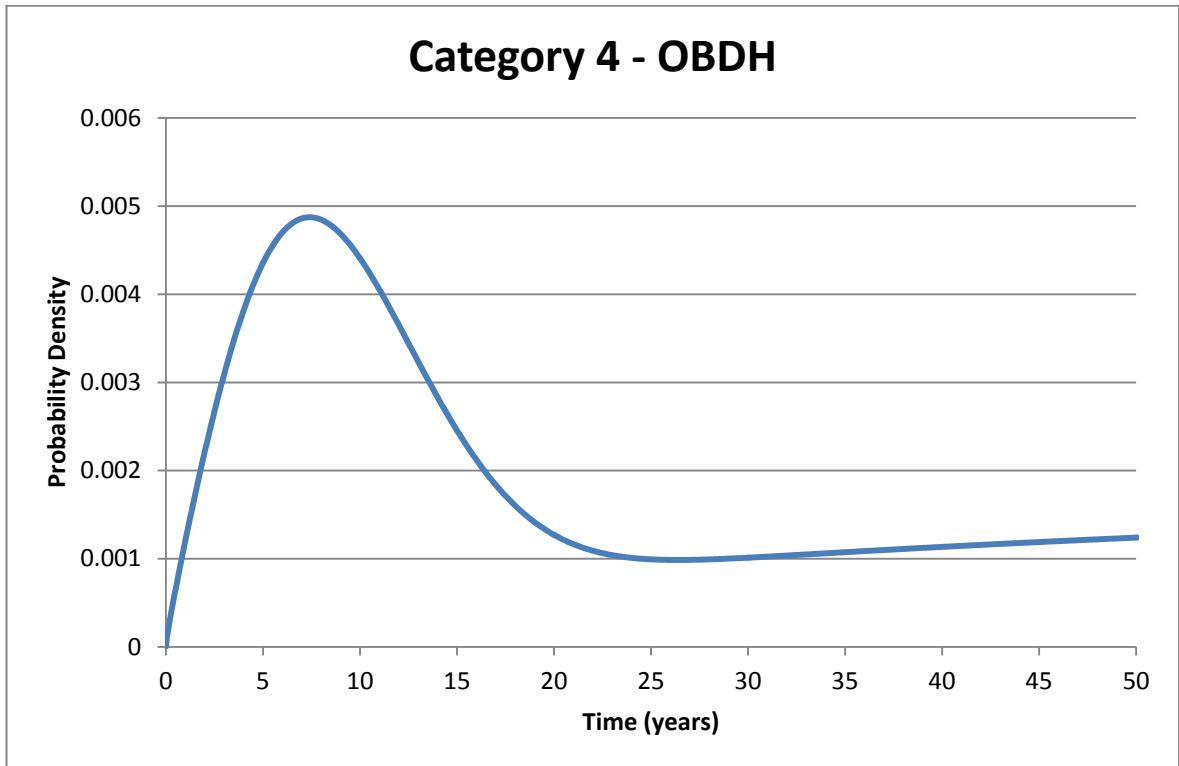


Figure 89 2-Mixed Weibull probability density function, Category 4

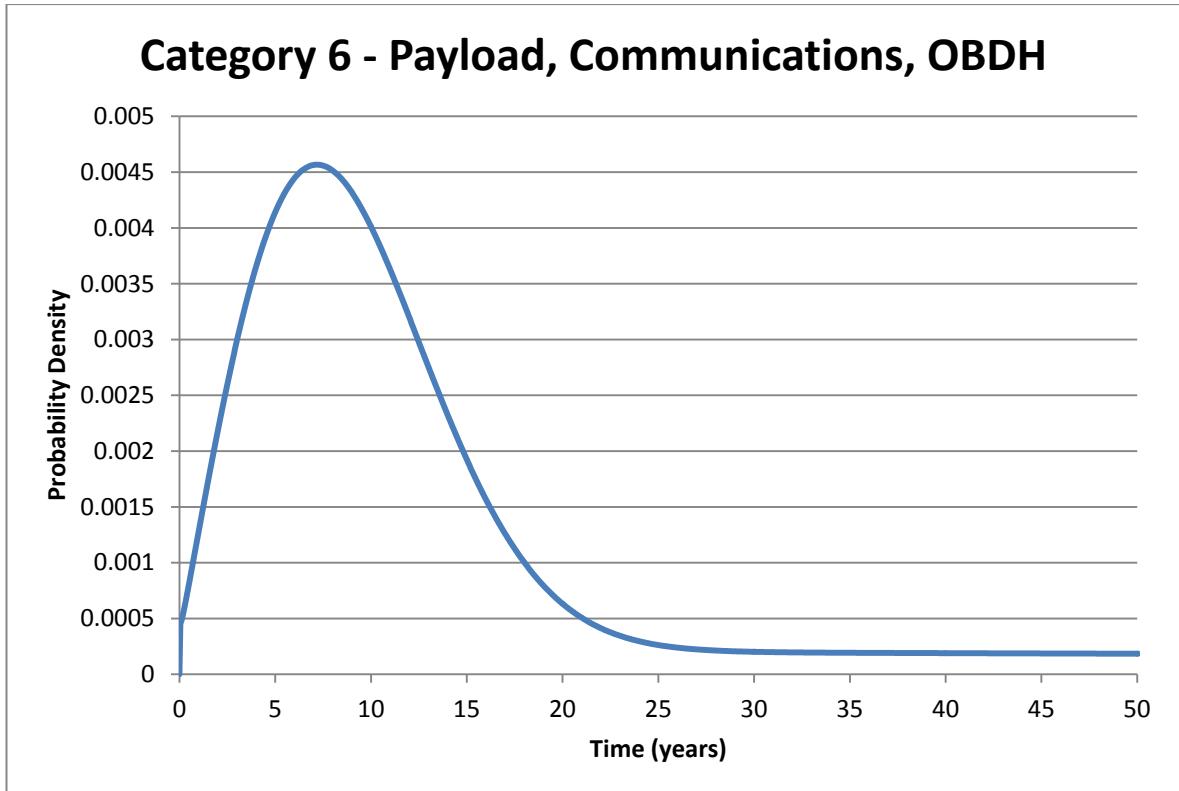


Figure 90 2-Mixed Weibull probability density function, Category 6

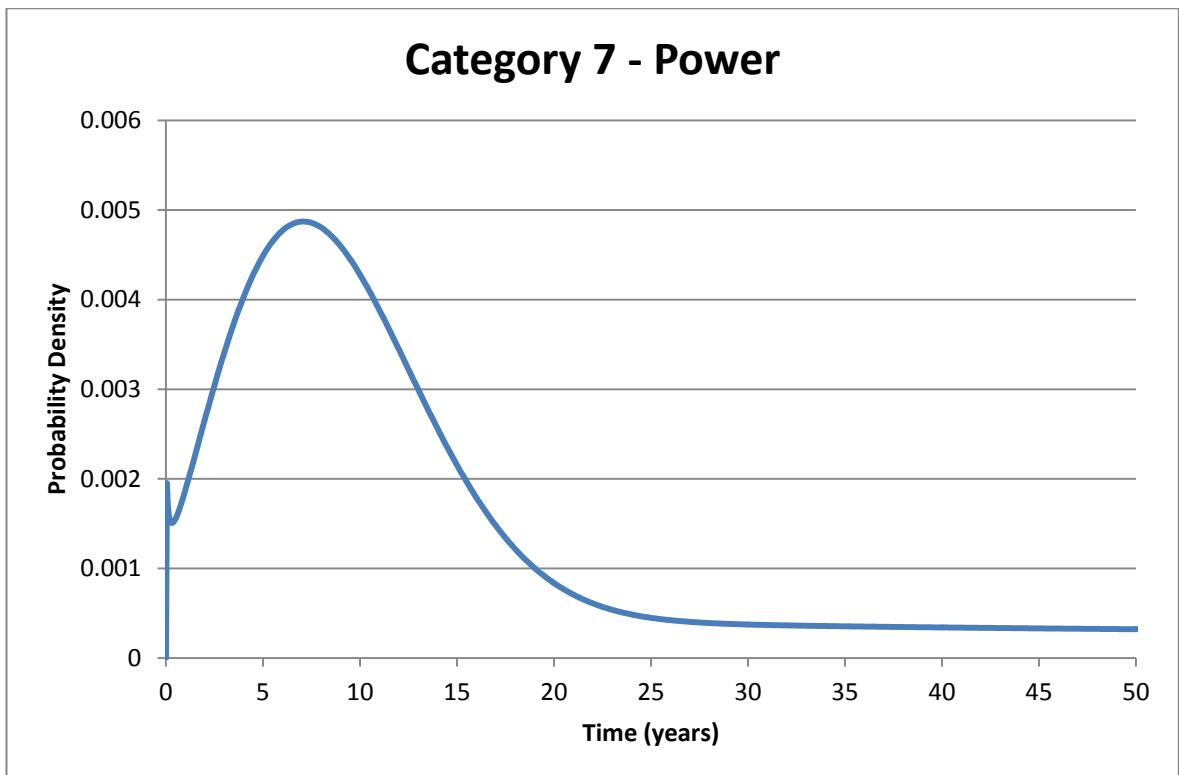


Figure 91 2-Mixed Weibull probability density function, Category 7

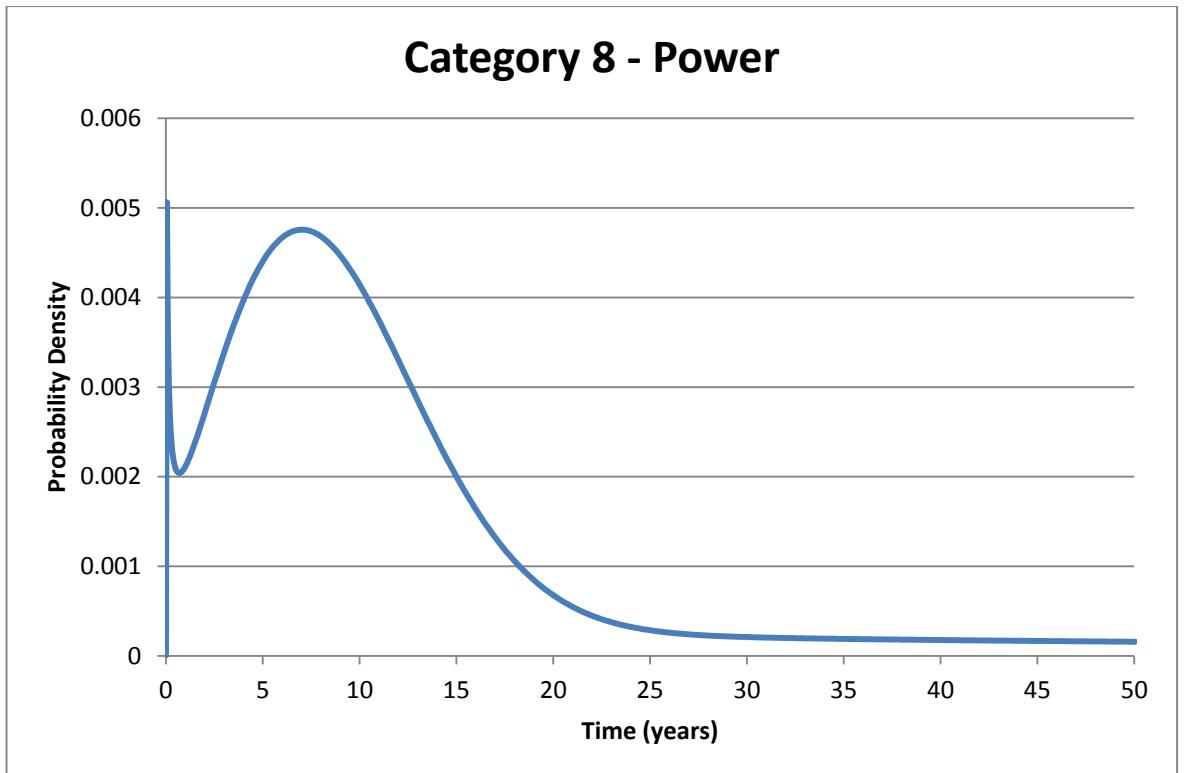


Figure 92 2-Mixed Weibull probability density function, Category 8

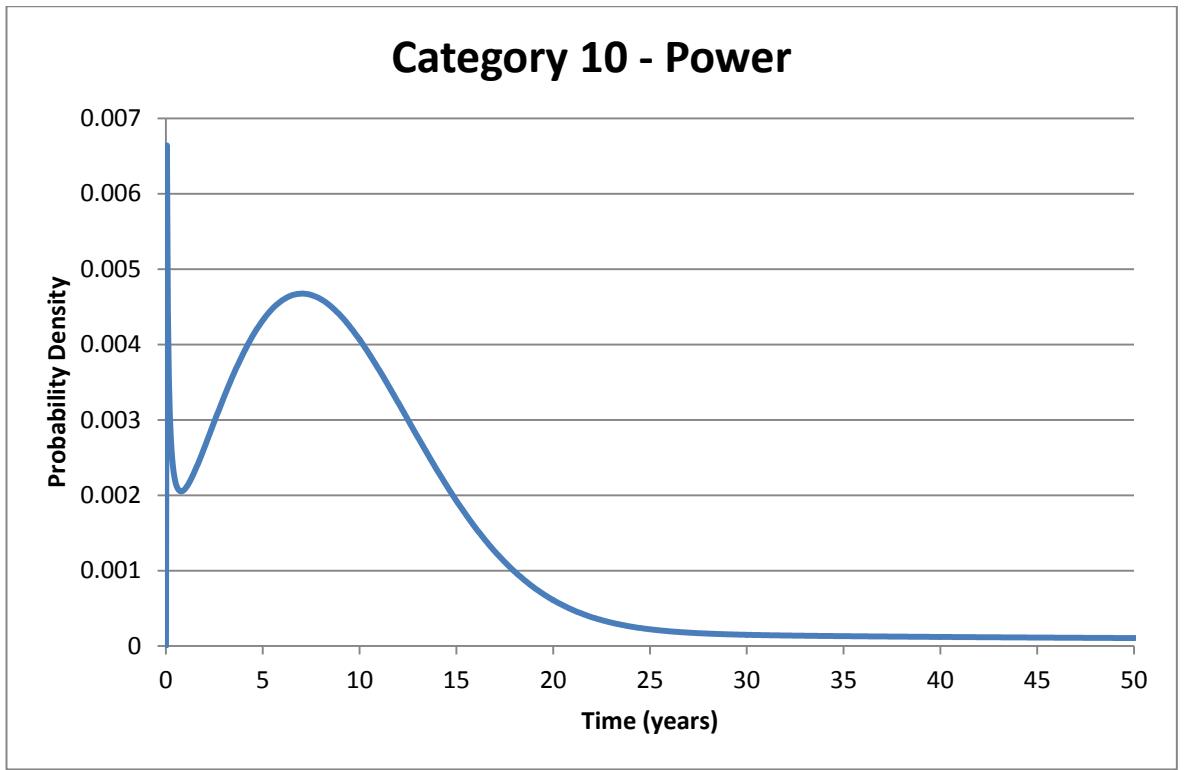


Figure 93 2-Mixed Weibull probability density function, Category 10

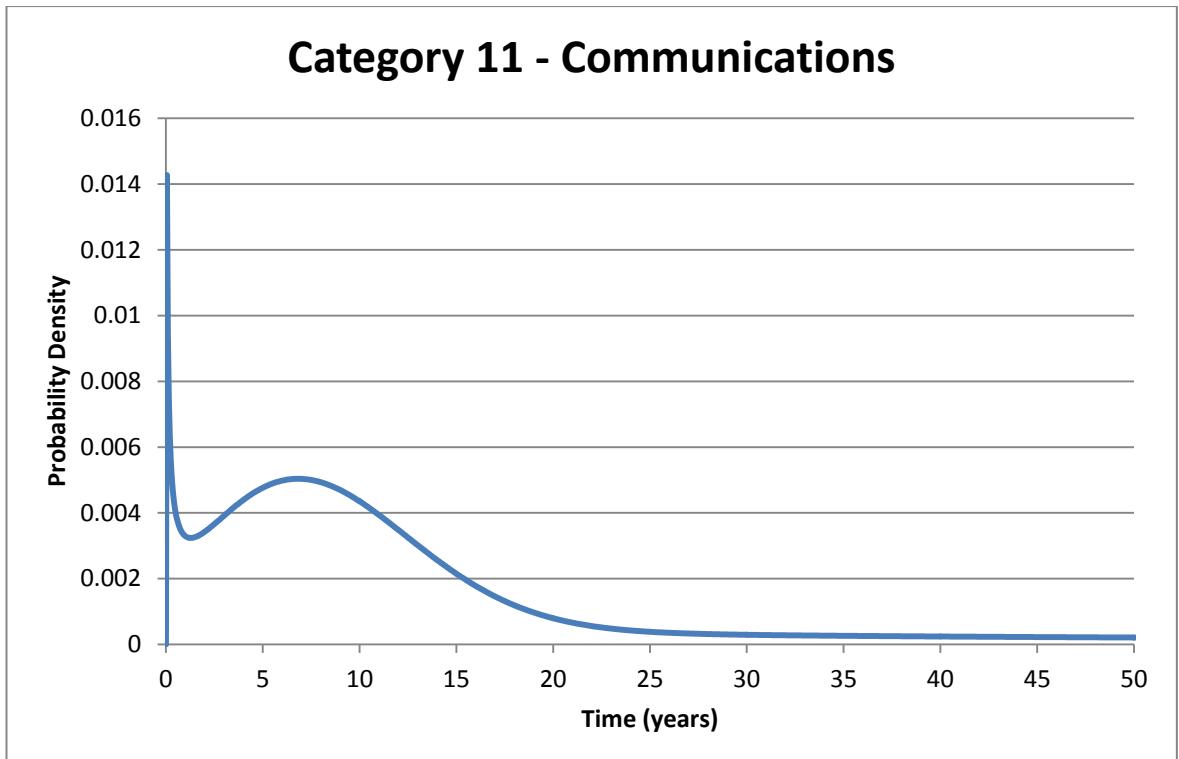


Figure 94 2-Mixed Weibull probability density function, Category 11

In order to determine the overall probability of failure for each curve, a trapezium rule was used to calculate the area under the curve in increments of one time-step. These were summed to give the overall probability of failure over the 50-year projection period. Table 28 shows the failure probability for each failure curve and the average failure probability for each subsystem if more than one curve was used.

Subsystem	Failure probability	Average Failure Probability
Payload	6.20%	-
Communications	6.20%	6.95%
	7.70%	
OBDH	9.50%	7.85%
	6.20%	
ADCS	9.30%	-
OCS	6.80%	-
Power	7.30%	6.80%
	6.70%	
	6.40%	

Table 28 Subsystem failure probabilities

4.2.1. Results and data analysis

This scenario is the main and final iteration of the model. In addition to the nominal replacement period of 0.7 years, two other replacement periods of 0.05 years and 1.5 years were investigated to understand the effect of very rapid response to a loss of capability and a replacement time closer to what is achievable at present.

4.2.1.1. 0.7 Year replacement time

Figure 95 gives the probability distribution for $P(Fitness)$ for this scenario and shows that 24% of architectures had the highest fitness values (greater than 1.0006). The bi-modal nature of this graph indicates that there is no “middle ground” for this dataset, 60% of architectures are effective at achieving the compromise between operational time and mass launched (greater than fitness ≥ 1.0004), and 40% were not.

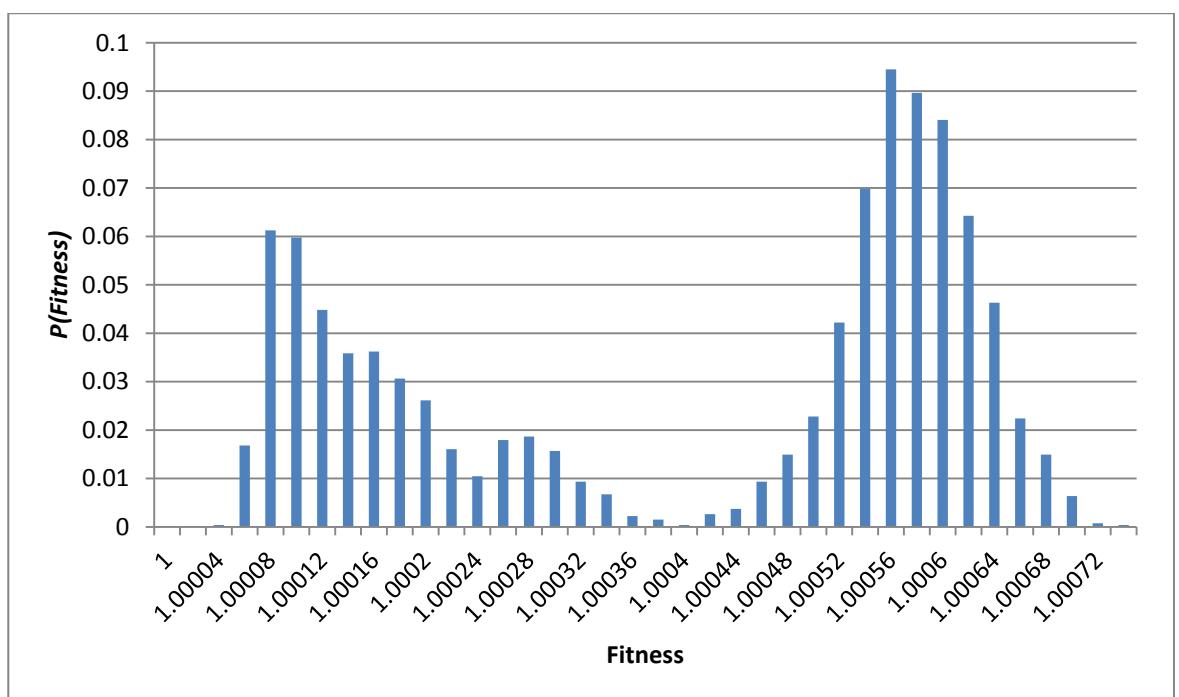
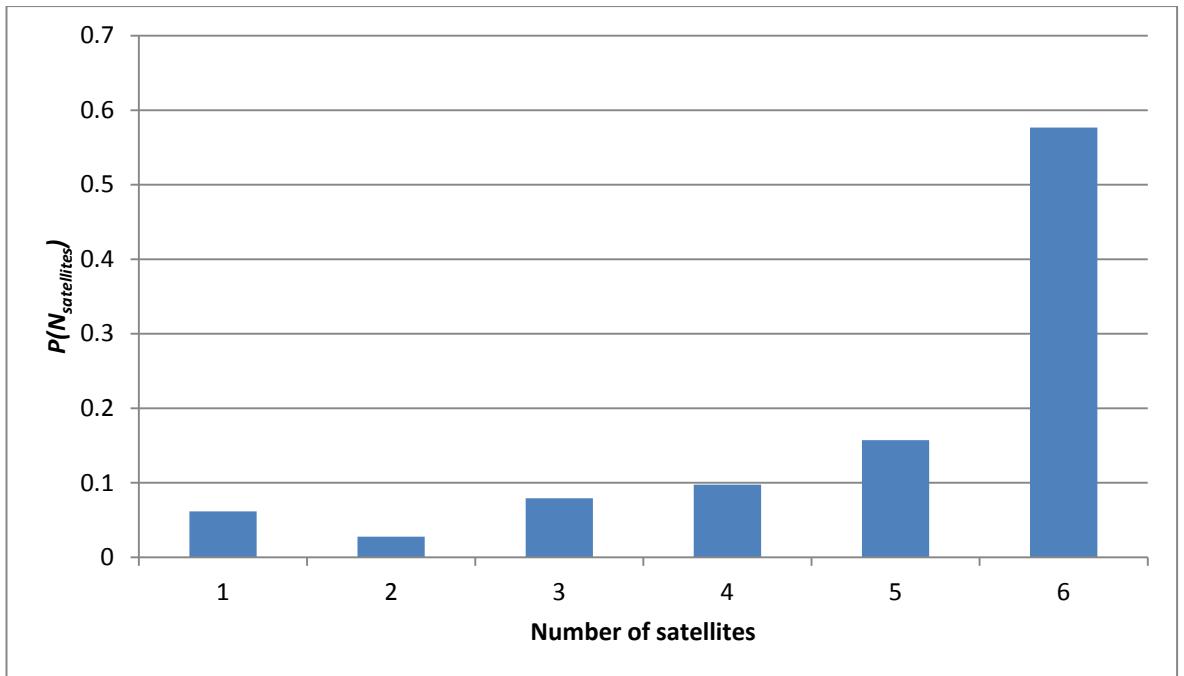
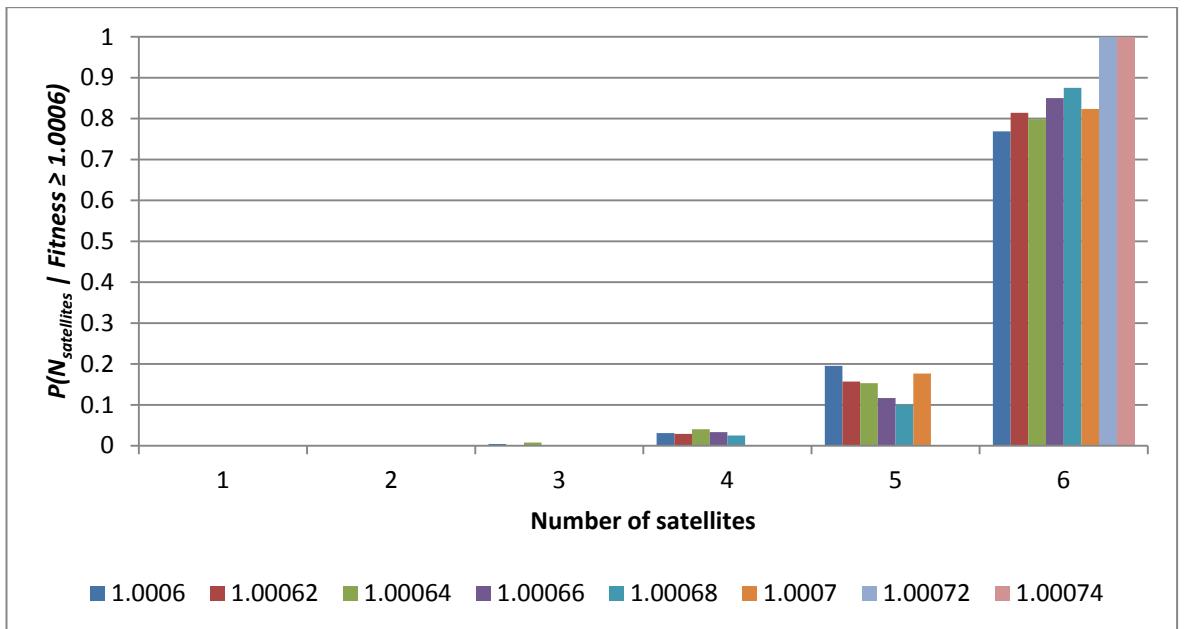


Figure 95 Probability distribution for $P(Fitness)$ for the 0.7 year replacement time scenario

Figure 96 shows that 58% of architectures explored in this dataset had $N_{satellites} = 6$. Figure 97 shows that the highest fitness architectures were at least four times more likely to have $N_{satellites} = 6$ than other values of $N_{satellites}$. The $P(N_{satellites} | Fitness \geq 1.0006) = 1$ values for fitness = 1.00072 and 1.00074 were the result of only one architecture achieving each of these fitness values. Figure 98 shows that the lowest fitness architectures (less than 1.0002) experienced all values of $N_{satellites}$ and that the most likely value of $N_{satellites}$ for the very lowest fitness architectures was $N_{satellites} = 1$.

Figure 96 Probability distribution for $P(N_{satellites})$ for the 0.7 year replacement time scenarioFigure 97 Histogram showing conditional probability distributions for $P(N_{satellites} | \text{Fitness} \geq 1.0006)$

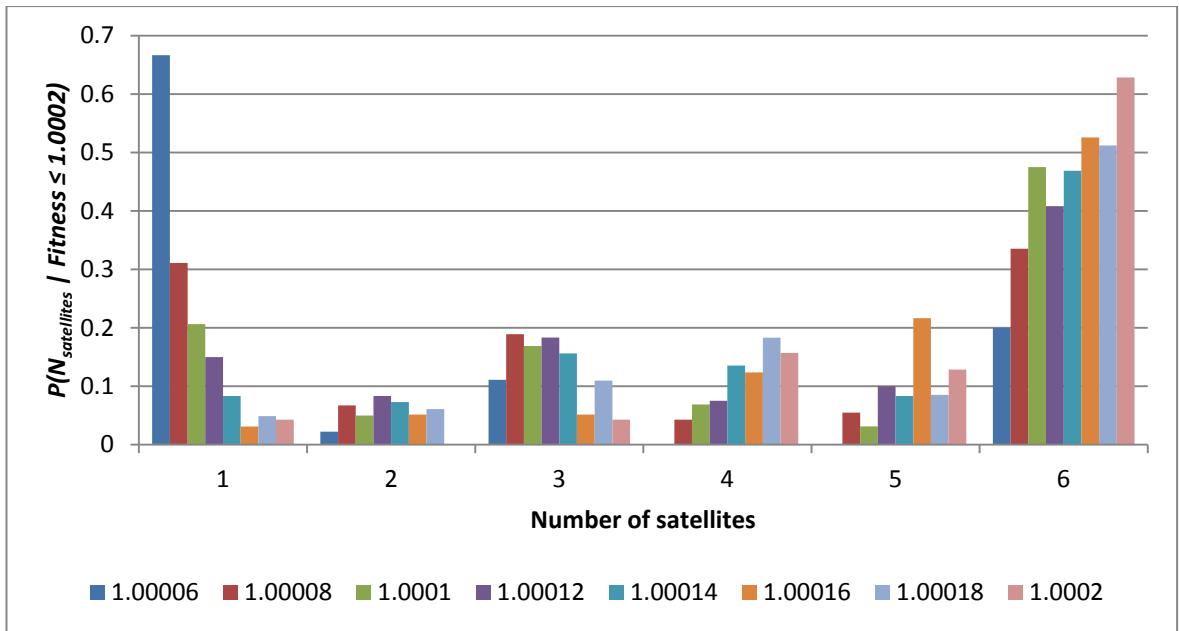


Figure 98 Histogram showing conditional probability distributions for $P(N_{satellites} | \text{Fitness} \leq 1.0002)$

Figure 99 shows that 43% of architectures explored in this dataset had $N_{subsystems} \geq 5$, and that 98% had $N_{subsystems} \geq 3$. Figure 100 and Figure 101 show similar distributions for both high and low fitness levels, in accordance with Figure 99. The two highest fitness architectures had $N_{subsystems} = 5.3$ and 6.

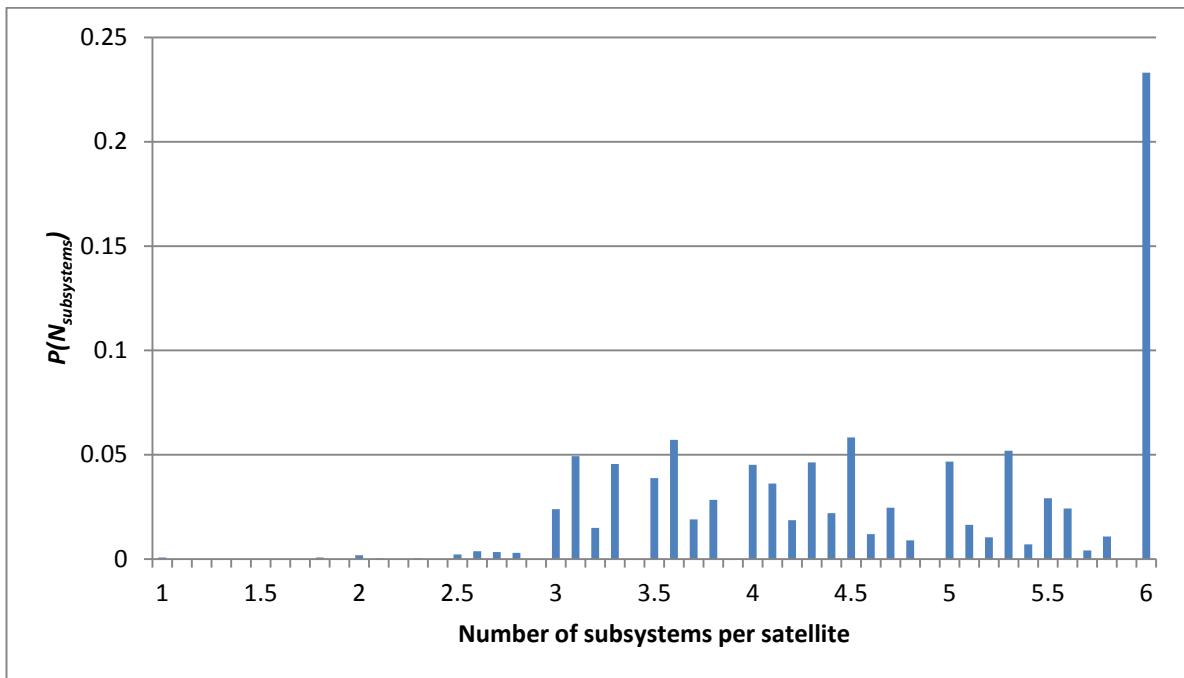


Figure 99 Probability distribution for $P(N_{subsystems})$ for the 0.7 year replacement time scenario

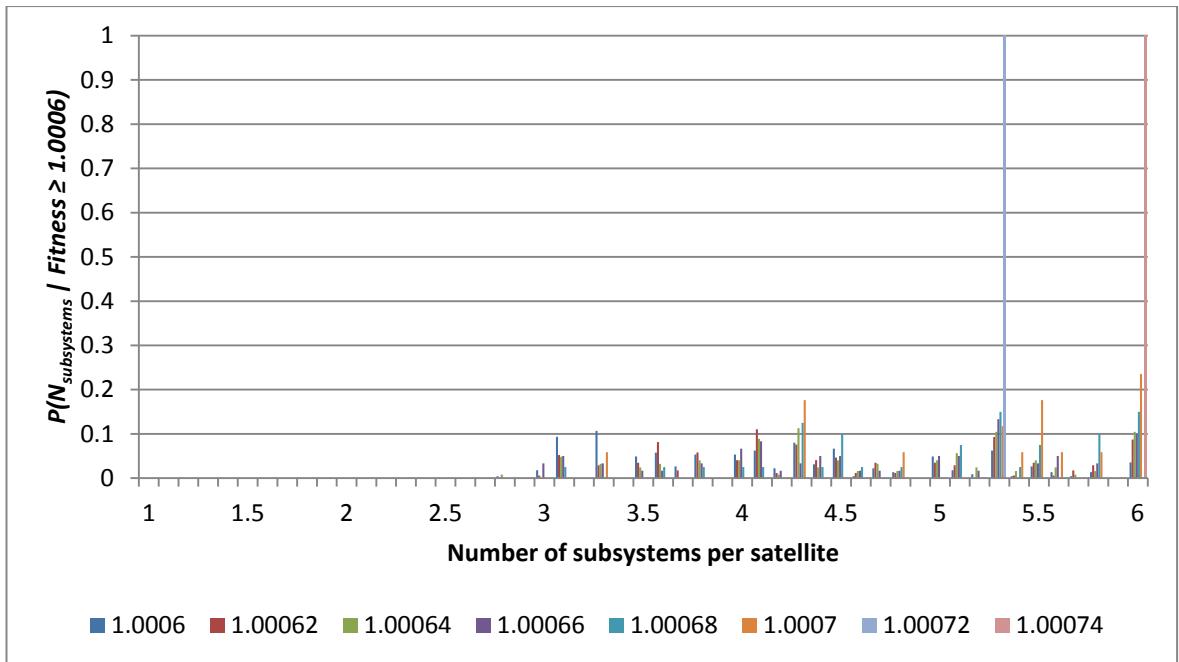


Figure 100 Histogram showing conditional probability distributions for $P(N_{\text{subsystems}} | \text{Fitness} \geq 1.0006)$

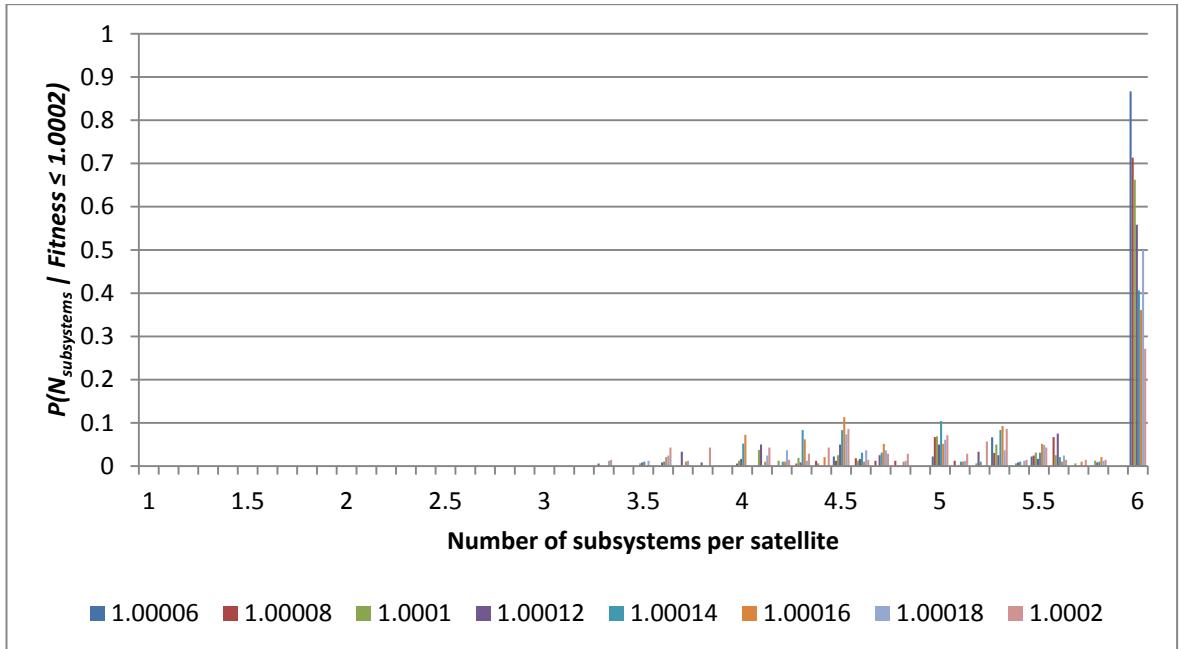


Figure 101 Histogram showing conditional probability distributions for $P(N_{\text{subsystems}} | \text{Fitness} \leq 1.0002)$

Figure 102 shows that only architectures with fitness values greater than 1.0004 had a degree of fractionation of six.

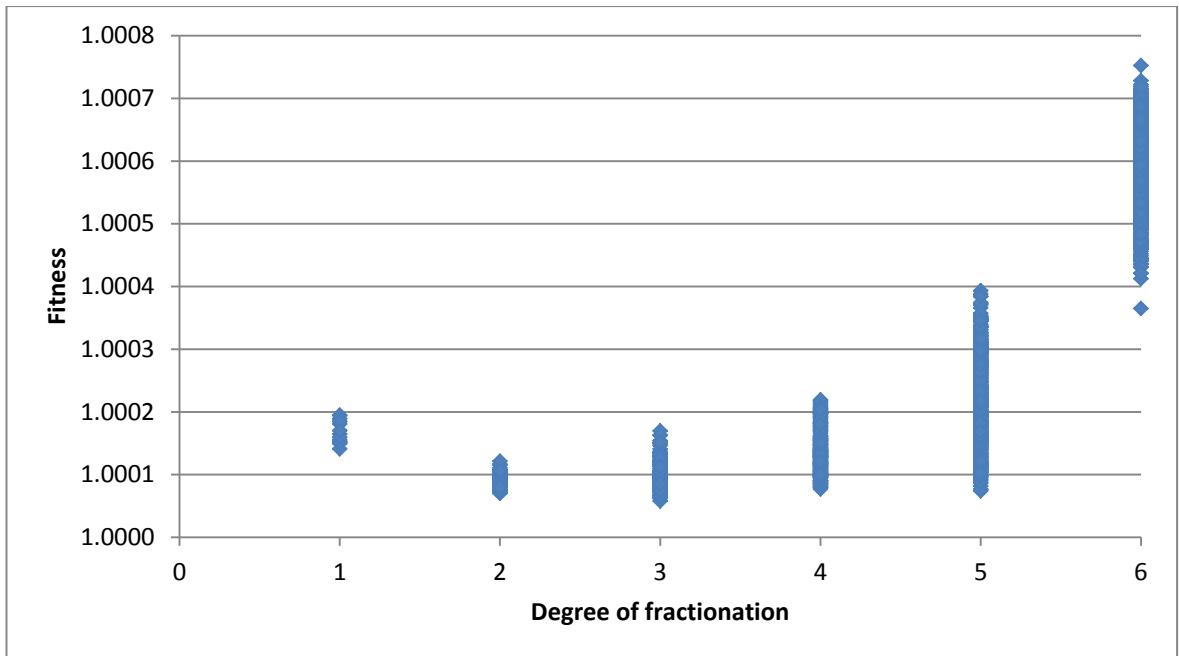


Figure 102 D versus fitness for the 0.7 year replacement time scenario

Figure 103 shows that 18% of architectures explored in this data set had $F \geq 5$, and that 90% had $F \geq 3$. Figure 104 shows that the highest fitness architectures had $F \geq 3$, with the very highest fitness architectures having $F = 5.3$ and 6. Figure 104 shows that the lowest fitness architectures had a wider range of values of F with the most likely values between 2.5 and 4.

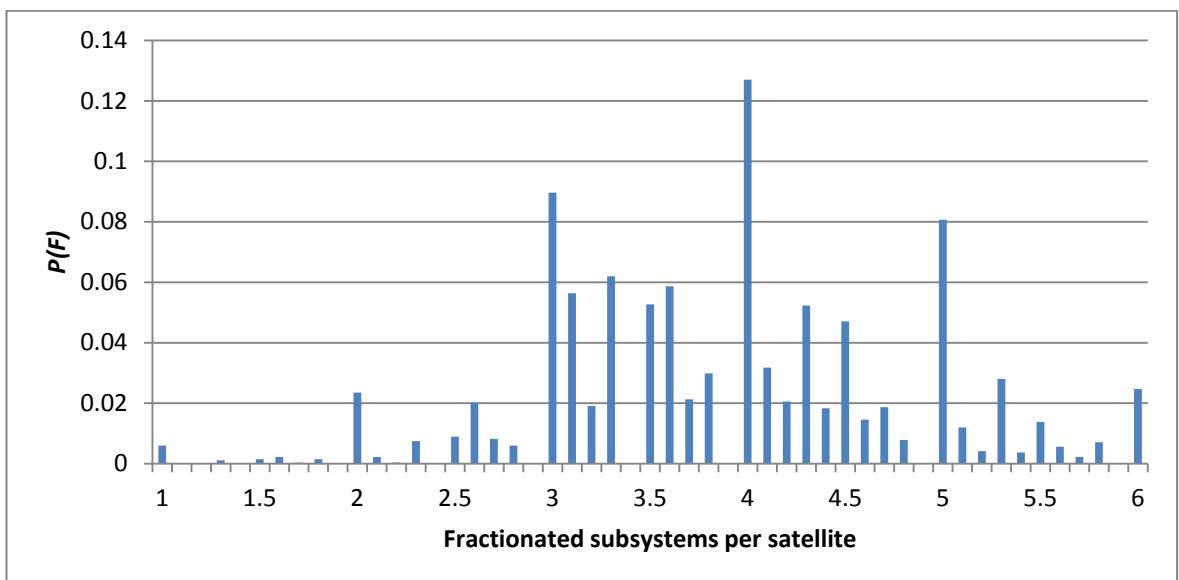


Figure 103 Probability distribution for $P(F)$ for the 0.7 year replacement time scenario

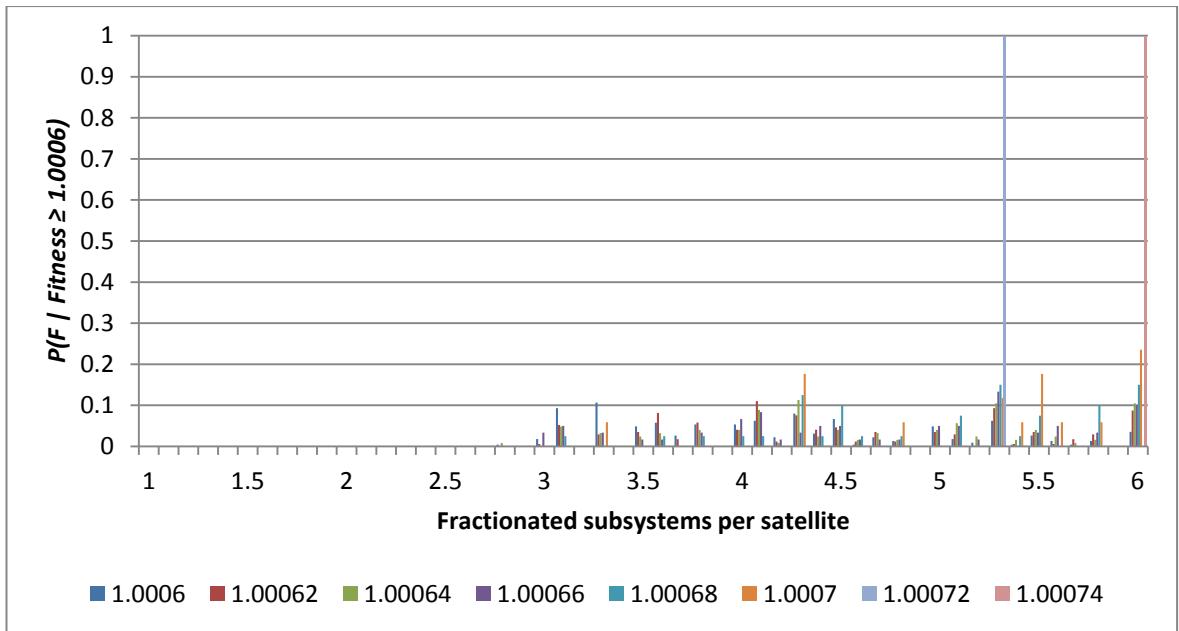


Figure 104 Histogram showing conditional probability distributions for $P(F | \text{Fitness} \geq 1.0006)$

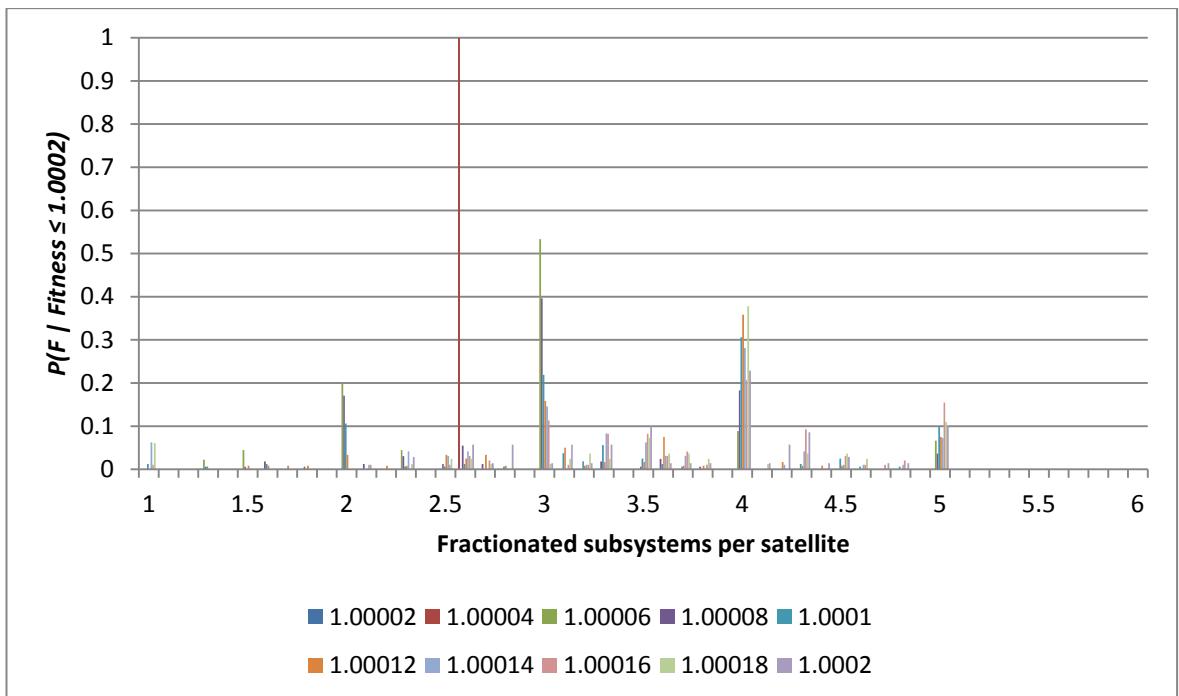


Figure 105 Histogram showing conditional probability distributions for $P(F | \text{Fitness} \leq 1.0002)$

To understand in more detail the architecture properties which produce high fitness architectures, the top scoring architecture from each model run was examined. 13 out of 15 architectures had $N_{satellites} = 6$ and the other two had $N_{satellites} = 4$ and 5. All 15 architectures had $D = 6$. Figure 106

shows that fitness is directly proportional to external redundancy in the fractionated subsystems for these top scoring architectures, while Figure 107 shows that high values for E increase the robustness to failures in the architecture.

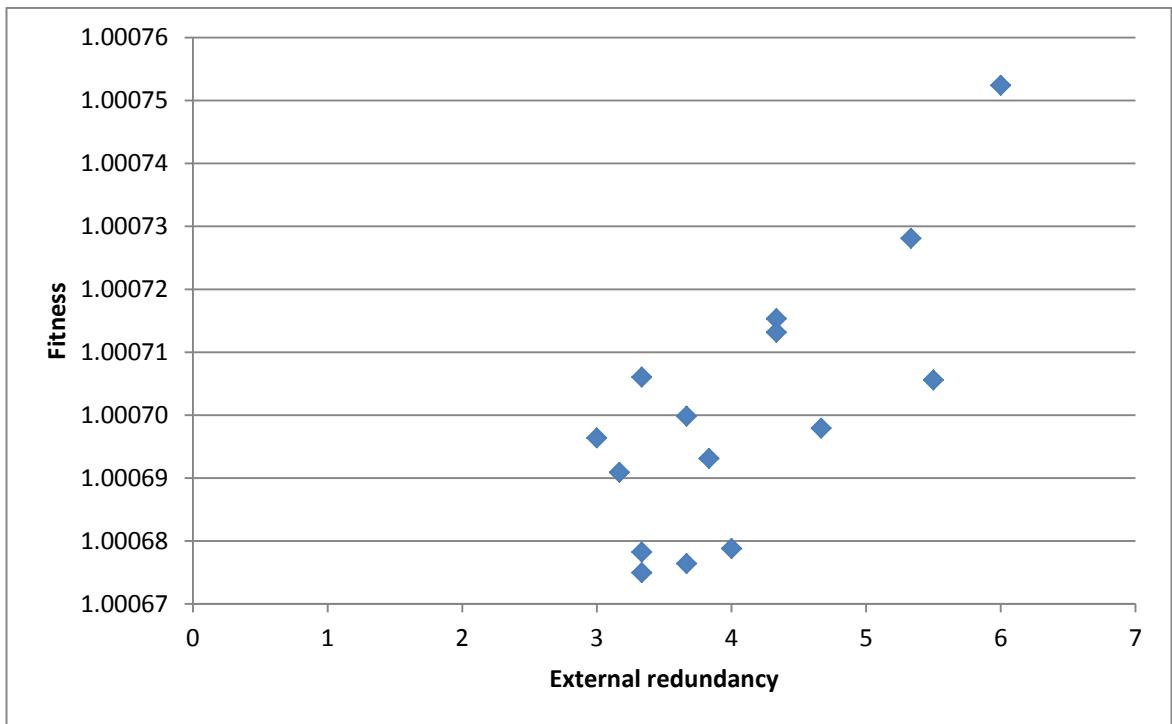


Figure 106 External redundancy versus fitness for the top 15 architectures for the 0.7 year replacement time scenario

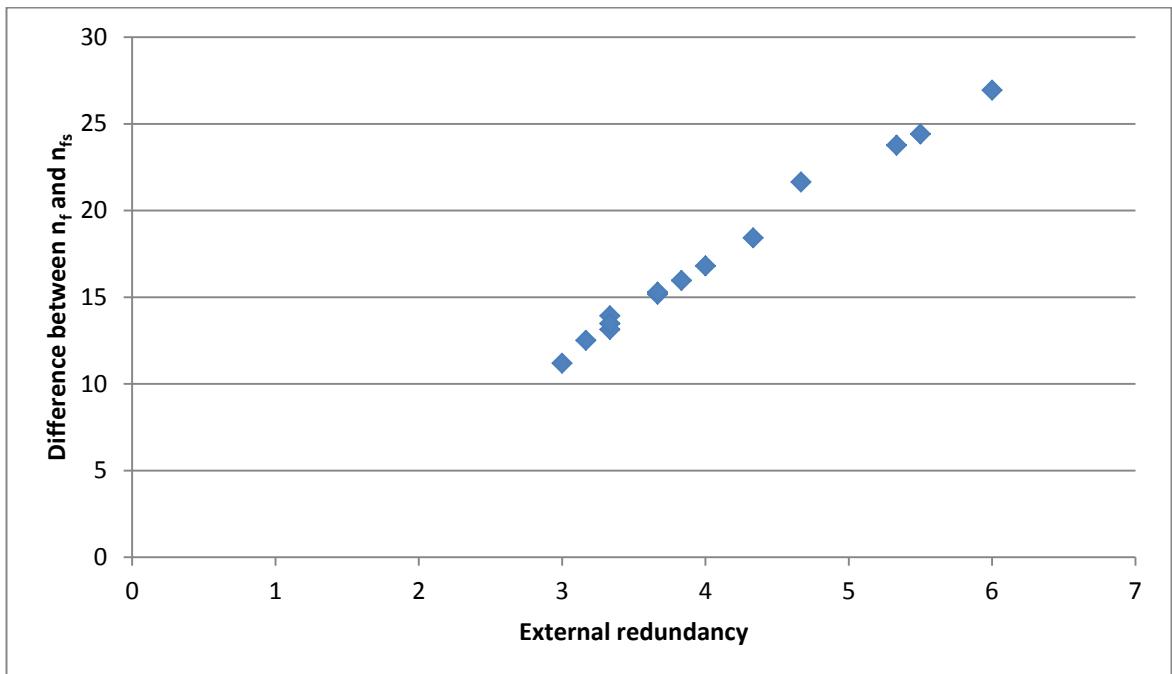


Figure 107 External redundancy versus the difference between n_f and n_{fs} for the top 15 architectures for the 0.7 year replacement time scenario

The results given in this section show that for a replacement period of 0.7 years and optimising for maximal operational time and minimal mass launched, a more homogenous approach to fractionation was preferable to more heterogeneous and monolithic solutions. This provided a high level of redundancy in the fractionated subsystems, which made the architecture more robust to failure. The best architecture found was one with six identical satellites with all six subsystems present on each satellite and all six subsystems fractionated.

This finding raised two interesting questions: is there any benefit to increasing the number of satellites beyond 6, and will there be a point at which the total launched mass becomes too high as the number of satellites increases reducing the benefit to operational time? To explore this issue, the architecture evaluation has been run for architectures with $D = 6$ and $E = 6$, with $N_{satellites}$ ranging between one and 50. Figure 108 shows that total launched mass and operational time are directly proportional to the number of satellites. However, the percentage operational time asymptotes to approximately 91%. This means that an optimum number of satellites can be found which may be more than the maximum of six satellites considered by the main model runs. To do this, the number of satellites is plotted against the fitness score, as shown in Figure 109. This shows that the maximum fitness can be found for an initial number of satellites of around four to 14 (circled in blue in Figure 109). The same analysis was performed for the 0.05 and 1.5 year replacement time scenarios.

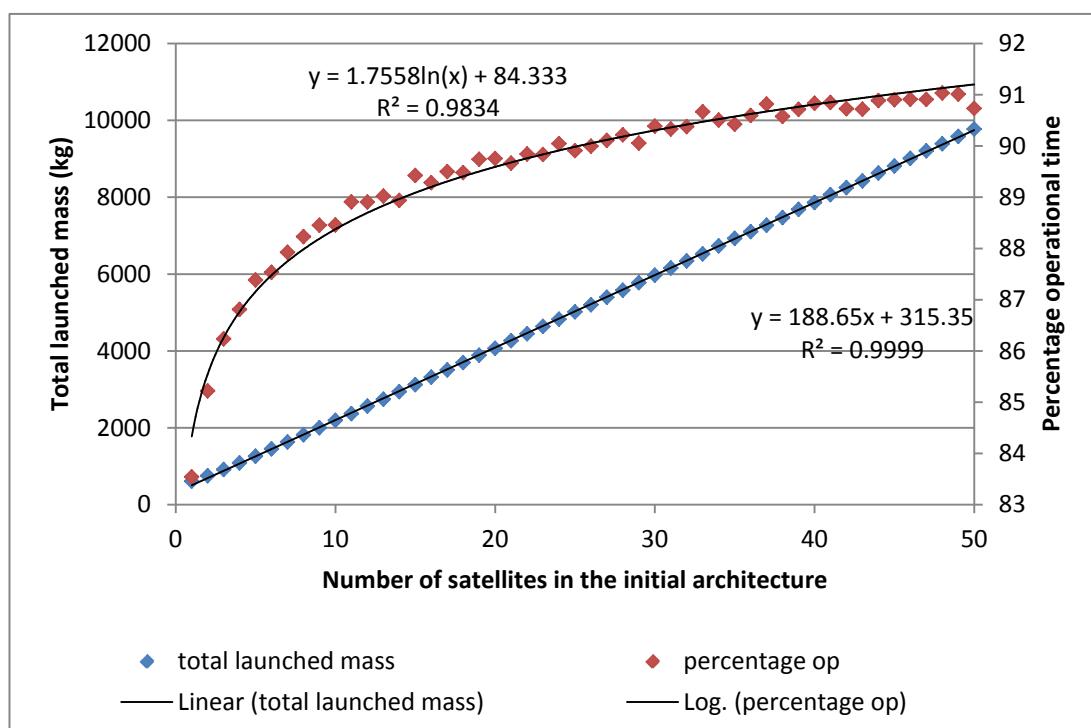


Figure 108 Number of satellites versus total launched mass and percentage operational time for $D = 6$ and $N_{subsystems} = 6$

$$N_{subsystems} = 6$$

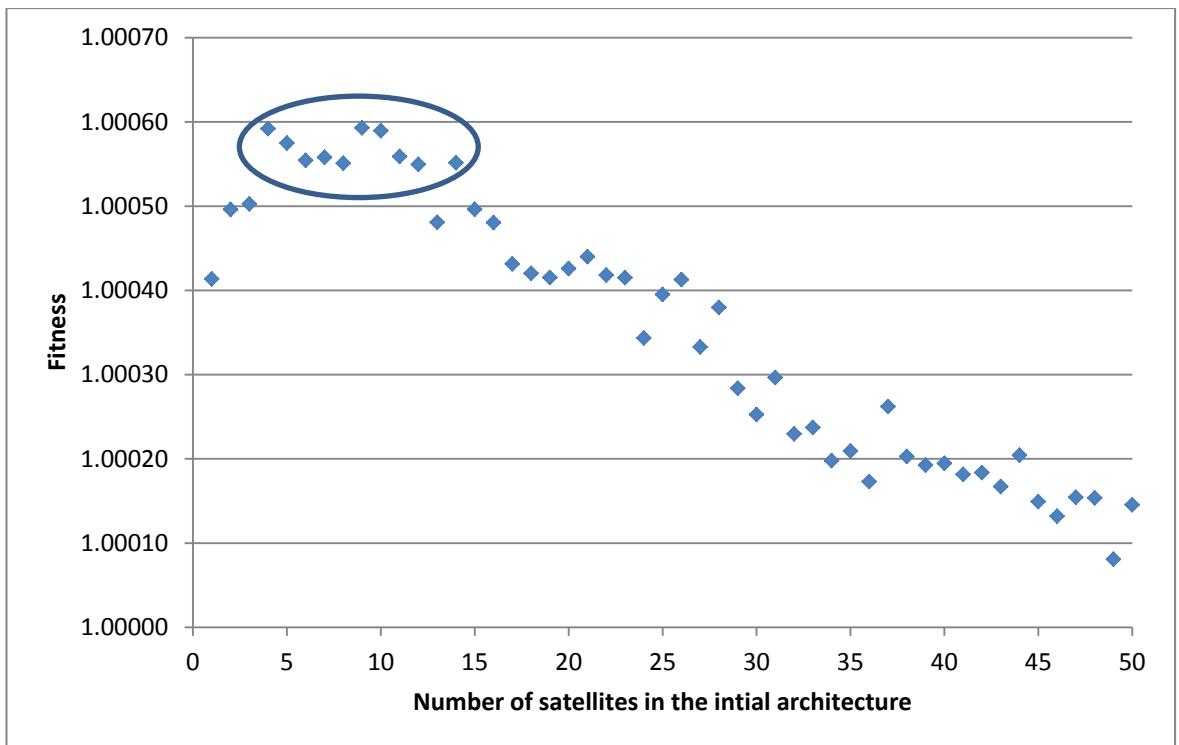


Figure 109 Number of satellites versus fitness for $D = 6$ and $N_{subsystems} = 6$ with the optimum numbers of satellites circled in blue

Using the subsystem-specific failure curves and a replacement time of 0.7 years, the results show that a high degree of fractionation and with external redundancy maximised across an architecture with four-15 satellites, will give the best compromise between operational time and the mass launched. The next scenario investigated is one with a 0.05 years replacement period, which represents the best possible scenario for a response to an on-orbit failure.

4.2.1.2. 0.05 Year replacement time

The 0.05 year replacement time scenario represents the best possible scenario for a response to an on-orbit failure. It was assumed that a range of tested, off-the-shelf-components are ready for assembly in the required combination and that they can be integrated into a spacecraft, tested and launched into the architecture in that period of time. This represents the highest level of responsiveness to failure that the model can examine, as the replacement will take place in the next time step following the failure. The results of the conditional probabilities analysis is the same as presented in section 4.2.1.1, so only the differences between this case and the 0.7 year replacement scenario will be presented. Figure 110 shows that the effect of using a shorter replacement period was to increase the operational times and reduce the range of values seen for this parameter. Figure 111 shows that there was little effect on the total launched mass. As with

the 0.7 year replacement period the highest scoring combination of architectures was $N_{satellites} = 6$, $N_{subsystems} > 5$, $D = 6$ and $F > 5$.

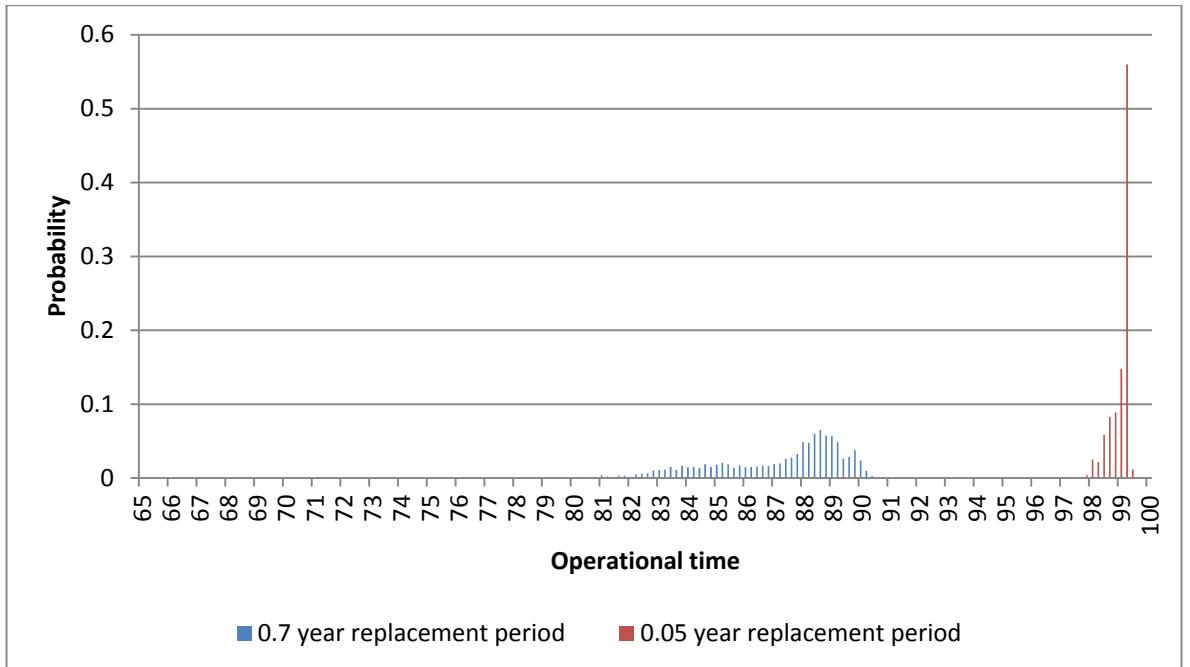


Figure 110 Histogram for operational times recorded in 0.05 and 0.7 year replacement time scenarios

By looking at the relationship between total launched mass and fitness for the two scenarios shown in Figure 111, we can understand the much larger range of fitness values found in the 0.05 year replacement period scenario. With the much smaller range of percentage operational times exhibited, the total launched mass becomes much more dominant and the lower end of the mass range produces much higher fitness values than for similar masses in the 0.7 year replacement period scenario.

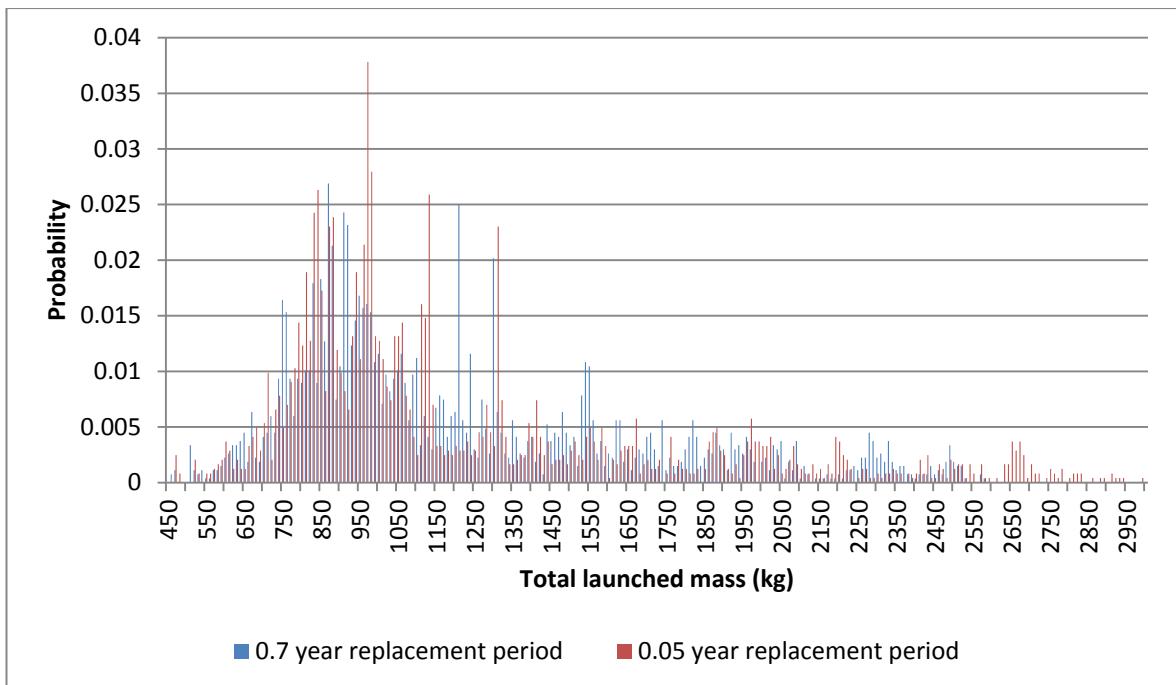


Figure 111 Histogram for total launched masses recorded in 0.05 and 0.7 year replacement time scenarios

The effect of increasing the $N_{satellites}$ for $D = 6$ and $F = 6$ systems has been investigated for the 0.05 year replacement period scenario. Figure 112 shows that the percentage operational time asymptotes to a value of approximately 99.30% while the total launched mass increases linearly with $N_{satellites}$. Figure 113 shows that the optimum fitness can be found for the case with between four to 15 satellites in the initial architecture (circled in blue in Figure 113).

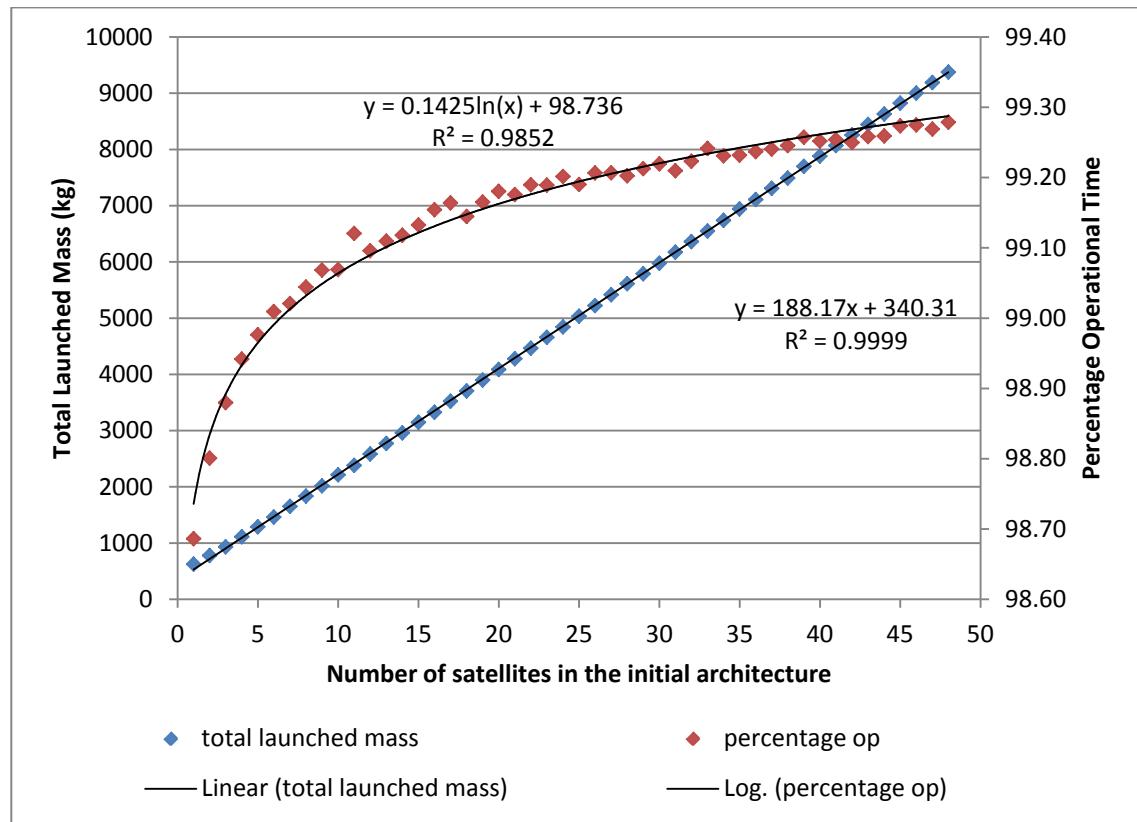


Figure 112 Number of satellites vs total launched mass and percentage operational time for $D = 6$ and $N_{subsystems} = 6$ using a 0.05 year replacement period

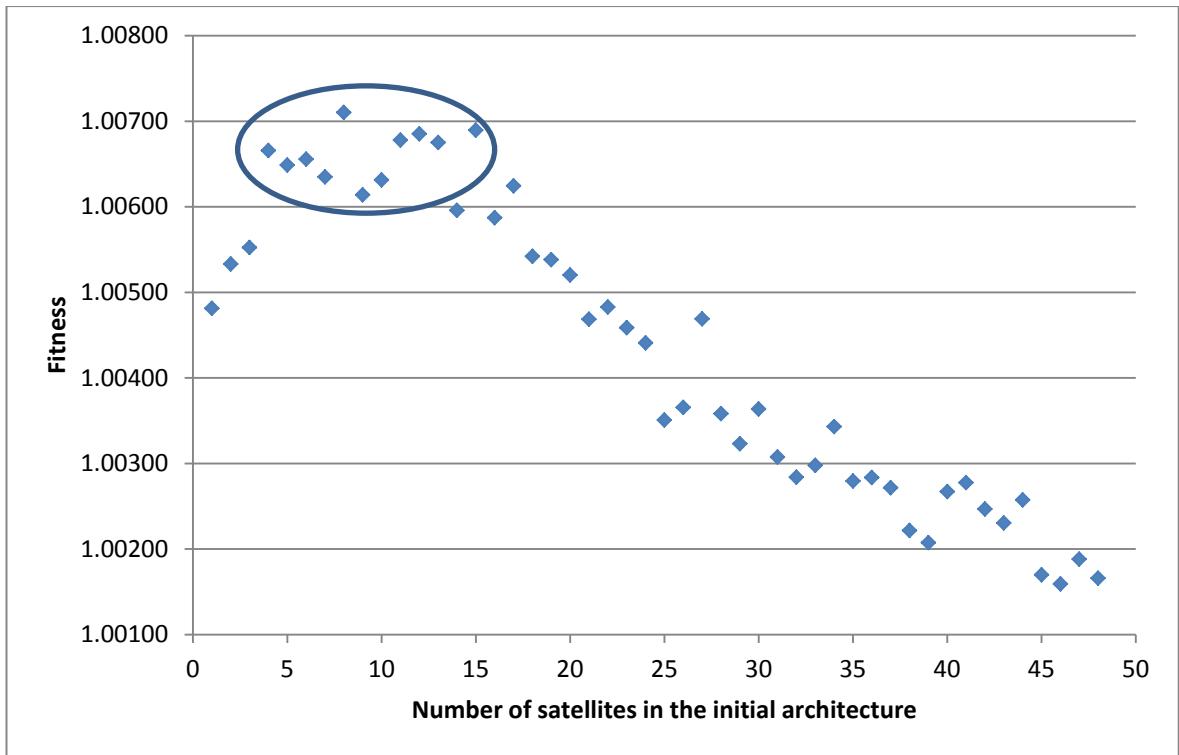


Figure 113 Number of satellites vs fitness for $D = 6$ and $N_{subsystems} = 6$ using a 0.05 year replacement period with the optimum numbers of satellites circled in blue

The results presented for this scenario show that, as for the nominal replacement time of 0.7 year, the optimal way to achieve a high operational time and minimise mass launched is to implement a highly redundant, highly fractionated system with between four and 15 satellites in the architecture. Decreasing the replacement time minimises the down-time caused by failures, and hence improves the percentage operational time as a whole across the dataset for this scenario when compared to the 0.7 year replacement time scenario. The next scenario takes a replacement time of 1.5 years, which represents a scenario closer to what is achievable today.

4.2.1.3. 1.5 Year replacement time

The 1.5 year replacement period represents an operational situation more akin to the current situation where launch opportunities take a long period of time to procure, and lead times on design, manufacture and test are also long. The results of the conditional probabilities analysis is the same as presented in section 4.2.1.1, so only the differences between this case and the 0.7 year replacement scenario will be presented. As with the 0.7 and 0.05 year replacement time scenarios the highest scoring combination of architectures was $N_{satellites} = 6$, $N_{subsystems} > 5$, $D = 6$ and $F > 5$.

Figure 114 shows that the increased replacement period reduced the operational times achieved by the architectures in this scenario. The range of total launched masses was very similar between the two scenarios (Figure 115), however for the same total launched mass, the operational time was reduced, giving the reduction in fitness.

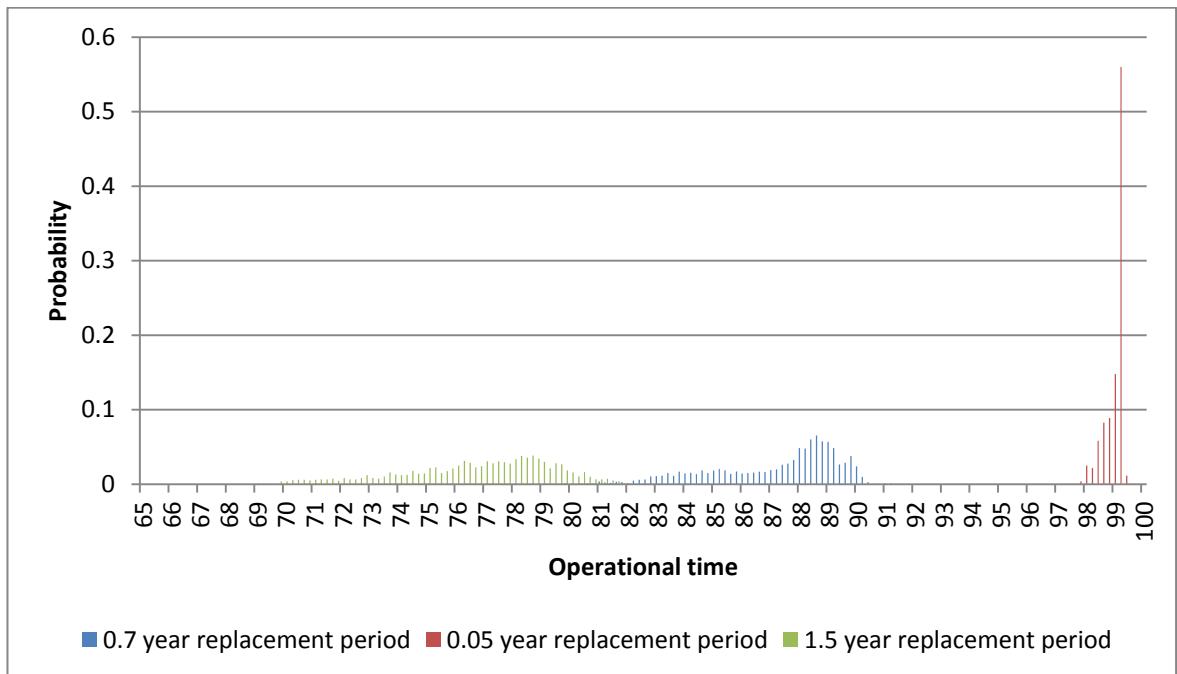


Figure 114 Histogram for operational times recorded in 0.05, 0.7 and 1.5 year replacement time scenarios

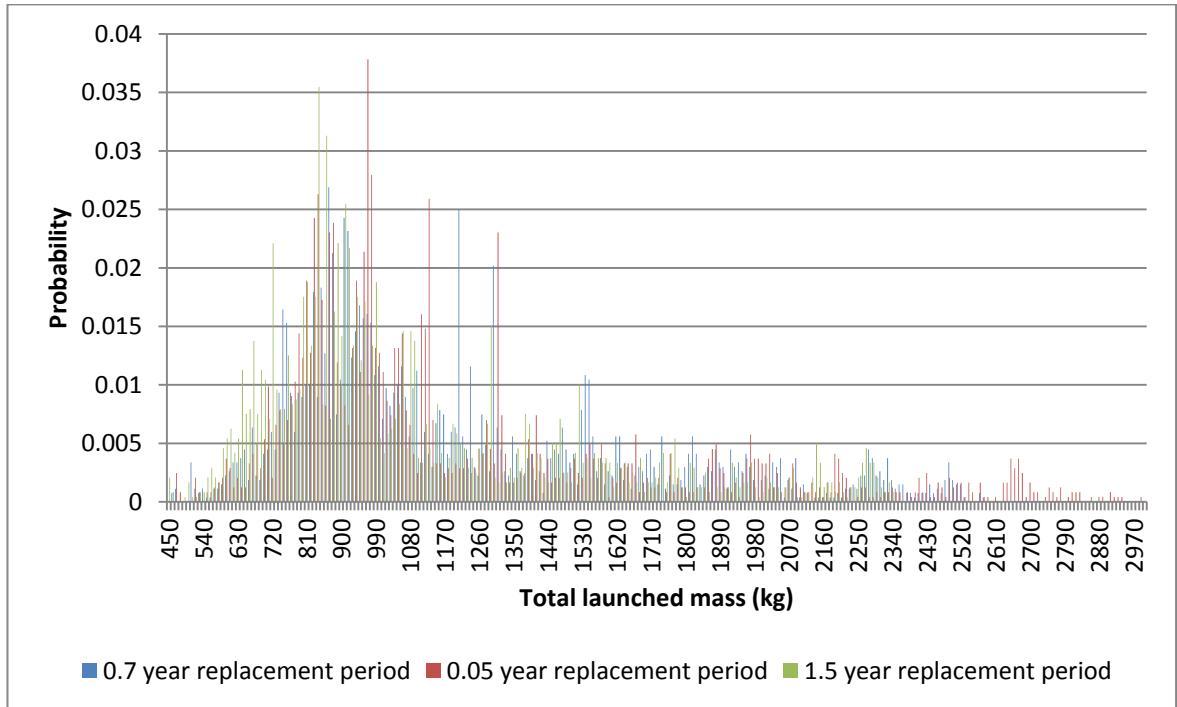


Figure 115 Histogram for total launched masses recorded in 0.05, 0.7 and 1.5 year replacement time scenarios

As with the 0.7 year replacement period scenario, the effect of increasing $N_{satellites}$ for $D = 6$ and $F = 6$ systems has been investigated. Figure 116 shows that the same relationships shown for the 0.7 year replacement period hold for this case, with the percentage operational time asymptoting to a value of approximately 83%. Figure 117 shows that the optimum fitness can be found for the case with around three to 17 satellites in the initial architecture (circled in blue in Figure 117).

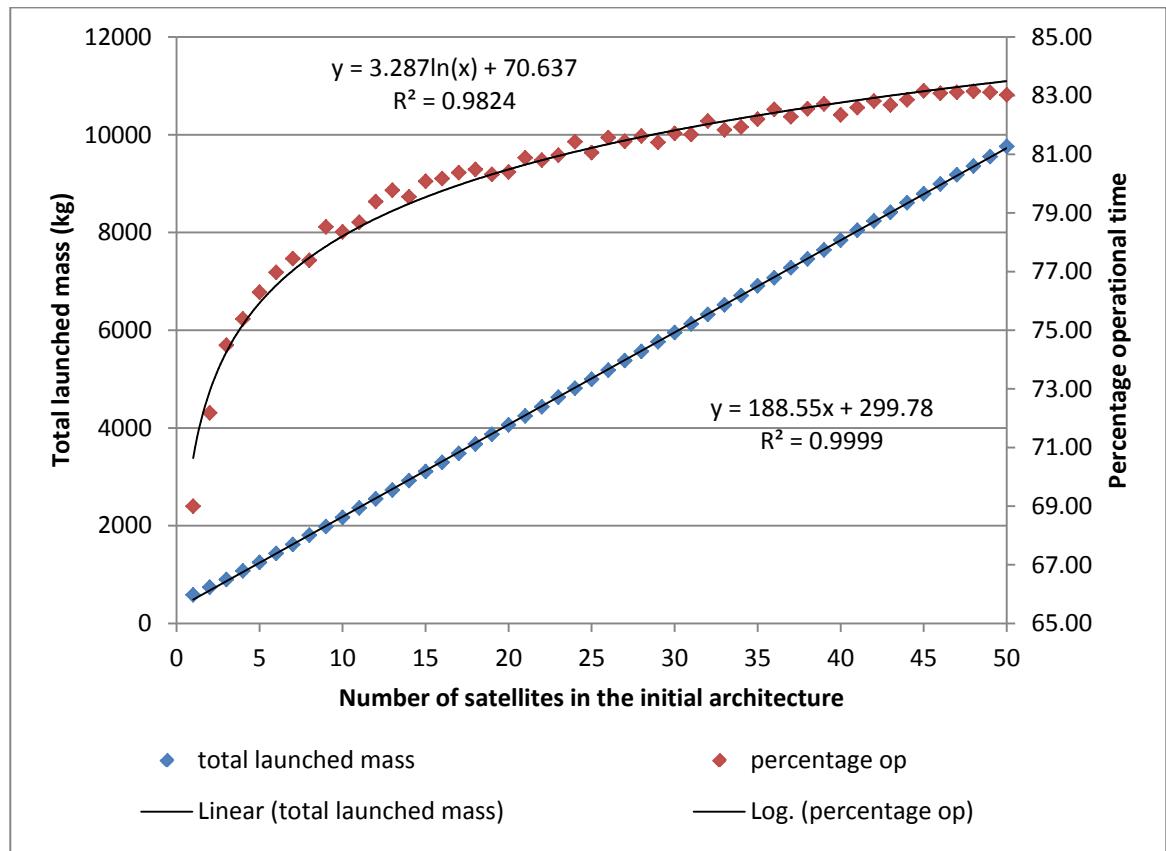
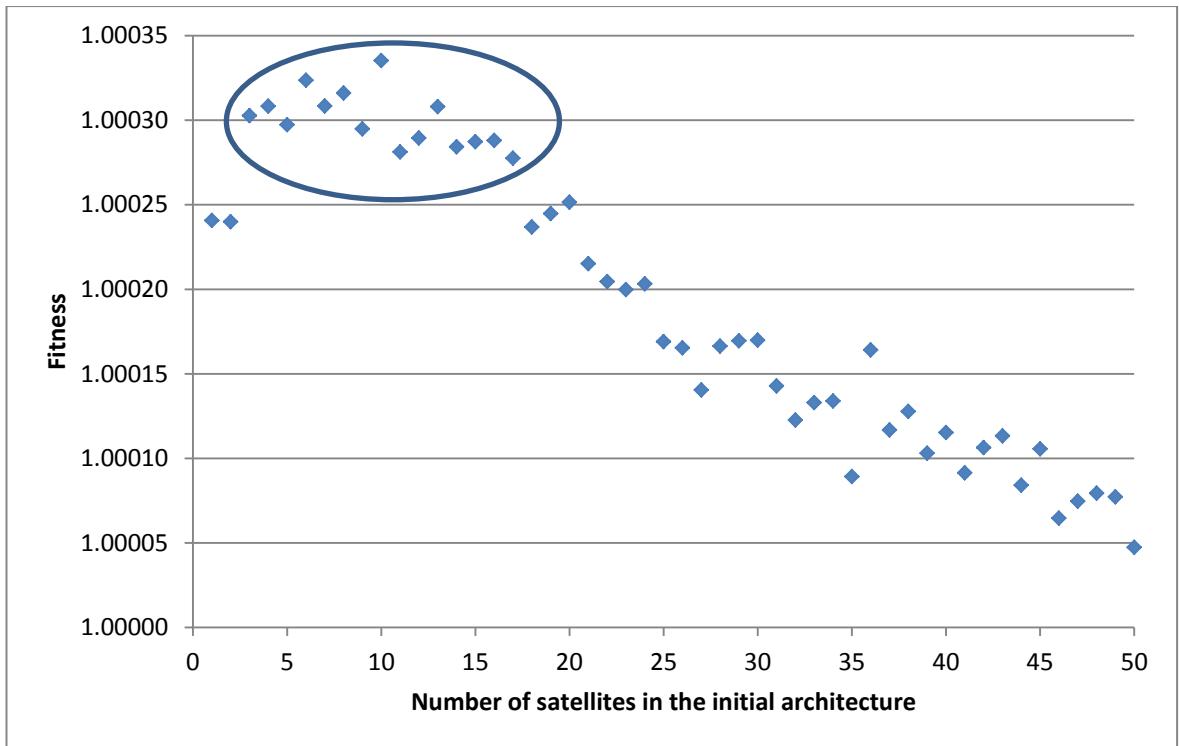


Figure 116 Number of satellites vs total launched mass and percentage operational time for $D = 6$ and $N_{subsystems} = 6$ using a 1.5 year replacement period



because the time to replace the failed subsystems is so short, but the question remains: do failures occur more frequently with short replacement periods?

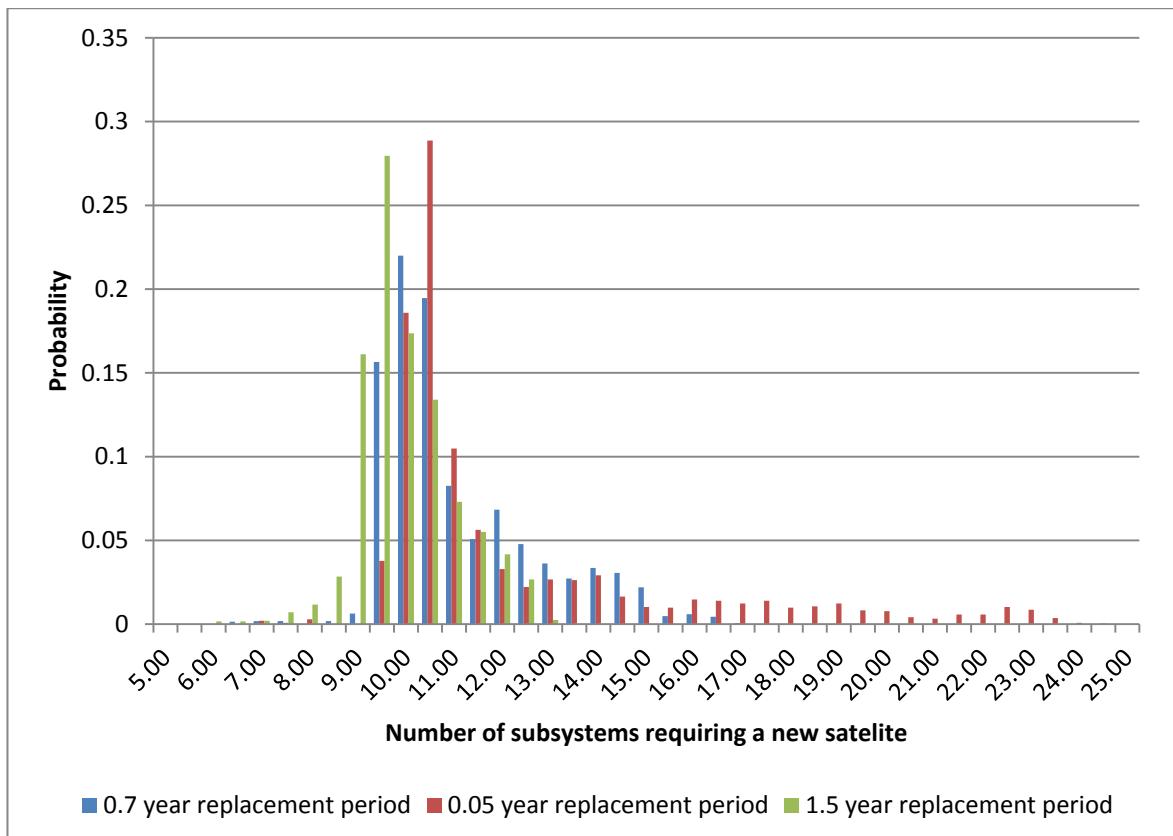


Figure 118 Number of failures requiring a new satellite comparison

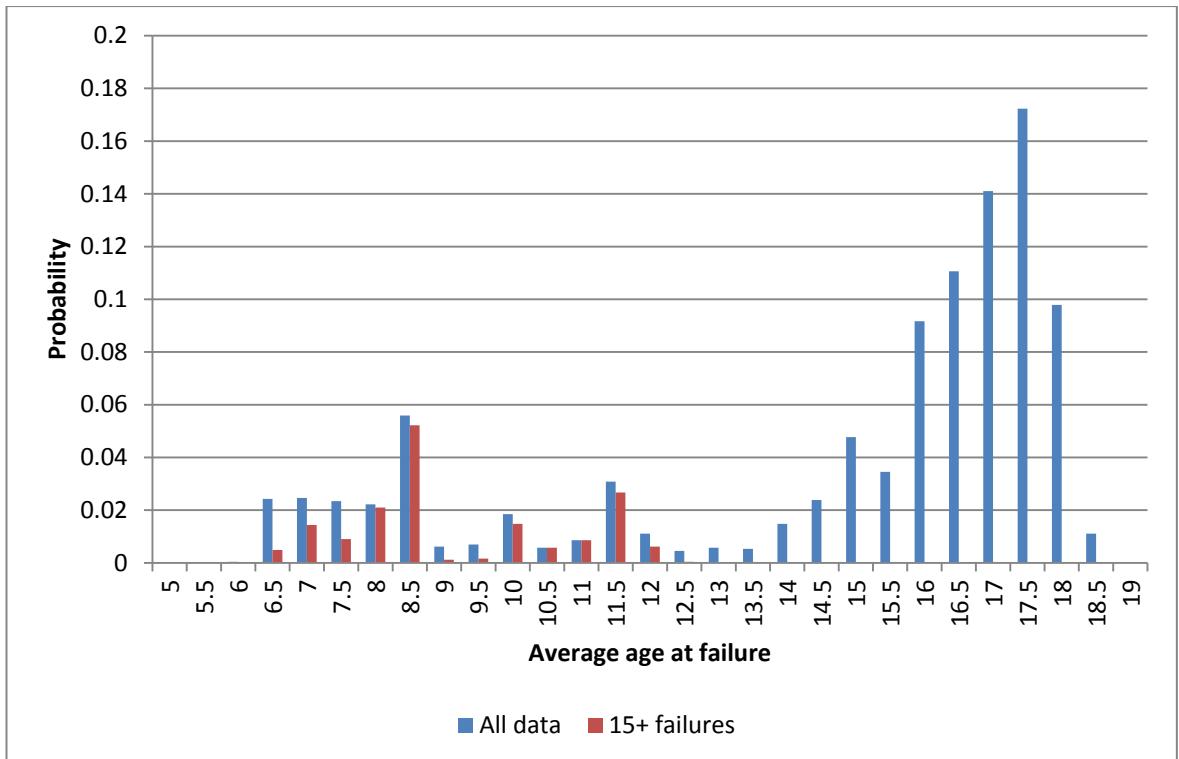


Figure 119 Average age at failure, 0.05 year replacement period

Figure 119 shows the a_f distribution for the 0.05 year replacement period scenario. The data points in blue are the full data set for this scenario, and those in red are the architectures which experienced 15 or more n_{fs} . This graph shows that a lower average age at failure is generally associated with higher numbers of failures. However, the average age at failure is not the only important factor, as some of the architecture properties will also contribute to the number of failures.

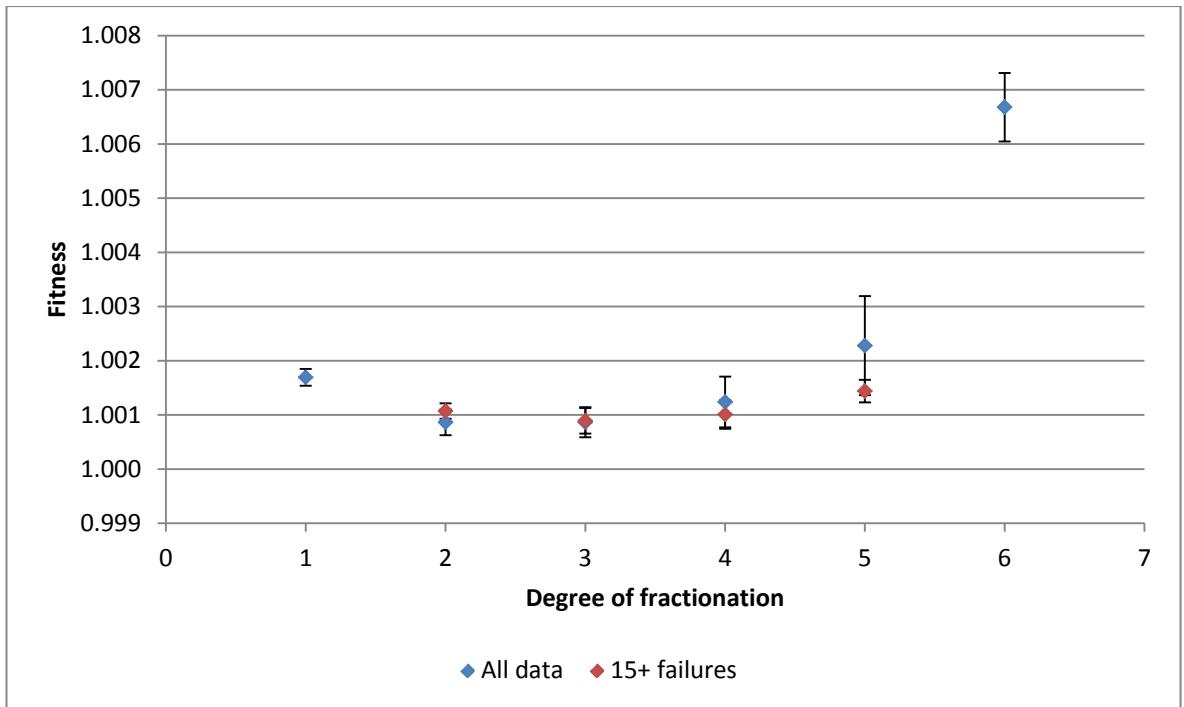


Figure 120 Degree of fractionation vs fitness, 0.05 year replacement period for all architectures and those that experienced 15 or more failures. Error bars show 1 standard deviation

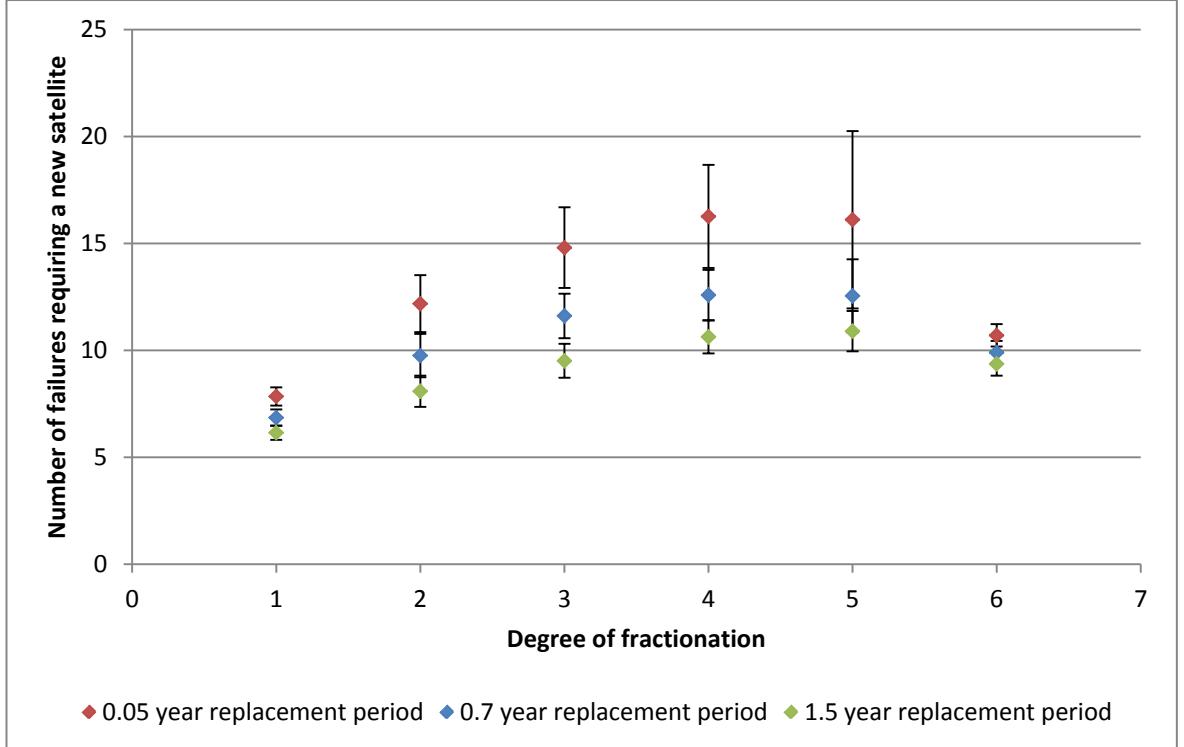


Figure 121 Degree of fractionation vs number of failures requiring a new satellite comparison. Error bars show 1 standard deviation

Figure 120 shows the degree of fractionation vs fitness for the 0.05 year replacement period using the same colour coding as above. Only architectures with $D = 1, 6$ were found to experience fewer than fifteen failures. This is confirmed by Figure 121 which shows that architectures with $D = 1, 6$ experienced the fewest failures, and had the smallest variance in the n_{fs} . This is due to the effects of the failures on the different architectures. $D = 1$ corresponded to a monolithic system with only the payload considered fractionated. Any failure of a non-payload subsystem in one of these architectures resulted in the host satellite being lost. The replacement strategy used limited the number of failures that a monolithic system can experience. A replacement was only launched if there was no payload in the architecture, and often this only occurred when no satellites were left in the architecture. The result was that after the initial satellites have failed, the architecture carried on with only one, monolithic spacecraft in it, which failed after a period of time, and was replaced by another monolithic spacecraft after the replacement period had elapsed, which then failed, and so on.

$D = 6$ was at the opposite end of the scale. With these architectures, a replacement was only required if the failed subsystem was the only one in the architecture (i.e. if $E = 1$). As shown in sections 4.2.1.1, 4.2.1.2 and 4.2.1.3, higher values of E gave more robustness to failure. For other architectures with $2 \leq D \leq 5$ there was a middle ground with between one and four non-fractionated subsystems present on each satellite in the architecture. In the event of a failure in one of these subsystems, the result was the loss of the host satellite and any fractionated resources on-board. This represented an inherent lack of robustness to failure in these architectures. The result was the high number of failures exhibited in the data presented, and the decrease in the number of failures between a degree of fractionation of 5 and 6.

The results presented in 4.2.1 all indicate that a fully fractionated ($D = 6$), fully redundant system ($E = 6$) with that redundancy distributed on five to 15 free flying satellites, is an effective way of maintaining continuity of service over an extended mission lifetime whilst minimising the mass launched. It has also been shown that the time taken to respond to failures should be minimised where possible.

4.3. Discussion

In this research, a compromise between operational time and mass launched for fractionated spacecraft was sought. The predominant factor found to improve the operational time was the time taken to launch a new replacement. The results presented in section 4.2.1.2 showed that using $t_r = 0.05$ years (approximately 18 days) allowed operational times of up to 99.2% to be achieved. This corresponds to a down-time of only 4.8 months over the 50 year lifetime of the

architecture. The predominant factor which acted to reduce the mass launched was the design of the architecture. Highly fractionated systems, with a high *external* redundancy provides robustness to subsystem failures.

For the initial data exploration presented in 4.1.1 the optimisation was performed with the objective of maximising the operational time. This enabled the model to be understood using a simple optimisation parameter and for the effect of operational time on the outputs to be determined. For the data exploration and results that followed the fitness parameter that combined both operational time and total launched mass was used. This sought to encapsulate a desired characteristic of a space system (maximal operational time) and a proxy for the “cost” of implementing the system to the user (mass launched). Although this fitness parameter is not a standard spacecraft characteristic, it allowed a relative “score” for each architecture to be determined allowing the search algorithm to progress through the search space. The large standard deviations observed in the fitness values such as those presented in section 3.4.4 are the result of the random nature of the subsystem failures. The use of a large number of Monte-Carlo runs (300) attempted to alleviate this issue, nevertheless the variation between fitness values for different architecture was smaller than the standard deviation on the fitness values for each architecture. While this could indicate that *any* of the architectures could give the highest fitness, it overlooks the fact that outputs from each individual run of the model consistently showed that the highly fractionated and redundant architectures produced the highest fitness values. It is this consistency over many runs of the model that gives confidence in its outputs and shows that the optimisation is not just driven by the random failure aspect of the model.

The conditional probability analyses produced individual distributions for individual model parameters. These were used together to draw some conclusions about what fractionated architectures should consist of. The individual conditional probabilities could not be combined mathematically as the model parameters are not independent of one another.

The current method of designing and operating monolithic space assets does not provide the best compromise between operational time and launched mass. The data generated for the $t_r = 1.5$ years scenario showed that the monolithic architectures produced launched mass values of approximately 1350kg and an operational time of approximately 80%. The mass was in the middle of the range of values (450-2300kg), while the operational time was near the high-end of the range of values (68-83%). Reducing the replacement time to 0.05 years improves operational time to approximately 99%, however the launched mass does not decrease. This results in the fully fractionated system still producing the highest fitness values in the dataset.

These findings are in contrast to the fractionated systems proposed by the DARPA studies [5], [11]–[14], in which each fractionated subsystem is a unique resource in the architecture, as it is hosted in individual satellites with little or no external redundancy. Employing the DARPA-proposed architecture would result in fractionated subsystems becoming single points of failure. If, for example, the only fractionated power system in the architecture fails, then the entire system is without power. However, if there were several power systems in the architecture, each sized to cope with the loss of one or two of the other power sources, this would avert the possibility of a single point failure in this fractionated subsystem. A fully externally redundant system would be simpler to implement, as any FDIR system used would be required to redistribute tasks rather than cope with an “absent” subsystem.

The main limitation of the approach taken in this research was the basic nature of the replacement strategy. In the model developed for this research the replacement process was only initiated when the architecture lost its operational capability, leading to a period of down-time corresponding to the value of t_r used. A potential option for an alternative replacement strategy is to try to anticipate an upcoming failure and launch a replacement satellite before the failure can occur. This strategy would, in theory, improve the operational times for all architectures up to 100%. The focus of the optimisation would then be the mass launched which is most strongly affected by the design of the architecture. The likely outcome of this scenario would be that homogenous, fully fractionated architectures produce the highest fitness values, but the number of satellites required in the initial architecture may be reduced.

There are important limitations to implementing a real, externally-redundant fully fractionated system. Not all subsystems can be easily fractionated using technology available today or in the near-term future, in particular, the power, attitude control and propulsion subsystems as described in section 3.3. This means that if a fractionated architecture is to be implemented in the near term, there will inevitably be some non-fractionated subsystems which will have to be present on every satellite in the architecture. In the event of a failure in one of these subsystems, the entire host satellite could be lost, together with any fractionated subsystems which share their resources with the rest of the architecture. Therefore, maximising the redundancy in the architecture for the fractionated subsystems makes the system robust to a failure in the non-fractionated subsystems. In the long-term, technologies for fractionating the power system, attitude control system and propulsion system should be developed so that the level of external redundancy and degree of fractionation can be increased.

It was assumed that in an architecture with a fractionated power system each satellite would carry a PCDU and battery to allow the satellites to function in the event of the primary power source being lost. This implied that the depletion of the battery is a recoverable failure. In reality, if a battery is depleted it is an irrecoverable failure. Therefore it is likely that failures which remove all primary power sources may result in the entire architecture being lost and having to be replaced, increasing the mass launched. This could be mitigated against by ensuring that the external redundancy in the primary power sources is maintained to a sufficient level by launching additional satellites hosting a power source even when no failure has occurred. Again this would increase the mass launched, but would be by a much smaller amount than that required to replace the entire architecture. A fractionated OCS was modelled taking the mass as that of a cold-gas propulsion system, due to the availability of data. The IBS concept was suggested in section 3.3.5 as a potential option for the fractionation of the OCS. If this were to be used, it would be likely that the launched mass for architectures with a fractionated OCS would have a higher total launched mass than the simulation results have shown, although the operational time will still be higher if the high level of redundancy is maintained. The 2-mixed Weibull probability density functions were open ended and didn't account for any consumables such as propellant. In reality spacecraft can only carry a finite amount of propellant, and once this runs out the mission is often terminated (or a portion of the propellant is reserved to de-orbit the spacecraft in accordance with space debris mitigation guidelines). In the results presented in 4.2.1.1 (0.7 year replacement period scenario) the average age at which the OCS subsystem failed was 15 years, with a maximum of 25 years. The SSTL OCS subsystem used in the model is only designed to support a mission duration of 7.5 years which was not accounted for in the model. If it were to be accounted for, non-fractionated OCS subsystem would fail approximately twice as often as observed in the results presented. If the chosen technology for a fractionated OCS subsystem uses a propellant (such as the IBS concept) then this too will limit the lifetime of this subsystem. Any effect that increases the number of failures adds weight to the case for employing externally redundant architectures and the use of predictive replacement strategies to ensure continual operations. This will come at the cost of an increased launched mass due to the increased number of replacement spacecraft that will be required.

In their 2010 paper, Dubos and Saleh presented results from a study in which they made a direct comparison between monolithic and fractionated spacecraft using failure and replacement Markov models [4]. Their study also considered fractionation for its potential to extend a mission beyond a single spacecraft lifetime. Their starting assumptions were very similar to those made here, but they used different simulation methods. However, two key differences exist between

the two approaches. Firstly, they only consider failures at a spacecraft level, and therefore cannot model any redundancy in the fractionated subsystems. Secondly, they consider the failure probability to be constant, rather than a probability of failure which changes with the age of the subsystem. Their study came up with similar conclusions to those presented here, with higher degrees of fractionation producing a lower cost (mass in this work) and higher utility (percentage operational time in this work). It is worth noting, however, with each satellite modelled as a black box rather than at the subsystem level, their term “degree of fractionation” refers to the number of fractionated satellites modelled. They also concluded that the lowest level of responsiveness (minimal replacement time) produced the highest utility, which is supported by the work presented in this thesis.

This work has focused on a mission with a single earth observation measurement objective, however many spacecraft host multiple payloads with different observation objectives. If these multi-payload missions are to be fractionated, the redundancy in the subsystems should still be present, however the level of redundancy in the payloads is more uncertain and future work should focus on this aspect of multi-payload fractionated architectures.

In section 2.5, a concept design was presented for a fractionated coastal salinity measurement system. Here, this design is discussed in the context of the outcomes of the simulation and optimisation work. The coastal salinity measurement system consisted of a primary payload consisting of 72 free flying spacecraft each carrying one Doppler Radiometer receiver, one central hub carrying 36 receivers and the correlator unit, used to combine the signals received by the antennas, a microwave radiometer satellite and a scatterometer satellite. The three satellite bus spacecraft and their main mission assets were as follows: an AODS spacecraft, carrying an attitude and orbit determination sensor subsystems, a communications subsystems carrying a high-speed data, telemetry and telecommand link and an OBDH spacecraft carrying the main mission processor and data storage. Taking into account the knowledge gained from the simulation and optimisation, the bus spacecraft should have the fractionated resources replicated across each of the three satellites to maximise redundancy. The primary payload will be robust to a small number of failures in the free flying spacecraft however the central hub with the correlator unit and 1/3 of the payload receivers will require an on-orbit spare to provide redundancy in case of failure in the primary satellite. Operationally, the architecture could cope with a failure in one of the two secondary payloads, providing that the SST or sea-surface roughness measurements could be obtained from another satellite, even if they are not made at the same time as the SSS measurements.

5. Conclusions

This thesis has investigated the concept of fractionated space systems, and reviewed the current literature on this subject. Two key questions were addressed in this research, firstly, is fractionation a mission enabling technology? Secondly, what is the best way to design a fractionated system, such that a continuity of measurements (or other payload service) can be maintained over a period of time that is greater than a “traditional” mission lifetime (10 to 12 years) whilst minimising the mass launched over the operational lifetime?

To address the first question, one possible mission was explored, which used the fractionated principles of physical decomposition and information integration in both the supporting subsystems and the payload to measure coastal salinity using a microwave radiometer. It was shown how such a potential mission could be achieved using a system of a few tens of payload satellites and two to three supporting satellites. Through this design process it was demonstrated that fractionation is indeed a mission enabling technology, as the measurement of coastal salinity from space is not possible using a conventional, monolithic system. An initial exploration of the fractionation of spacecraft subsystems was conducted which allowed a foundation to be built for the more general exploration of the concept that followed. It also allowed the questions which were to be explored to be refined, with some expectations about how fractionated spacecraft would function.

Following on from this design, the more fundamental concept of exactly how a fractionated system might be implemented for a simpler, single payload mission was explored in the context of the second key research question. This was achieved using a simulation and optimisation tool developed as a part of this research. It was shown that a highly fractionated system, with a high level of redundancy spread across five to 14 satellites, gives the best compromise between operational time and launched mass in order to allow a continuity of measurements to exist over an extended mission period, exceeding that of a normal satellite mission. This finding contrasted with the heterogeneous architectures proposed by DARPA. A key finding from the simulation analysis was that, in order to achieve the best possible operational time, the time taken to replenish lost resources should be minimised. Maximising the redundancy in the fractionated subsystems is the dominant factor in reducing the mass launched (and therefore cost).

A key assumption in the development of the model was that the technology was available to fractionate every subsystem in the architecture. In reality, the technology readiness levels (TRL)

for the fractionation of each subsystem will be different. Therefore, a fractionated architecture may be implemented with only some of the subsystems fractionated. In order to avoid any fractionated resources being lost due to a failure in a non-fractionated subsystem, the fractionated subsystems should be replicated across several subsystems.

On the 16th of May 2013, DARPA announced that the System F6 program had ended, and that they would undertake no further work on this project [97]. A key reason cited for this decision was the lack of a business case for heterogeneous fractionated systems. The USAF is continuing to study the wider fractionated concept under the term “Disaggregated Satellites”. This concept treats spacecraft fractionation in much the same way as the research presented here [6]. Much like the internet started as a defence research project into fault-tolerant computer networks, spacecraft fractionation has the potential to develop into a “global” infrastructure for Earth-orbiting satellites. One can imagine that such an infrastructure could provide power, data handling and communications, attitude and orbital control and all the operator would have to do is to launch a satellite consisting of their payload and the required hardware to “plug into” the infrastructure.

Although this PhD was unconnected with System F6 or the USAF study, the DARPA studies were taken as a starting point. The purpose of the work undertaken in this research was to complement the value-orientated approach by analysing the fractionated spacecraft concept from a systems engineering point of view. The outcomes of the analysis presented here provide a contrasting view on the way to design and implement fractionated space systems when compared to the DARPA concept.

There is strong potential for taking the work presented in this thesis further. In particular, it would be valuable to analyse a system which used multiple payloads, and even explore the fractionated payload concept described in section 2.5 further using the simulation and optimisation tool developed. This would provide insight into operational strategies required to implement these mission types and highlight any potential architecture designs that should be aimed for or avoided. Further exploration of the fractionated payload concept could also highlight other examples of missions that are enabled by the use of fractionation. The impact of more novel fractionation technologies (such as IBS) should be investigated within the context of the simulation and optimisation work to understand their impact as enabling technologies for fractionation. This would enable designers and operators to fully appreciate whether it is worth implementing a fractionated system with only some of the technologies at the required maturity, or whether it is better to wait for technologies such as IBS or wireless power transmission to be developed before

a fractionated system. Finally, a more realistic approach to dealing with failures should be investigated, to include the development of an algorithm which could anticipate the failure of a satellite or subsystem and launch a new satellite to boost the redundancy and ensure continuous operations and further minimise down-time. This would enable operators of potential fractionated systems to implement operational strategies which will allow them to maximise the operational time without having to launch unnecessarily large spacecraft.

6. References

- [1] ASD-EUROSPACE, “Space Trends: Global Space Activity Overview 1986-2011,” 2012. [Online]. Available: <http://eurospace.org/Data/Sites/1/pdf/spacetrends/eurospacetrends2011.pdf>. [Accessed: 03-Dec-2013].
- [2] H. Apgar, D. Bearden, and R. Wong, “Cost Modeling,” in *Space Mission Analysis and Design*, 3rd ed., W. J. Larson and J. R. Wertz, Eds. Space Technology Library, 2006, pp. 783–820.
- [3] M. N. Sweeting and C. I. Underwood, “Small-Satellite Engineering And Applications,” in *Spacecraft Systems Engineering*, Third Edit., P. Fortescue, J. Stark, and G. Swinerd, Eds. Wiley, 2006, pp. 581–612.
- [4] G. F. Dubos, J.-F. Castet, and J. H. Saleh, “Statistical reliability analysis of satellites by mass category: Does spacecraft size matter?,” *Acta Astronaut.*, vol. 67, no. 5–6, pp. 584–595, Sep. 2010.
- [5] O. Brown and P. Eremenko, “Fractionated space architectures: A vision for responsive space,” in *4th Responsive Space Conference April 24-27, 2006, Los Anglese, CA*, 2006.
- [6] USAF, “Global Horizons: United States Air Force Global Science and Technology Vision,” 2013. [Online]. Available: <http://www.defenseinnovationmarketplace.mil/resources/GlobalHorizonsFINALREPORT6-26-13.pdf>. [Accessed: 03-Dec-2013].
- [7] P. Molette, C. Cougnet, H. Saint-Aubert, P. W. Young, R, and D. Helas, “Technical and Economical Comparison Betwen a Geostationary Space Platform and a Cluster of Satellites,” *Acta Astronaut.*, vol. 12, no. 11, pp. 771–784, 1984.
- [8] D. M. LoBosco, G. E. Cameron, R. A. Golding, and T. M. Wong, “The Pleiades fractionated space system architecture and the future of national security space,” in *Proc. of AIAA SPACE 2008 Conference & Exposition, AIAA-2008-7687, San Diego, CA*, 2008, pp. 9–11.
- [9] A. da Silva Curiel, L. Boland, J. Cooksley, M. Bekhti, P. Stephens, W. Sun, and M. Sweeting, “First results from the disaster monitoring constellation (DMC),” *Acta Astronaut.*, vol. 56, no. 1–2, pp. 261–271, Jan. 2005.
- [10] O. Brown, “Reducing risk of large scale space systems using a modular architecture,” 2004.
- [11] O. Brown, A. Long, N. Shah, and P. Eremenko, “System lifecycle cost under uncertainty as a design metric encompassing the value of architectural flexibility,” *AIAA Pap. 6023*, 2007.

- [12] O. Brown and P. Eremenko, “The value proposition for fractionated space architectures,” *AIAA Pap. 7506*, pp. 1–22, 2006.
- [13] O. C. Brown, P. Eremenko, and P. D. Collopy, “Value-Centric Design Methodologies for Fractionated Spacecraft: Progress Summary from Phase 1 of the DARPA System F6 Program,” *AIAA Pap. 6540*, no. September, pp. 1–14, 2009.
- [14] O. Brown, P. Eremenko, and C. Roberts, “Cost-Benefit Analysis of a Notational Fractionated SATCOM Architecture,” *AIAA Pap. 5328*, 2006.
- [15] C. Cougnet, A. Satellites, T. France, and B. Gerber, “IAC-10.D1.4.1 FRACTIONATED SATELLITE: A STEP TOWARDS FLEXIBILITY AND RESPONSIVENESS C.Cougnet,” pp. 1–8.
- [16] CNES, “SPOT 5.” [Online]. Available: <http://spot5.cnes.fr/gb/satellite/satellite.htm>. [Accessed: 22-Aug-2011].
- [17] P. Snoeij, E. Attema, M. Davidson, N. Flouri, G. Levrini, B. Rosich, and B. Rommen, “Sentinel-1, the GMES radar mission,” *2008 IEEE Radar Conf.*, pp. 1–5, May 2008.
- [18] M. Martin-Neira, J. Font, M. Srokosz, I. Corbella, and A. Camps, “Ocean salinity observations with SMOS mission,” in *Geoscience and Remote Sensing Symposium, 2000. Proceedings. IGARSS 2000. IEEE 2000 International*, 2000, vol. 6, pp. 2552–2554.
- [19] C. Woody, E. Shih, J. Miller, T. Royer, L. P. Atkinson, and R. S. Moody, “Measurements of Salinity in the Coastal Ocean: A review of Requirements and Technologies,” *Mar. Technol. Soc. J.*, vol. 34, no. 2, pp. 26–33, 2000.
- [20] J. L. Miller and S. Payne, “Remote sensing of coastal salinity: naval needs and developing capability,” *IGARSS 2000. IEEE 2000 Int. Geosci. Remote Sens. Symp. Tak. Pulse Planet Role Remote Sens. Manag. Environ. Proc. (Cat. No.00CH37120)*, pp. 2534–2536, 2000.
- [21] L. D. T. and A. U. Bindoff, N.L., J. Willebrand, V. Artale, A. Cazenave, J. Gregory, S. Gulev, K. Hanawa, C. Le Quéré, S. Levitus, Y. Nojiri, C.K. Shum, “2007: Observations: Oceanic Climate Change and Sea Level. In: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change,” 2007.
- [22] S. Zine, J. Boutin, P. Waldteufel, J.-L. Vergely, T. Pellarin, and P. Lazure, “Issues About Retrieving Sea Surface Salinity in Coastal Areas From SMOS Data,” *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 7, pp. 2061–2072, Jul. 2007.
- [23] I. S. Robinson, “SOES6017 Lecture 1 Notes: Introductory Ocean Remote Sensing (University of Southampton).” 2008.

- [24] H. J. C. Blume and B. M. Kendall, “Passive microwave measurements of temperature and salinity in coastal zones,” *Geosci. Remote Sensing, IEEE Trans.*, no. 3, pp. 394–404, 1982.
- [25] A. Camps, J. Font, J. Etcheto, V. Caselles, A. Weill, I. Corbella, M. Vall-Ilossera, N. Duffo, F. Torres, R. Villarino, L. Enrique, A. Julia, C. Gabarro, J. Boutin, E. Rubio, S. C. Reising, P. Wursteisen, M. Berger, and M. Martfn-Neira, “Sea surface emissivity observations at L-band: first results of the Wind and Salinity Experiment WISE 2000,” *IEEE Trans. Geosci. Remote Sens.*, vol. 40, no. 10, pp. 2117–2130, Jan. 2002.
- [26] R. Oliva, E. Daganzo-eusebio, Y. H. Kerr, S. Member, S. Mecklenburg, S. Nieto, P. Richaume, and C. Gruhier, “SMOS Radio Frequency Interference Scenario : Status and Actions Taken to Improve the RFI Environment in the 1400 – 1427-MHz Passive Band,” vol. 50, no. 5, pp. 1427–1439, 2012.
- [27] J. Font, G. Lagerloef, D. Le Vine, A. Camps, and O. Zanife, “The determination of surface salinity with SMOS-recent results and main issues,” in *Geoscience and Remote Sensing Symposium, 2003. IGARSS'03. Proceedings. 2003 IEEE International*, 2003, vol. 1, no. C, pp. 7–9.
- [28] G. Lagerloef, “Final Report of the First Workshop of the Salinity Sea Ice Working Group,” 1998. [Online]. Available: <http://www.esr.org/lagerloef/ssiwg/ssiwgrep1.v2.html>. [Accessed: 03-Dec-2013].
- [29] I. S. Robinson, “Microwave Radiometry,” in *Measuring the Oceans from Space*, 2004, pp. 318–383.
- [30] H. Park and Y.-H. Kim, “Microwave motion induced synthetic aperture radiometer using sparse array,” *Radio Sci.*, vol. 44, no. 3, pp. 1–12, Jun. 2009.
- [31] D. M. Le Vine, “Synthetic aperture radiometer systems,” *IEEE Trans. Microw. Theory Tech.*, vol. 47, no. 12, pp. 2228–2236, 1999.
- [32] CNES, “smos_final_new1.jpg.” [Online]. Available: http://smsc.cnes.fr/IcSMOS/photos/satellite/smox_final_new1.jpg. [Accessed: 05-Jul-2010].
- [33] A. Camps, J. Bara, I. C. Sanahuja, and F. Torres, “The processing of hexagonally sampled signals with standard rectangular techniques: application to 2-D large aperture synthesis interferometric radiometers,” *IEEE Trans. Geosci. Remote Sens.*, vol. 35, no. 1, pp. 183–190, 1997.
- [34] K. D. McMullan, M. A. Brown, M. Martín-neira, S. Member, W. Rits, S. Ekholm, J. Marti, and J. Lemanczyk, “SMOS : The Payload,” *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 3, pp. 594–605, 2008.
- [35] A. Camps and C. T. Swift, “A two-dimensional Doppler-Radiometer for Earth observation,” *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 7, pp. 1566–1572, Jul. 2001.

- [36] H. Park, S.-H. Kim, H.-J. Lee, N.-W. Moon, and Y.-H. Kim, “A Rectangular Array for Motion Induced Synthetic Aperture Radiometer,” *IGARSS 2008 - 2008 IEEE Int. Geosci. Remote Sens. Symp.*, pp. II–1160–II–1163, 2008.
- [37] H. Park and Y. Kim, “Improvement of a Doppler-Radiometer Using a Sparse Array,” *IEEE Geosci. Remote Sens. Lett.*, vol. 6, no. 2, pp. 229–233, Apr. 2009.
- [38] Y. Kerr, J. Font, P. Waldteufel, A. Camps, J. Bara, I. Corbella, F. Torres, N. Duffo, M. Vallillossera, and G. Caudal, “New radiometers: SMOS-a dual pol L-band 2D aperture synthesis radiometer,” *2000 IEEE Aerosp. Conf. Proc. (Cat. No.00TH8484)*, no. 33, pp. 119–128, 2000.
- [39] D. Burrage, M. . Heron, J. . Hacker, J. L. Miller, T. . Stieglitz, C. R. Steinberg, and A. Prytz, “Structure and influence of tropical river plumes in the Great Barrier Reef: application and performance of an airborne sea surface salinity mapping system,” *Remote Sens. Environ.*, vol. 85, no. 2, pp. 204–220, May 2003.
- [40] E. J. D’Sa, J. B. Zaitzeff, and R. G. Steward, “Monitoring water quality in Florida Bay with remotely sensed salinity and in situ bio-optical observations,” *Int. J. Remote Sens.*, vol. 21, no. 4, pp. 811–816, 2000.
- [41] ESA, “Australia_Calib_H_cl2_L.jpg.” [Online]. Available: http://www.esa.int/images/Australia_Calib_H_cl2_L.jpg. [Accessed: 08-Apr-2010].
- [42] NASA, “Aquarius overview.” [Online]. Available: <http://aquarius.nasa.gov/images/overview.pdf>. [Accessed: 30-Mar-2010].
- [43] S. Cole and A. Buis, “NASA’S ‘Age Of Aquarius’ dawns with launch from California,” 2011. [Online]. Available: http://www.nasa.gov/home/hqnews/2011/jun/HQ_11-181_Aquarius.html. [Accessed: 03-Dec-2013].
- [44] J. W. Johnson, R. W. Lawrence, M. C. Bailey, C. T. Swift, D. M. Pozar, S. D. Targonski, and D. M. LeVine, “Technology Requirements in Advanced Microwave Radiometry,” *A Technol. Defin. Study NASA ESTO*, 1998.
- [45] S. H. Yueh, R. West, W. J. Wilson, F. K. Li, E. G. Njoku, and Y. Rahmat-Samii, “Error sources and feasibility for microwave remote sensing of ocean surface salinity,” *IEEE Trans. Geosci. Remote Sens.*, vol. 39, no. 5, pp. 1049–1060, May 2001.
- [46] K. McMullan, “Personal Communication: Distributed Satellites.” 2007.
- [47] B. S. Schwarz, “Distributed Satellites,” 2008.
- [48] A. Moccia and G. Fasano, “Analysis of Spaceborne Tandem Configurations for Complementing COSMO with SAR Interferometry,” *EURASIP J. Adv. Signal Process.*, vol. 2005, no. 20, pp. 3304–3315, 2005.

- [49] J. K. Eyer, “PhD Thesis: A Dynamics and Control Algorithm for Low Earth Orbit Precision Formation Flying Satellites,” 2009. [Online]. Available: https://tspace.library.utoronto.ca/bitstream/1807/19186/6/Eyer_Jesse_K_200911_PhD_thesis.pdf. [Accessed: 03-Dec-2013].
- [50] B. Schwarz, A. R. L. Tatnall, and H. G. Lewis, “A Fractionated Satellite Approach to Coastal Salinity Measurement,” *61st Int. Astronaut. Congr.*, pp. 1–8, 2010.
- [51] NASA, “AMSR-E Instrument.” [Online]. Available: http://aqua.nasa.gov/about/instrument_amsr.php. [Accessed: 07-Sep-2012].
- [52] “EO Portal: Aqua.” [Online]. Available: <https://directory.eoportal.org/web/eoportal/satellite-missions/a/aqua>. [Accessed: 12-Apr-2014].
- [53] EUMETSAT, “MetOp Satellite.” [Online]. Available: <http://www.eumetsat.int/Home/Main/Satellites/Metop/index.htm?l=en>. [Accessed: 22-Aug-2011].
- [54] “EO Portal: Metop.” [Online]. Available: <https://directory.eoportal.org/web/eoportal/satellite-missions/m/metop>. [Accessed: 12-Apr-2014].
- [55] J. M. Lafleur, J. H. Saleh, D. Current, and R. Frequency, “Feasibility Assessment of Microwave Power Beaming for Small Satellites,” *Eng. Conf.*, no. July, pp. 28 – 30, 2008.
- [56] S. S. T. SSTL, “Microwheel 100SP-M,” *Technology*, vol. 44, no. 0. 2010.
- [57] SSTL, “Small Satellite Reaction Wheel,” 2013. [Online]. Available: <http://www.sstl.co.uk/getattachment/8ccf207d-1ada-4bfc-a21b-5c400a1f5d52/SmallWheel>. [Accessed: 03-Dec-2013].
- [58] SSTL, “MTR-30 Magnetorquer,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/189dd4b3-214d-4e4b-b52e-d666a7ea1410/MTR-30>. [Accessed: 03-Dec-2013].
- [59] SSTL, “Xenon Propulsion System,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/28baa271-9c8d-4d35-ae16-847cc024dd6e/Xenon>. [Accessed: 03-Dec-2013].
- [60] SSTL, “Power System 150,” 2013. [Online]. Available: <http://www.sstl.co.uk/getattachment/80c25f64-d725-4907-b942-a3df147cac16/Power-System-150>. [Accessed: 03-Dec-2013].
- [61] W. J. Larson and J. R. Wertz, *Space Mission Analysis and Design*. 2006.
- [62] C. Pearson and N. Russel, “ABSL Space Products: Product Update Space Batteries.” 2007.

- [63] J. P. W. Stark, “Electrical Power Systems,” in *Spacecraft Systems Engineering*, 2003, pp. 325–354.
- [64] SPECTROLAB, “Space Solar Panels,” 2010. [Online]. Available: <http://www.spectrolab.com/DataSheets/Panel/panels.pdf>. [Accessed: 03-Dec-2013].
- [65] S. Persson, S. Veldman, and P. Bodin, “PRISMA—A formation flying project in implementation phase,” *Acta Astronaut.*, vol. 65, no. 9–10, pp. 1360–1374, Nov. 2009.
- [66] SSTL, “S-Band Downlink Transmitter,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/e4c92d26-ee24-40d8-b21b-36822e93d10f/S-Band-4-Watt-Downlink-Transmitter>. [Accessed: 03-Dec-2013].
- [67] SSTL, “S-Band Transmit Filter.” 2009.
- [68] SSTL, “S-Band Receiver,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/689c7c33-52f7-4cad-9f30-8f60b54a61db/S-Band-Receiver>. [Accessed: 03-Dec-2013].
- [69] SSTL, “S-Band Receive Filter,” 2009.
- [70] SSTL, “S-Band Patch Antenna,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/6c88ca87-0802-45e7-bfb6-356336ff710a/S-Band-Patch-Antenna>. [Accessed: 03-Dec-2013].
- [71] SSTL, “S-Band High Power Amplifiers,” 2007.
- [72] SSTL, “X-Band Downlink Transmitter,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/39811ce4-df75-4456-8ffd-71eb2c63f729/X-Band-Transmitter>. [Accessed: 03-Dec-2013].
- [73] SSTL, “Antenna Pointing Mechanism High Gain X-Band Antenna,” 2013. [Online]. Available: <http://www.sstl.co.uk/Products/Subsystems/Downloads/Antenna-Pointing-Mechanism---Sales-Brochure>. [Accessed: 03-Dec-2013].
- [74] SSTL, “OBC750 LEO Flight Computer,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/d2595f55-5579-44a4-8408-058e7ef63b52/OBC750>. [Accessed: 03-Dec-2013].
- [75] SSTL, “High Speed Data Recorder,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/6d57a930-9854-4ba2-9197-5d193a6ee4bd/High-Speed-Data-Recorder--HSDR->. [Accessed: 03-Dec-2013].
- [76] C. W. Reynolds, “Flocks, herds and schools: A distributed behavioral model,” *ACM SIGGRAPH Comput. Graph.*, vol. 21, no. 4, pp. 25–34, Aug. 1987.
- [77] R. J. Hamann, J. Bouwmeester, and G. F. Brouwer, “Delfi-C3 Preliminary mission results,” in *Proc. Of the 23rd annual AIAA/Utah State University Small Satellite Conference, SSC09-IV-7, Logan, Utah, USA*, 2009.

- [78] “Network Topology.” [Online]. Available: http://en.wikipedia.org/wiki/Network_topology. [Accessed: 02-Aug-2011].
- [79] SSTL, “SGR-10 – Space GPS Receiver,” 2013. [Online]. Available: <http://www.sstl.co.uk/getattachment/cda26002-f017-4d5a-9920-2174e9c8c907/SGR-10>. [Accessed: 03-Dec-2013].
- [80] S. S. T. SSTL, “MIRAS-01 Inertial Sensor,” *IEEE Sensors Journal*, vol. 1, no. 4, pp. 332–339, 2001.
- [81] SSTL, “2-axis DMC Sun Sensor.” [Online]. Available: <http://www.sstl.co.uk/getattachment/dadecfcb-e1e0-4ddf-91c6-f046a5441be1/Sun-Sensor>. [Accessed: 03-Dec-2013].
- [82] SSTL, “Magnetometer Magnetometer,” 2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/3b64658e-a81d-46da-97c1-ee43c602b3da/Magnetomer>. [Accessed: 03-Dec-2013].
- [83] SSTL, “Rigel-L Star Tracker,” 15-Aug-2012. [Online]. Available: <http://www.sstl.co.uk/getattachment/15ebab75-8c95-4e32-865b-2e1a360728dd/Rigel-L>. [Accessed: 03-Dec-2013].
- [84] G. F. Dubos and J. H. Saleh, “Comparative cost and utility analysis of monolith and fractionated spacecraft using failure and replacement Markov models,” *Acta Astronaut.*, vol. 68, no. 1–2, pp. 172–184, Jan. 2011.
- [85] ESA, “Sentinels Overview,” 2013. [Online]. Available: http://www.esa.int/Our_Activities/Observing_the_Earth/Sentinels_overview. [Accessed: 26-Nov-2013].
- [86] ESA, “ESA declares end of mission for Envisat,” 2012. [Online]. Available: https://earth.esa.int/web/guest/news/-/asset_publisher/G2mU/content/good-bye-envisat-and-thank-you. [Accessed: 26-Nov-2013].
- [87] SSTL, “WIDE SWATH DMC MSI,” 2010. [Online]. Available: <http://www.sstl.co.uk/Products/Subsystems/Downloads/DMC-MSI>. [Accessed: 03-Dec-2013].
- [88] R. H. Nansen, “Wireless power transmission: the key to solar power satellites,” *IEEE Aerosp. Electron. Syst. Mag.*, vol. 11, no. 1, pp. 33–39, 1996.
- [89] H. Schaub, G. G. Parker, L. B. King, and J. June, “Challenges and Prospects of Coulomb Spacecraft Formation Control of the Astronautical Sciences Challenges and Prospects of Coulomb Spacecraft Formation Control,” *J. Astronaut. Sci.*, vol. 52, pp. 169–193, 2004.
- [90] C. Bombardelli and J. Pel’aez, “Ion Beam Shepherd for Contactless Space Debris Removal,” *J. Guid. Control Dyn.*, vol. 34, no. 3, pp. 916–920, 2011.

- [91] T. A. Meaker, “Product Assurance,” in *Spacecraft Systems Engineering*, John Wiley & Sons Ltd, 2003, pp. 549–579.
- [92] J.-F. Castet and J. H. Saleh, “Single versus mixture Weibull distributions for nonparametric satellite reliability,” *Reliab. Eng. Syst. Saf.*, vol. 95, no. 3, pp. 295–300, Mar. 2010.
- [93] G. Green, “GPS 21 primary satellite constellation,” *Navig. J. Inst. Navig.*, vol. 36, no. 1, pp. 9–24, 1989.
- [94] D. A. Coley, *An Introduction to Genetic Algorithms for Scientists and Engineers*. World Scientific, 1999.
- [95] J. C. Spall, *Introduction to Stochastic Search and Optimisation: Estimation, Simulation and Control*. Wiley-Interscience, 2003.
- [96] J.-F. Castet and J. H. Saleh, “Satellite and satellite subsystems reliability: Statistical data analysis and modeling,” *Reliab. Eng. Syst. Saf.*, vol. 94, no. 11, pp. 1718–1728, Nov. 2009.
- [97] W. Ferster, “DARPA Cancels Formation-flying Satellite Demo,” *Space News*, 2013. [Online]. Available: <http://www.spacenews.com/article/military-space/35375darpa-cancels-formation-flying-satellite-demo#.Ub8CgfmkpyW>. [Accessed: 25-Nov-2013].

Appendix A: Journal paper and conference abstracts

COASTAL SALINITY MEASUREMENT USING A DOPPLER RADIOMETER

Benjamin S Schwarz, Adrian RL Tatnall and Hugh G Lewis

Published in Advances in Space Research 50(2012) 1138-1149

Abstract:

Coastal salinity is characterized by large and variable salinity contrasts on relatively small scales. Measurements of salinity at a resolution compatible with these coastal regions on a regular basis would provide a rich source of information that could be used for a number of applications that have a fundamental bearing on the world's lifestyle. Doppler radiometry offers an approach to capture such measurements, as it reduces the number of required antennas needed to form an image, compared with an Interferometer type instrument.

In this work, a Doppler Radiometer type instrument on free-flying satellites is introduced. This approach removes the need for a physical connection between all the antennas, affords the system a degree of reconfigurability, yet is still able to provide data of sufficient resolution. A Y-shaped central hub (similar to the SMOS configuration) is employed with additional antennas mounted on free flying platforms surrounding the central hub. The additional baselines formed between the antennas of the free flying satellites and central hub as well as between the free flying satellites extend the u-v coverage beyond that of just the central hub. The spatial resolution of a Doppler Radiometer system with a Y-shaped hub with a SMOS configuration of antennas, with each arm extended by five 6m spaced free flying antennas would be of the order of 5km, when imaging from 800km.

This paper will present some initial results from a study into an instrument concept that could provide coastal salinity measurements at microwave wavelengths. The study focuses on antenna array design and on quantifying the improvement in spatial resolution available by using this method, and includes an investigation into the effects of the relative motion between the hub and the free flying satellites on the imaging. Further, whilst this paper focuses on the application of the Doppler Radiometer to salinity measurement, the techniques described are applicable to other applications of a passive microwave instruments.

COASTAL SALINITY MEASUREMENT USING A DOPPLER RADIOMETER

Benjamin S Schwarz, Adrian RL Tatnall and Hugh G Lewis

Presented at the 38th COSPAR Scientific Assembly 2010

Abstract:

Coastal salinity is characterized by large and variable salinity contrasts on relatively small scales. Measurements of salinity at a resolution compatible with these coastal regions on a regular basis would provide a rich source of information that could be used for a number of applications that have a fundamental bearing on the world's lifestyle. Doppler radiometry offers an approach to capture such measurements, as it reduces the number of required antennas needed to form an image, compared with an Interferometer type instrument.

In this work, we introduce a Doppler Radiometer type instrument on free-flying satellites. This approach removes the need for a physical connection between all the antennas, affords the system a degree of reconfigurability, yet is still able to provide data of sufficient resolution. A Y-shaped central hub (similar to the SMOS configuration) is employed with additional antennas mounted on free flying platforms surrounding the central hub. The additional baselines formed between the antennas of the free flying satellites and central hub as well as between the free flying satellites extend the u-v coverage beyond that of just the central hub. The spatial resolution of a Doppler Radiometer system with a Y-shaped hub with a SMOS configuration of antennas, with each arm extended by five 6m spaced free flying antennas would be of the order of 5km, when imaging from 800km.

This paper will present some initial results from a study into an instrument concept that could provide coastal salinity measurements at microwave wavelengths. The study focuses on quantifying the improvement in spatial resolution available by using this method, and includes an investigation into the effects of the relative motion between the hub and the free flying satellites on the imaging. Further, whilst this paper focuses on the application of the Doppler Radiometer to salinity measurement, the principles are equally applicable to any passive microwave instrument.

A FRACTIONATED SATELLITE APPROACH TO COASTAL SALINITY MEASUREMENT

Benjamin S Schwarz, Adrian RL Tatnall and Hugh G Lewis

Presented at the 61st International Astronautical Congress 2010

Abstract:

Coastal salinity is characterized by large and variable salinity contrasts on relatively small spatial scales. Measurements of salinity at a resolution compatible with these coastal regions on a regular basis would provide new data for the monitoring of a key climate change variable in an as yet unobserved region. This paper describes an approach to measuring coastal salinity using the concept of fractionated satellites. The fractionated concept, developed by Owen Brown and Paul Eremenko of the Defence Advanced Research Projects Agency (DARPA), uses physical decomposition of the spacecraft and information integration to achieve a space system that is much more flexible and adaptable than its monolithic equivalent. Soil Moisture and Ocean Salinity (SMOS), the current European Space Agency (ESA) mission to measure sea surface salinity, already utilises information integration at a payload level to combine the information from the individual receiving antennas to form the image onboard the spacecraft, before it is transmitted to the ground. A coastal salinity measurement system that requires flexibility and reliability can be achieved through physical decomposition of the payload and information integration. This presents significant challenges in terms of the data handling, onboard computing and communications architectures employed as well as the orbits employed by the individual free flying modules and the accuracy to which they are maintained. We present some initial discussion and results from a concept exploration of a space system to measure coastal salinity.

OPTIMISING FRACTIONATED SPACECRAFT

Benjamin S Schwarz, Adrian RL Tatnall and Hugh G Lewis

Presented at the 62nd International Astronautical Congress 2011

Abstract:

The concept of fractionated satellites has been explored for the last six years, with DARPA's System F6 program the most prominent of the research efforts. Research in this area has been focused largely on the programmatic and economic benefits of fractionation. However, there are also significant technical and operational challenges to fractionation that have not been studied as extensively. Consequently this paper looks at fractionated satellites from this underexplored point of view. In this work, a variety of fractionated architectures have been simulated taking into account the satellite and subsystem failures based upon different failure rate assumptions. The architectures are characterised by a few simple parameters, such as the extent of the fractionation of the system, the number of satellites in the formation, and the location of subsystems within the formation. Operational replacement and redundancy strategies for failed satellites have also been studied. The architectures are simulated to investigate the effects and operational implications of maintaining the fractionated system in orbit for period of time that far exceeds the lifetime of a conventional, monolithic spacecraft. This is done with the aim of maximising the percentage of the lifetime that the architecture is operational for. The outputs of these simulations are used to highlight particular fractionated architectures and operational concepts that will enable space systems to maintain a high level of operational time on orbit. Results show that this will be strongly influenced by the redundancy strategy employed, the failure rate used and the degree of fractionation of the system.

THE DESIGN OPTIMISATION OF FRACTIONATED SATELLITE SYSTEMS

Benjamin S Schwarz, Adrian RL Tatnall and Hugh G Lewis

Presented at the 63rd International Astronautical Congress 2012

Abstract:

The concept of fractionated satellites involves the decomposition of the traditional monolithic satellite into a system of free flying satellites with specific functions, such as to provide a high speed data relay to the ground. These satellites share resources to achieve the mission. To date, research into fractionated satellites has largely focused on quantifying the programmatic and economic benefits of implementing this concept for the next generation of space systems. However the significant technical and operational challenges to satellite fractionation have not been studied as extensively. This paper examines the many different configurations a fractionated satellite might take and how they might be operated once in space. A variety of fractionated satellite architectures have been simulated taking into account the satellite and subsystem failures based upon a bath-tub type failure rate curve. The architectures are characterised by parameters, such as the degree to which the system is fractionated, the number of satellites in the system, and the distribution of subsystems on the satellites. There are a very large number of combinations and so simulations have been performed to identify the optimum architecture over a long period of continuous operation. The objective of these simulations is to characterise how each fractionated architecture performs over a 50 year lifetime with respect to the failures. Each architecture is assessed with respect to two traits: the percentage of the lifetime that the payload and system operations can be maintained, and the total mass launched. A local search optimisation based on this assessment is used to highlight particular characteristics of fractionated architectures that maximise operational time whilst minimising the mass. Results show that the outcome of this optimisation will be strongly influenced by the redundancy strategy employed, the distribution of subsystems throughout the fractionated system and the degree of fractionation of the system.