CRANFIELD UNIVERSITY

ROSS BURGON

MANOEUVRE PLANNING ARCHITECTURE FOR THE OPTIMISATION OF SPACECRAFT FORMATION FLYING RECONFIGURATION MANOEUVRES

SCHOOL OF ENGINEERING

PhD

2010

SUPERVISOR – DR PETER ROBERTS

25TH MAY 2010

CRANFIELD UNIVERSITY

SCHOOL OF ENGINEERING

PhD THESIS

Academic Year 2009-2010

ROSS BURGON

Manoeuvre Planning Architecture for the Optimisation of Spacecraft Formation Flying Reconfiguration Manoeuvres

Supervisor – Dr Peter Roberts

May 2010

ABSTRACT

Formation flying of multiple spacecraft collaborating toward the same goal is fast becoming a reality for space mission designers. Often the missions require the spacecraft to perform translational manoeuvres relative to each other to achieve some mission objective. These manoeuvres need to be planned to ensure the safety of the spacecraft in the formation and to optimise fuel management throughout the fleet. In addition to these requirements is it desirable for this manoeuvre planning to occur autonomously within the fleet to reduce operations cost and provide greater planning flexibility for the mission. One such mission that would benefit from this type of manoeuvre planning is the European Space Agency's DARWIN mission, designed to search for extra-solar Earth-like planets using separated spacecraft interferometry.

This thesis presents a Manoeuvre Planning Architecture for the DARWIN mission. The design of the Architecture involves identifying and conceptualising all factors affecting the execution of formation flying manoeuvres at the Sun/Earth libration point L2. A systematic trade-off analysis of these factors is performed and results in a modularised Manoeuvre Planning Architecture for the optimisation of formation flying reconfiguration manoeuvres. The Architecture provides a means for DARWIN to autonomously plan manoeuvres during the reconfiguration mode of the mission. The Architecture consists of a Science Operations Module, a Position Assignment Module, a Trajectory Design Module and a Station-keeping Module that represents a multiple multi-variable optimisation approach to the formation flying manoeuvre planning problem. The manoeuvres are planned to incorporate target selection for maximum science returns, collision avoidance, thruster plume avoidance, manoeuvre duration minimisation and manoeuvre fuel management (including fuel consumption minimisation and formation fuel balancing). With many customisable variables the Architecture can be tuned to give the best performance throughout the mission duration. The implementation of the Architecture highlights the importance of planning formation flying reconfiguration manoeuvres. When compared with a benchmark manoeuvre planning strategy the Architecture demonstrates a performance increase of 27% for manoeuvre scheduling and fuel savings of 40% over a fifty target observation tour.

The Architecture designed in this thesis contributes to the field of spacecraft formation flying analysis on various levels. First, the manoeuvre planning is designed at the mission level with considerations for mission operations and station-keeping included in the design. Secondly, the requirements analysis and implementation of Science Operation Module represent a unique insight into the complexity of observation scheduling for exo-planet analysis missions and presents a robust method for autonomously optimising that scheduling. Thirdly, in-depth analyses are performed on DARWIN-based modifications of existing manoeuvre optimisation strategies identifying their strengths and weaknesses and ways to improve them. Finally, though not implemented in this thesis, the design of a Station-keeping Module is provided to add station-keeping optimisation functionality to the Architecture.

ACKNOWLEGEMENTS

Much of the content of this research thesis would not have been possible without the support and assistance of a number of people. I'd first like to thank Jenny Roberts for the initial spark of inspiration and hard work in setting up the co-funding agreement that enabled the research to be performed and to her husband Peter Roberts for filling her shoes in her absence. I'd also like to acknowledge Finn Ankersen from the European Space Agency (ESA) and Stephen Kemble from EADS Astrium who also have provided support throughout and helped keep the direction of the research steady and relevant. Malcolm Fridlund (ESA), Anders Karlsson (ESA) and Glenn White (Open University) are also thanked for their time and feedback.

No research of this scale and over this length of time would be possible without the continued support from family and friends. I'd especially like to mention my long-suffering wife, Rachel and wonderful children, Isaac and Ava, who have helped me maintain my drive and motivation throughout the PhD process.

ACRONYMS

The list below contains all the most commonly used acronyms throughout this thesis. Other, less frequently used, acronyms can be found within the text and are not on the list.

(B)PAM	(Benchmark) Position Assignment Module
(B)SOM	(Benchmark) Science Operations Module
(B)TDM	(Benchmark) Trajectory Design Module
BCS	Beam Combiner Spacecraft
B-MPA	Benchmark Manoeuvre Planning Architecture
CCD	Charge Coupled Device
CR3BP	Circular Restricted Three Body Problem
DOF	Degrees of Freedom
DSS	Distributed Spacecraft Systems
ER3BP	Elliptical Restricted Three Body Problem
ESA	European Space Agency
ESKM-MPA	Embedded Station-keeping Module Manoeuvre Planning Architecture
FEEP	Field Emission Electric Propulsion
FOV	Field of View
GA	Genetic Algorithm (optimisation algorithm)
GADS	Genetic Algorithm and Direct Search (MATLAB toolbox)
GNC	Guidance, Navigation and Control
GPS	Global Positioning System
IPAM-MPA	Integrated Position Assignment Module Manoeuvre Planning Algorithm
ISKM-MPA	Integrated Station-keeping Module Manoeuvre Planning Architecture
ISOM-MPA	Integrated Science Operations Module Manoeuvre Planning Architecture
L_2	2nd Lagrange point in any three-body system
LEO	Low Earth Orbit
linTTN	Linear Three Telescope Nuller
LQG	Linear Quadratic Gaussian (control / controller)
LQR	Linear Quadratic Regulator (control / controller)
MID(M)	Manoeuvre Information Dissemination (Module)
MPA	Manoeuvre Planning Architecture
NASA	National Aeronautics and Space Administration
ODL	Optical Delay Line
PS	PatternSearch (optimisation algorithm)
RF	Radio Frequency
RT	Reference Trajectory
SepM-MPA	Separate Modular Manoeuvre Planning Architecture
S_{FP}	Spacecraft with Formation Planning capability
SimM-MPA	Simplified Modular Manoeuvre Planning Architecture

S_{IP}	Spacecraft with Individual Planning capability
SKM	Station-keeping Module
S_{NP}	Spacecraft with No Planning capability
STA	Science Task Assignment
TPF	Terrestrial Planet Finder
TPS	Target Point Strategy
triTTN	Triangular Three Telescope Nuller
TS	Telescope Spacecraft
WFC	Wide Field Camera

NOMENCLATURE

Science Operations Module

a	edge index number $\{1\ldots k\}$
c_j	task completion duration for task *j* (days)
D	mean task duration of a tour (days)
E	set with infinite number of elements and range $\{0\ldots3\}$ representing a number line
f	task prioritization flag $\{1\ldots\infty\}$
i, j	index numbers for sets $\{1\ldots\infty\}$
I	number of iterations of the science operations planning algorithm
k	number of edges in U_i
$l_{i,j}$	duration of the edges $\mathbf{v_{i,j}}$ (days)
L	tour duration (days)
$m_{i,j}$	manoeuver duration from task *i* to task *j* (days)
M_i	set of nodes ($\subset N$) representing all achievable tasks starting from n_i
n_i	node $i \in N$
N	set of nodes representing all science tasks
$\dot{O}$	mean rate of extra observation time gained using the SOM instead of the BP to plan tours (hrs/day)
O_i	set of nodes ($\subset N$) representing all tasks within field of view defined by β_i
p	number of edges in *T*
p_{waits}	number of 'wait' edges in *T*
r	random number generated from a normal distribution with zero mean and variable σ
R	task/time ratio of the tour (tasks/day)
R_{BP}	task/time ratio of a tour generated by the Benchmark Planner (tasks/day)
R_{SOM}	task/time ratio of a tour generated by the Science Operations Module (tasks/day)
$s(n_i)$	ecliptic longitude of star associated with node n_i (deg)
S	set of star positions represented in ecliptic longitude
t_i	time of task completion for node n_i
T	set of edges in tour
T_{target}	user-defined minimum tour duration (days)
$\mathbf{v_{i,j}}$	edge from n_i to n_j
V_i	set of all edges from n_i to all other nodes
$V_{i,sort}$	set of edges, U_i, sorted by increasing edge duration $l_{i,j}$
*V**	set of edges, $V_{i,sort}$, arranged with prioritized edges at the beginning
U_i	set of all edges from n_i to all nodes in M_i
$\mathbf{v}^{*k}_{a}$	edge representing the a^{th} element of *V**
W_i	set of edges connecting n_i to $n_j \in O_i$
$\dot{\alpha}$	angular rate of the field of view (deg/day)
$\beta(t_i)$	ecliptic longitude of anti-sun vector at time t_i (deg)

σ	standard deviation of the random number generator

Position Assignment Module

b	baseline distance (m)
F_0	inertial reference frame
F_F	final formation reference frame
F_S	initial formation reference frame
$\mathbf{F}$	unit vector of target star's position in F_0
g	acceleration due to gravity at sea-level = 9.80665 ms^{-2}
I_{sp}	specific impulse (s)
i	spacecraft identifier $i=\{BCS,TS1,TS2,TS3\}$
$\mathbf{i, j, k}$	orthonormal basis vectors for reference frames
J	cost of a PAM manoeuvre
m_i	mass of spacecraft i (kg)
mF_i	fuel consumed by spacecraft i during manoeuvre
mf_i	fuel mass of spacecraft i (kg)
$\dot{mf}_i$	fuel consumption rate of spacecraft i (kgs^{-1})
N	number of spacecraft in the formation
$\mathbf{r_F}$	translation vector of F_F from F_0
$\mathbf{r_{0,i}}$	initial position vector of spacecraft i in F_0
$\mathbf{r_{f,i}}$	final position vector of spacecraft i in F_0
$\mathbf{r_{traj,i}}$	trajectory of spacecraft i
$\ddot{\mathbf{r}}_\mathbf{i}$	acceleration vector of spacecraft i in F_0 (ms^{-1})
$\mathbf{S}$	until vector of starting star's position in F_0
t	time within a manoeuvre (s),$t_0 \le t \le t_f$
t_0	manoeuvre start time (s)
t_i	time for spacecraft i to complete a trajectory based on a bang-bang thrust profile (s)
t_f	manoeuvre end time / formation manoeuvre time (s)
t_β	manoeuvre duration of spacecraft β / formation manoeuvre time (s)
$t_{\omega,i}$	thrust pulse width of spacecraft i (s)
$\mathbf{u_i}$	unit force vector (N)
X	independent variable for the PAM optimisation routine, $X=\{r_F, \theta_1\}$
$y_i(t)$	position of spacecraft i along trajectory at time t
α_i	thrust saturation limit of spacecraft i (N)
β	spacecraft identifier indicating the spacecraft with the longest manoeuvre time
γ	a proportionality constant = $1/I_{sp}g$ (sm^{-1})
ΔV_i	change in velocity capability of spacecraft i (ms^{-1})
θ_1	relative angular position slot of TS1 around the BCS in F_F
θ_F	angle between F_0 and F_F (rads)
θ_S	angle between F_0 and F_S (rads)
μ_{PAM}	fuel minimising / fuel balancing weight for the PAM
Π_S	initial formation frame

Π_F final formation frame

Trajectory Design Module

F_C spacecraft reference frame

J_1 fuel optimisation cost of the TDM manoeuvre

J_2 total cost of the TDM manoeuvre

$J_C(i,j)$ collision cost for the TDM manoeuvre between spacecraft *i* and spacecraft *j*

$J_P(i,j)$ plume impingement cost for the TDM manoeuvre between spacecraft *i* and spacecraft *j*

$mF_{par,i}$ fuel consumed during the manoeuvre by the parallel thrust profile (kg)

$mF_{per,i}$ fuel consumed during the manoeuvre by the perpendicular thrust profile (kg)

$\mathbf{r}_j^{nom}(t)$ parallel thrust plume cone direction vector for spacecraft *j* in F_0 at time *t*

$\mathbf{r}_j^{per}(t)$ perpendicular thrust plume cone direction vector for spacecraft *j* in F_0 at time *t*

$\mathbf{T_{par,i}}$ the thrust component parallel to the spacecraft trajectory for spacecraft *i* (N)

$\mathbf{T_{per,i}}$ the thrust component perpendicular to the spacecraft trajectory for spacecraft *i* (N)

$t_{per,i}$ the execution time of the perpendicular thrust component for spacecraft *i* (s)

$t_{\omega per,i}$ pulse width of the perpendicular thrust component for spacecraft *i* (s)

Y independent variable for the TDM optimisation, $Y = \left\{\theta_{T,i}, \left|\mathbf{T_{per,i}}\right|, t_{\omega per,i}, t_{per,i}\right\} : i = \{1 \ldots N\}$

$\alpha_{per,i}$ thrust saturation limit in the direction perpendicular to the spacecraft trajectory (N)

$\theta_{T,i}$ angle from the reference axis to the perpendicular thrust component for spacecraft *i* (rads)

μ_{TDM} fuel minimising / fuel balancing weight for the TDM

Stationkeeping Module

Libration Point Definition

A_x, A_y, A_z amplitudes of a trajectory about a collinear libration point (using the 1st order solution to the equations of motion)

D distance between the primaries in the CR3BP

D1 distance from the barycentre to M1 in the CR3BP

D2 distance from the barycentre to M2 in the CR3BP

G the gravitational constant = 6.67300×10^{-11} m^3 kg^{-1} s^{-2}

M1 mass of the larger primary in the CR3BP

M2 mass of the smaller primary in the CR3BP

m mass of the spacecraft in the CR3BP

n constant angular velocity between *M1* and *M2* in the CR3BP

R position vector of the spacecraft from the barycentre in the non-dynamic rotating reference frame

R1 position vector of the spacecraft from M1

R2 position vector of the spacecraft from M2

r position vector of the spacecraft relative to a libration point in a rotating reference frame linearised about the libration point

t time (s)

X,Y,Z components of the vector **R**

x,y,z	components of the vector **r**
x	six-dimensional state vector of the spacecraft in the non-linear CR3BP
ρ	non-dimensionalised mass of M2 (i.e. ratio of M2 to total system mass)
ω_{xy}	in-plane frequency of a trajectory about a collinear libration point (using the 1st order solution to the equations of motion)
ω_z	out-of-plane frequency of a trajectory about a collinear libration point (using the 1st order solution to the equations of motion)

Reference Trajectory Generation

o, p, f	trajectory segment start and end points
p^*	intermediate point in trajectory segment
t_o, t_p, t_f	time at segment positions o, p and f respectively
U	pseudo-potential within the CR3BP
$\bar{u}$	vector describing the position and segment duration corrections to make
$\delta\mathbf{h}$	vector describing required corrections to the target point positions and segment durations
ΔV_P	velocity difference between points p and p* in the trajectory segment
$\Phi(t,t_0)$	state transition matrix

Target Point Strategy

J	cost of the station keeping manoeuvre
p_1, p_2, p_3	position deviations of the spacecraft from the reference trajectory at t_1, t_2, and t_3 respectively if the stationkeeping manoeuvre is not executed
$Q, R, R_v, S, S_v, T, T_v$	3x3 weighting matrices
t_0	execution time of the stationkeeping manoeuvre
t_1, t_2, t_3	future time along the reference trajectory
v_1, v_2, v_3	velocity deviations of the spacecraft from the reference trajectory at t_1, t_2, and t_3 respectively if the stationkeeping manoeuvre is not executed
ΔV	size of the stationkeeping manoeuvre executed at t_0

CONTENTS

1. INTRODUCTION

This thesis reports three years of research work into optimal manoeuvre planning for separated spacecraft interferometry missions. The initial remit for the research was to study optimal path planning techniques and strategies for formation flying manoeuvres at L_2 and in particular develop an engineering software simulator in support of the DARWIN mission design. The aim of this thesis is to provide a detailed report on this research and detail the requirements, development, implementation and output of the software simulator. Within the scope of this thesis a brief study of exo-planet finding techniques and mathematical optimisation techniques is also given. The research project builds on Cranfield University's interests in the fields of formation flying mission design and spacecraft dynamics and control but refines the focus of this interest towards L_2 and the formation flying manoeuvre planning environment.

The core of this research project is found in Chapters 6-0 which describe background, requirements, design and analysis of the various spacecraft formation flying manoeuvre planning issues and the methods developed to include them within the manoeuvre planning environment. The preliminary chapters of this thesis provide important background information to help assess the motivation, scope and objectives for the research which will be introduced in the following sub-sections of this chapter.

1.1 Research Motivations

The motivations for this research project are formed from the number of parallel interests of the stakeholders. The stakeholders for this research project are the author, Cranfield University, the European Space Agency (ESA) and EADS Astrium. Although the research remained wholly the responsibility of the author, the influences of the other stakeholders must be introduced to fully understand how the direction of the research progressed. This sub-section details the motivations of the stakeholders of the project and their interest in co-funding the research.

1.1.1 Formation Flying Research at Cranfield University

The Space Research Centre at Cranfield University operates a common theme of Distributed Space Systems in its research interests. This theme encompasses spacecraft design and miniaturisation, formation flying dynamics and control, spacecraft autonomy, space applications and spacecraft end-of-life technologies (Cranfield University Space Research Centre website, 2009). Since 2000 the Centre has been researching formation flying dynamics and control initially through the MUSTANG project. The Multi-University Space Technology Advanced Nano-satellite Group (MUSTANG), collaboration between Cranfield University, the University of Southampton and EADS Astrium, designed a technology demonstrator mission, MUSTANG-2, that involved a two-spacecraft formation with an aim to demonstrate formation flying techniques and space-qualify various experimental payloads (Roberts, Bowling and Hobbs, 2002). Initial formation flying dynamics modelling, at the Space Research Centre, centred on the low Earth orbit (LEO) MUSTANG-2 mission (Izzo,

2002). Roberts and Roberts (2004) continued this theme creating high fidelity models for relative motion and control algorithms to simulate formation flying in LEO. Roberts also expanded the theme with the development of a relative motion model for formation flying around L_2 (Roberts, 2004) and an eventual publication of a thesis "*assessing the feasibility of achieving high precision formation flying of a fleet of satellites in both the Low Earth Orbit and Lagrange point (L_2) environments*" (Roberts, 2005). It is at this point the author joined the research group at the Space Research Centre to continue the formation flying research within the group.

1.1.2 PhD Co-funding Agreements

This research project was co-funded under ESA's Networking and Partnership Initiative (ESA News website, 2006) by ESA, EADS Astrium and Cranfield University. A common programme of research was defined to examine "*optimal path planning techniques and strategies for formation flying reconfiguration manoeuvres in L_2"* (Cranfield University Statement of Work, 2006). The research proposal covered many aspects of the spacecraft formation flying research field including investigating solutions to multi-manoeuvre and formation reconfiguration problems, relative dynamics models for spacecraft motion, optimal guidance strategies for fuel balancing, discrete and continuous actuation systems, fault tolerant system design, sensor data fusion, GNC system error propagation and frequency domain limiting perturbations. In addition, the ESA component of the statement of work required the development of a software simulator, the Optimal Path Planner for Formation Flying (OPAFF), "*that aims at the development, synthesis and optimisation of robust guidance and control algorithms for performing optimal reconfiguration manoeuvres of multi-spacecraft flying in formation such as DARWIN."* An additional requirement of this simulator was its implementation within the MATLAB® software environment (The Mathworks, 2006). The statement of work was influential in guiding the author towards a specific area of the spacecraft formation flying research field from the beginning. Through a combination of the literature review, the author's personal interests and regular contact with the co-funding stakeholders, the research project has evolved from the initial statement of work to what is presented in this thesis.

1.1.3 Author's Personal Interest

The final stakeholder in the research project is the author himself. Prior to joining the research group at the Space Research Centre the author knew little about the field of spacecraft formation flying and its complexities. Thus part of the author's motivation for taking part in this research project was to learn about this field and investigate the spacecraft formation flying manoeuvre planning obstacles and how to overcome them. Another driver concerns the research's potential application in support of the DARWIN mission. The author has a long-held interest in astronomy and in particular extra-solar planet hunting (Burgon, 1997).

1.2 DARWIN

The motivations of the various stakeholders to the research project have been introduced in the previous sub-section. Before the problem statement for the research project is presented the baseline mission for the research, DARWIN, must be briefly introduced. A more in-depth account of the various mission concepts, and of formation flying, can be found in Chapters 2 and 0. This sub-section simply gives a summary of the concepts required to understand the problem.

DARWIN is a European Space Agency (ESA) mission designed for the search, detection and characterisation of Earth-like planets around other stars. It uses the technique of nulling interferometry to achieve this aim and requires the use of multiple spacecraft. These spacecraft must fly in tightly controlled formations to achieve the resolutions required. For each different star a different formation pattern is required so the spacecraft must be able to change their positions relative to each other. Furthermore, the spacecraft are to be placed at the L_2 point of the Sun/Earth system. L_2 is a dynamically unstable point in space, 150 000 km from the Earth where the centripetal and gravitational forces of the system are in equilibrium. For DARWIN to be successful requires the autonomous and timely generation of safe and optimal spacecraft manoeuvres.

1.3 Problem Statement

In this sub-section the problem relating to formation flying manoeuvre planning is characterised so that a set of aims and objectives for the research project can be defined.

1.3.1 Spacecraft Formation Flying Manoeuvres

One of the benefits of using multiple spacecraft to perform a mission is their ability to change their positions relative to each other and hence the way they operate as a unified system. These relative translational manoeuvres pose a significant risk to the safety of the individual spacecraft and to the longevity of the mission. The risks to the spacecraft include collisions between spacecraft and thruster exhaust damage from nearby spacecraft. The risks to the mission longevity include the length of time it takes to perform manoeuvres, the amount of fuel consumed during each manoeuvre and the distribution of fuel amongst the fleet. The first two mission risks are self-explanatory. Manoeuvre duration, and the frequency of manoeuvres, can directly affect the size of the science return for a mission when science cannot be performed whilst the manoeuvres are taking place. Manoeuvre fuel consumption is also an important factor since running out of fuel will terminate the mission. Wasting spacecraft fuel on unnecessary, or unnecessarily large, manoeuvres may prematurely end the mission. The final mission risk involves the distribution of the fuel amongst the formation spacecraft. Though the spacecraft will work together within the formation to achieve the mission goals they manoeuvre independently from each other and are likely to use differing amounts of fuel performing manoeuvres. Over time the fuel imbalance between formation members will increase and may lead to single/multiple spacecraft fuel starvation whilst the other spacecraft have plenty of fuel remaining. Since the formation relies on multiple

spacecraft the loss of one or more may impair the functionality of the formation or prematurely end the mission. In order to mitigate the risks posed by formation flying manoeuvres it is essential that the manoeuvres are planned and optimised to ensure the safety and durability of the mission.

1.3.2 Formation Flying at Lagrange Points

Station-keeping is a common phenomenon for geostationary spacecraft that require periodic east-west and north-south manoeuvres to maintain their position within their orbital slots. LEO spacecraft require station-keeping manoeuvres to counteract atmospheric drag and other orbital perturbations. Station-keeping is also required by any spacecraft attempting to remain within the vicinity of a collinear Lagrange point since the dynamic environment around these points are unstable. These periodic manoeuvres ensure that the spacecraft can track their planned trajectory through space and various methods exist for planning and optimising these types of manoeuvre. However, formation flying spacecraft at Lagrange points faces additional problems.

In addition to the station-keeping manoeuvres the spacecraft must also perform formation-keeping manoeuvres. Depending on the separation distances of the spacecraft the local dynamical environment at L_2 may cause the spacecraft relative positions to drift significantly. If the mission requires the relative positions of the spacecraft to remain fixed then formation-keeping manoeuvres by each spacecraft need to be performed to counteract the drift. The imbalance of fuel consumption by individual spacecraft performing formation-keeping manoeuvres compounds the problem of managing the fuel distribution amongst fleet members.

Current Lagrange point station-keeping planning and optimisation methods have been designed for the single spacecraft environment. For a formation flying mission, the required formation and formation-keeping manoeuvres will perturb the spacecraft far more than the natural dynamic environment. This means the spacecraft will move through space along trajectories not planned by the station-keeping planners leading to the requirement for larger station-keeping manoeuvres as the mission progresses. In order to mitigate the risks posed by the dynamic environment of L_2 it is essential that manoeuvres are planned and optimised to ensure the spacecraft remain within the vicinity of the L_2 point.

1.3.3 Autonomy

For spacecraft the term autonomy describes its ability to act independently from operator control. This allows the flight operations phase of a space mission to be much less costly and more streamlined as less operations support on the ground is required. One area where autonomy has been used for a long time is within the Attitude Determination and Control System (ADCS) where closed-loop feedback is used to autonomously maintain a desired spacecraft pointing direction with no input from human operators on the ground. More recently, autonomy is being employed within the operations sub-system with the autonomous scheduling of payload and other sub-systems operations to optimise mission operations.

For formation flying missions, autonomy is a required element for the guidance, navigation and control (GNC) sub-system (see Chapter 5). Analogous to the ADCS, closed-

loop feedback systems are required for the formation to maintain desired relative positioning with no human input from the ground. Autonomy can also be ascribed to the reconfiguration manoeuvres and their planning and optimisation. If the formation can plan its own manoeuvres and manoeuvre schedule then ground operations are significantly simplified, restricted communications windows can devote more bandwidth to the science data and the complexities of one-way-light-time (OWLT) can be ignored.

On-board GNC autonomy for formation flying missions is a much more complex area than ADCS requiring inputs from multiple internal and external sensors, inter-spacecraft and (periodically) ground station communications links and a dedicated planning and optimisation platform. This requires computer resources that can tax even ground-based computer systems. Any algorithms therefore that enable levels of autonomy within the GNC sub-system must be complex enough to be able to process all the available data whilst also simple enough produce actionable manoeuvre commands within a short time frame.

1.4 Research Aims and Objectives

The ultimate aim of this research project is

To design optimal manoeuvre planning algorithms for use with separated spacecraft interferometry missions at L_2 (but specifically in support of the DARWIN mission) to enable the safe execution of formation flying reconfiguration manoeuvres. Planning these manoeuvres should allow the maximum science return to be realised for the mission through a combination of schedule optimisation, manoeuvre optimisation and optimal fuel management across all spacecraft in the formation. The planning algorithms should also be of sufficient simplicity to enable their inclusion as part of an on-board autonomous guidance, navigation and control sub-system.

Within this aim there are a number of objectives stated below:

- Investigate spacecraft formation flying manoeuvre planning and related spacecraft formation flying fields including planning architectures, control co-ordination and control strategies.
- Examine the dynamic environment for single and multiple spacecraft within the vicinity of L_2 in the Sun/Earth-Moon system.
- Assess the field of optimisation strategies and techniques.
- Analyse how these investigations can be made applicable for separate spacecraft formation flying manoeuvre planning with DARWIN as the reference mission.
- Design a manoeuvre planning architecture for a formation flying manoeuvre planning algorithm that incorporates observation scheduling, fuel managed manoeuvre planning and manoeuvre safety planning.
- Evaluate the algorithms using a number of DARWIN-like manoeuvre examples and comparison planning algorithms to assess the requirements, optimality and expediency in relation to manoeuvre planning.

- Develop the algorithms into a software simulator that can be used as a mission design tool for separated spacecraft interferometry at L_2 formation flying manoeuvre planning.

1.5 Research Contributions

Spacecraft formation flying research has been very active for over the past 10 years with much effort directed towards the guidance, navigation and control aspects of formation flying missions. The research performed over the course of this project has provided a greater insight into formation flying manoeuvre planning at mission level through the development of a multi-objective optimal manoeuvre planning architecture that encompasses many of the single objective manoeuvre planning optimisation problems tackled thus far in the literature. In addition, this architecture includes a module to optimally plan observation schedules (specifically for DARWIN) incorporating a complex set of temporal pointing and mission constraints. Finally, the proposed inclusion of L_2 station-keeping within the manoeuvre planning architecture adds the unique dynamic environment of the Lagrange point to the planning process and provides a platform to gain insight into the affect of station-keeping manoeuvres on future formation manoeuvres and vice-versa.

1.6 Thesis Overview

The content of this thesis is broadly separated into three sections. In the first section, covering Chapters 1-5, the background of the various aspects of this research project is presented. Chapter 2 begins with a brief synopsis of the search for extra-solar planets, one of the driving motivations for this project, before giving an in-depth look at the DARWIN mission. Chapter 3 provides an overview of the field of spacecraft formation flying first through clarifications of the terms used in the literature and then through a review of formation flying missions and enabling concepts. Chapter 4 includes a précis of mathematical optimisation and the different types of optimisation strategies and techniques employed today. The first section of this thesis ends with Chapter 5 where the various concepts introduced in the previous three chapters are discussed and a selection is made regarding the formation flying concepts and constraints that will be implemented within the design of the manoeuvre planning architecture.

The middle section of this thesis (Chapters 6-10) details the design of a manoeuvre planning architecture for the optimisation of formation flying reconfiguration manoeuvres. Chapter 6 re-iterates the formation flying manoeuvre planning problem statement before introducing a number trade-offs that can be implemented to simplify the problem. These simplifications result in four separate ‘optimisation modules’ that are crafted into a Manoeuvre Planning Architecture (MPA). Chapters 7-10 detail the development of each MPA ‘module’ introduced in Chapter 6. Each of these chapters begins with a background and literature review of the optimisation problem they are designed to solve before providing full developmental details, comparison and analysis and a discussion of the limitations and improvements that could be made. The middle section of the thesis ends with Chapter 11, analysing the performance of the MPA ‘modules’ within the chosen MPA.

The final part of this research thesis, Chapter 12, provides a summary of the main research outcomes and conclusions are drawn. Chapter 12 concludes with a discussion into future work proposals to further this research field and enhance the manoeuvre planning software.

1.7 Publication History

Prior to publication of this thesis some of the research documented has been published and presented at conferences. Below is a list of these papers.

1.7.1 Journal Papers

- Burgon, R., Roberts, P.C.E., Roberts, J.A. and Ankersen, F. "Manoeuvre Planning Optimisation for Spacecraft Formation Flying Missions", *Journal of the Astronautical Sciences*, Vol. 56, Issue 4, pp. 545-571, Oct-Dec 2009.
- Burgon, R., Roberts, P.C.E., Roberts, J.A. and Ankersen, F. "Science Operations Planning Optimization for Spacecraft Formation Flying Maneuvres" *Journal of Spacecraft and Rockets*, Vol. 46, No. 3, pp. 634-644, May-June 2009.

1.7.2 Conference Papers

- Burgon, R., Roberts, P.C.E. and Ankersen, F. "Two-Stage Optimal Manoeuvre Planning for Spacecraft Formation Flying Missions", *The 3rd International Symposium on Formation Flying, Missions and Technologies*, Noordwijk, The Netherlands, 23–25 April, 2008.
- Burgon, R., Roberts, P.C.E. and Ankersen, F. "Optimal Autonomous Manoeuvre Planning for Spacecraft Formation Flying – Position Assignment", *Advances in the Astronautical Sciences*, Vol. 130, 2008, pp. 1015-1032.
- Burgon, R., Roberts, J.A. and Roberts, P.C.E., "Optimal Path Planning for Spacecraft Formation Flying: Planning Architecture and Operations", *Advances in the Astronautical Sciences*, Vol. 129, 2007, pp. 2685 – 2704.

2. THE SEARCH FOR EXTRA-SOLAR PLANETS AND THE DARWIN MISSION

One of the driving motivations for this research project is its support of a mission that is designed to advance the search for Earth-like extra-solar planets. In this chapter a review of the search for extra-solar planets is presented including scientific motivation, techniques and current and envisaged space missions. The DARWIN mission is also introduced in this chapter and specific details relating to the formation flying guidance, navigation and control (GNC) are given.

2.1 The Search for Extra-Solar Planets

For over 15 years ground-based and space-based astronomical instruments have been using the latest detector technology and astronomy techniques to discover planets around other stars in our galaxy. These extra-solar planets (or exoplanets) have forced astrophysicists to re-evaluate their theories of planetary formation and given astrobiologists the hope that life indicating molecules may be discoverable within exoplanetary atmospheres. This sub-section will briefly introduce the ideas surrounding the search for exoplanets, provide a review of the current techniques being employed in that search and introduce the current and future space missions tasked with continuing this endeavour.

2.1.1 Why Search for Exo-planets?

It has long been believed that our Solar System was not unique within the Universe but, up until recently, the technology to prove this had not been developed. Finding exoplanets can provide astronomers with vital clues to understanding the history and evolution of the planets in our own solar system but more interestingly provide clues and evidence of life elsewhere in the Galaxy.

The nebular model of planetary formation (Zeilik and Smith, 1973) has been the accepted planetary formation model for some time. This model states that the origins of the solar system began in a giant molecular cloud. Due to density variations within the cloud part of it began collapsing under its own mass. This gravitational collapse, coupled with the conservation of angular momentum, resulted in the formation of the Sun surrounded by a disk of matter that eventually coalesced into the Solar System. This model was developed entirely on observations of our own planetary system. With the discovery of over 300 exoplanets within the last 15 years (Schneider, 2009), the nebular model is in need of revision since there are a number of anomalous planets that the model does not predict. One of these anomalies is the existence of ‘hot-Jupiters’, large gas giants orbiting very close to the parent star. The nebular model predicts that such large planets cannot form so close to their parent stars however there have been at least 50 planets found with masses greater than 1 Jupiter mass orbiting within 1 A.U. of the parent star. The continuing search for exoplanets will provide astrophysicists with many more planetary systems to examine and advance their theories of solar system formation and evolution.

Another motive for the search for exo-planets relates to extra-solar life detection. Life has evolved on the Earth (and possibly elsewhere in the Solar System) and so it would seem probable that life would evolve elsewhere in the Universe as well. Astrobiologists have coined the term 'habitable zone' to describe the area around a star where water could be present in its liquid form. The distance from the parent star of this habitable zone depends on the size of the star as in Figure 2-1. The life detection process starts with the search for terrestrial planets within the habitable zone of Sun-like stars. These stars are limited to types F, G, K and M with luminosity types IV and V (sub-giants and main sequence stars respectively). Any terrestrial planets detected within a habitable zone could have liquid water on its surface or in the atmosphere and this condition makes the planet ideal for harbouring life. Spectral analysis of a detected planet may reveal the presence of bio-marker molecules within the atmosphere. The existence of CH_4 (at 7-8μm), O_3 (at 9.6μm), CO_2 (at 15μm) and H_20 (at >17μm) (amongst others) serves as a compelling indicator for the existence of life on the planet.

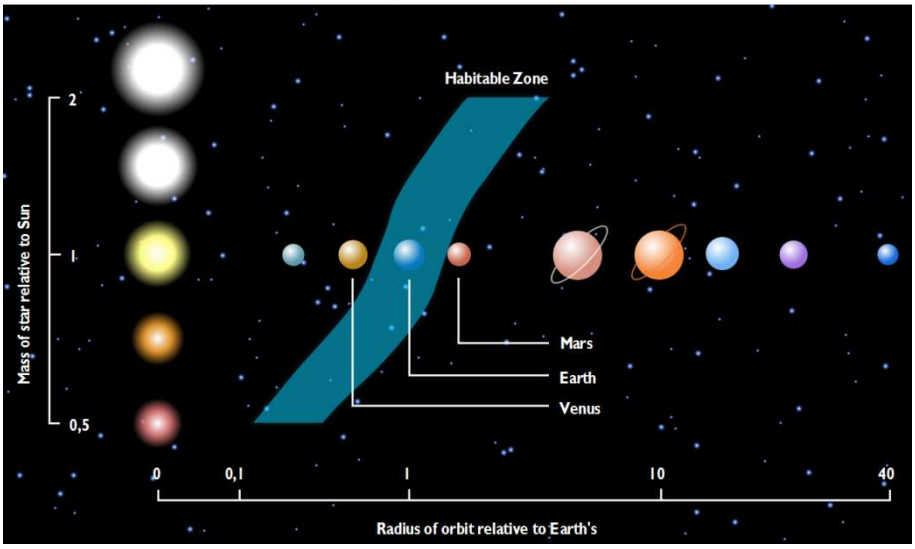

Figure 2-1 The stellar habitable zone for different stellar masses. The blue area shows the stellar habitable zone that moves further away from a star as the star's mass increases. (DLR website, 2009)

2.1.2 Main Exo-planet Detection Methods

The desire to search for extra-solar planets has been around since at least the days of Isaac Newton. The technology to directly observe these planets or observe their effects on their parent stars has not been available until fairly recently. Though only 15 years old the search for extrasolar planets has revealed an abundance of data and confirmed the realisation that these objects are very common indeed. Some of these methods are shown in Table 2-1.

Table 2-1 Exoplanet detection methods and associated space missions

Name	Method	Advantages	Disadvantages	Associated Space Missions
Radial velocity	measuring Doppler shift of star spectra	good at finding large planets close to the star with edge on orbital plane	no inclination data	GAIA (ESA GAIA website, 2009)
Astrometry	precise measurement of star's position in sky	good at finding large planets with face on orbital plane	no inclination data	GAIA (ESA GAIA website, 2009)
Transit	observed dimming of a star as a planet transits across the star's disk	estimates of the planet's size and true mass, atmosphere can be analysed during the transit, planet's infra-red radiation can be analysed during planetary eclipse	planetary alignment with observer's line-of sight needs to be perfect	HST (Charbonneau, et al., 2002) MOST (MOST Science website, 2009) CoRoT (CoRoT website, 2009) Kepler (Kepler website, 2009) GAIA (ESA GAIA website, 2009)
Microlensing	gravitational influence of a foreground star magnifies the light coming from a background star, if the lensing star has a planet then its gravitational influence can contribute to the lensing effect	good at identifying small planets, can show the position of the planet on the sky, gives accurate mass estimates	required star alignment cannot be repeated and can only identify planets several kilo-parsecs away	n/a
Direct Imaging	direct imaging	can see the planet	extremely difficult with today's technology	HST (Smith, 2008) TPF_C (TPF-C website, 2009) TPF-I (TPF-I website, 2009) DARWIN

2.2 The DARWIN Mission

DARWIN was ESA's proposed infra-red space interferometer consisting of a number of spacecraft required to fly in formation to achieve the desired resolution of the telescope. The mission was initially a proposal for ESA's Cosmic Vision 2015-2025 but failed to be selected for further review in October 2007 (ESA Cosmic Vision Website, 2009). The mission concept, however, remains important for ESA and studies are continuing due to a likely collaboration with NASA (Lawson, et al., 2007). Throughout its development the DARWIN mission has evolved from its initial proposals (Fridlund, 1999) to an ESA mission assessment study (Karlsson, et al., 2004). The most recent iterations however have been in the form of industrial mission assessment studies (Ruilier, Sghedoni and Krawczyk, 2007) and (Wallner, 2006). The most recent study at the commencement of this research project was Karlsson, et al. (2004) and it is this study that has been used as the baseline DARWIN mission for this research. The following sub-sections introduce the DARWIN mission and highlight the areas of the mission that affect the development of the manoeuvre planning algorithms found later in this thesis.

2.2.1 The Mission

The science objectives of the DARWIN mission are two-fold:

1. Search nearby stars for the existence of terrestrial planets orbiting within the star's habitable zone.
2. Analyse the detected planet's orbital, geophysical and atmospheric characteristics and look for the existence of bio-marker molecules within the atmosphere.

One of the difficulties faced with trying to directly observe extra-solar planets is the requirement to resolve the planet from the star. The DARWIN mission uses a technique called nulling interferometry to effectively cancel the stellar component of the radiation from the star system. After nulling all that remains in the signal are the emission and reflection signatures of the planet. Interferometry involves combining the signal (from the same source) from two (or more) spatially separated telescopes. The combined signals interfere with each other to produce an interference pattern. With nulling interferometry, by manipulating the position of the telescopes in three-dimensional space, the interference pattern can be changed. With the correct positioning of the telescopes the interference pattern can be manipulated so that the stellar component of the signal is removed. A more in-depth study of nulling interferometry and formation flying can be found in Roberts (2005).

The necessary separation distance of the telescopes for exo-planet interferometry is ~10-100 m. Using a conventional monolithic spacecraft to achieve these separations with multiple telescopes would be prohibitively costly and so multiple separate spacecraft are envisioned. The use of separate spacecraft also increases the flexibility of the interferometer. Since the spacecraft can be manoeuvred anywhere, any interferometry positioning requirements can be satisfied. The science requirement to perform spectroscopy in the infra-red assists in the

performance required by the nulling interferometer. The brightness contrast between a star and a planet is typically ~10^9 in the visible part of the spectrum. However, at infra-red frequencies this reduces to ~10^6. This means the size of the telescopes can be smaller as well as the required scale of the interferometer. Infra-red detectors require cooling to 8 °K and so a cryogenic cooling system is needed to achieve this.

Another aspect of the mission that assists the detector cooling is the placement of the spacecraft in trajectory about the L_2 libration point in the Sun/Earth-Moon system (see sub-section 10.1.1 for more information on libration points). With its position 1.5 million km away from the Earth along the Sun/Earth line, the L_2 point provides an ideal environment for astronomy missions. The DARWIN pointing constraints require the telescopes to always point within ±45° of the anti-Sun vector. At any one time therefore, 15% of the sky is available for observation, any one source stays within the field of view for ~91 consecutive days once a year and throughout the entire year 71% of the sky is available for observation. Compared to the Earth orbital environment the dynamic environment for trajectories around L_2 is relatively benign with very little manoeuvring required to maintain attitude pointing for very long periods of time. This will assist in maintaining the telescope pointing to within the required 24 milli-arcseconds. This benign dynamic environment will also help the spacecraft to maintain their relative positioning to within the required 1 cm and 250 μms^{-1} during observations.

2.2.2 Science Observations for the DARWIN mission

The two main types of science observation required for the DARWIN mission are planet detection and planet atmospheric spectroscopy. With this in mind Karlsson, et al. (2004) identified two formation configurations that would achieve this to the required resolution. For the planetary detection observations the preferred configuration is the linear three telescope nuller (linTTN) as shown in Figure 2-2. In the linTTN configuration the formation forms a co-planar ‘T’ shape. The Telescope Spacecraft (TS) are separated from each other by the desired baseline distance. This same baseline distance is used between the Beam Combiner Spacecraft (BCS) and the central TS. The size of the baseline depends on the target star’s spectral type and its distance from the formation. During the detection observation the entire formation is required to rotate 180° within the formation plane whilst maintaining the accuracy of the baseline to within 1 cm and remaining pointing at the target star within 24 milli-arcseconds. To achieve this requirement the manoeuvre is restricted to a maximum rotation speed of 360° per day.

For the planet atmospheric spectroscopy observations the formation forms a co-planar ‘Y’ shape dubbed the triangular three telescope nuller (triTTN) as shown in Figure 2-3. The BCS sits in the centre with the TSs separated from the BCS by the desired baseline distance. The baseline depends on the target star’s type, its distance from the formation and the wavelength of the spectroscopy being performed. The TS have a 120° separation within the formation plane in this example (and in this thesis), however Karlsson, et al. (2004) note that the TSs can be separated by any arbitrary angle.

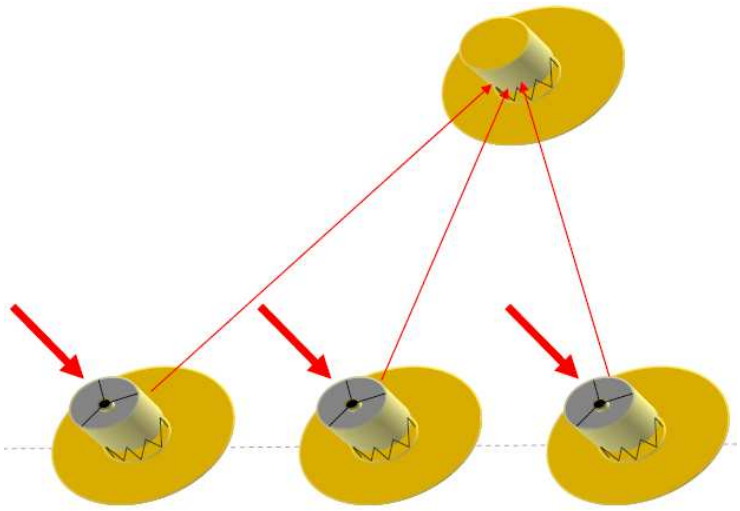

Figure 2-2 The linTTN configuration for DARWIN (Karlsson, et al., 2004)

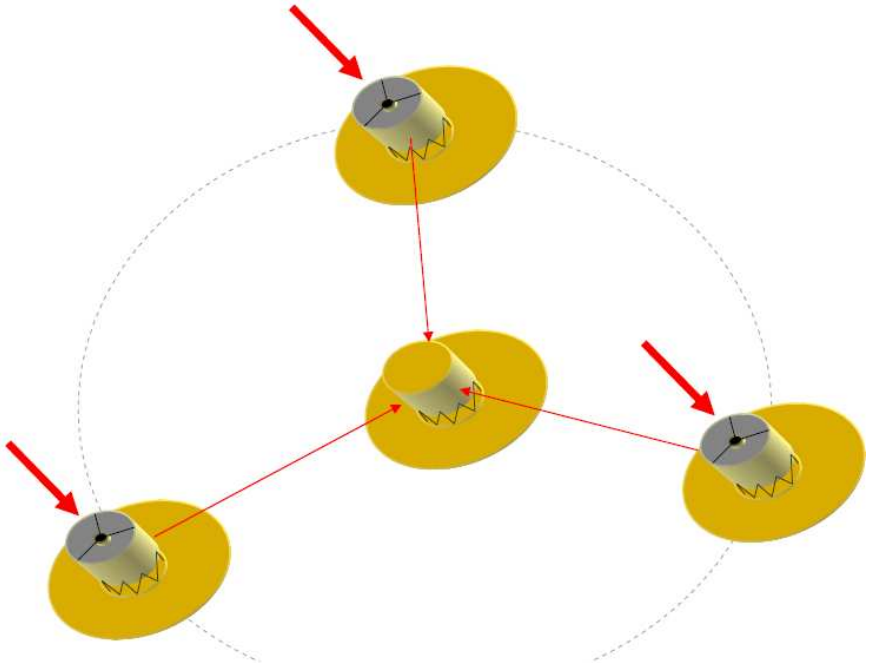

Figure 2-3 The triTTN configuration for DARWIN (Karlsson, et al., 2004)

As well as formation configuration there are other parameters that differ between the detection and spectroscopy observations. These are detailed in Table 2-2. The linTTN configuration is preferred for planet detection as it has higher signal modulation efficiency across the entire habitable zone allowing unambiguous planet detection using a 180° rotation. For planet detection to be confirmed, three separate detection observations must be performed on the star. These observations must be separated by enough time for any potential planet to appreciably move within its orbit. Three planet detections at three spatially different locations around the star provide confirmation of the planet's existence and gives insight into its orbital characteristics. Each detection observation requires an integration time of 1-4 days depending on the spectral type of the target star and the distance to the target star.

Table 2-2 Differences between science observations for DARWIN

Science Phase	Detection	Spectroscopy
Objectives	To detect exoplanets within the habitable zones of nearby stars	To analyse the 7-20 µm spectral lines of the planet atmosphere
Number of visits per star	3	1-6
Observation time per star	⅓-4 days	up to ~91 days
Formation configuration	Linear TTN	Triangular TTN
On-observation manoeuvres	180° rotation	None

For the spectroscopy observations the triTTN is preferred as it provides better imaging capabilities. Six separate spectroscopy observations are required corresponding to the six spectrum bands to be analysed. Staring spectroscopy (where the configuration remains stationary within the formation plane) is planned as it gives lower integration time than the alternative rotating spectroscopy (where the formation rotates within the formation plane during the observation). The integration times however are significantly greater than those for the detection observations with certain star/spectroscopy combinations requiring integration times longer than the star can remain within the field of view (FOV)[1]. Further details on the integration times for both detection and spectroscopy observations can be found in Table 7-2.

2.2.3 DARWIN Spacecraft

The version of the DARWIN mission introduced in Karlsson, et al. (2004) requires the use of four separate spacecraft. One spacecraft is the Beam Combiner Spacecraft (BCS) whilst the other three spacecraft are identical Telescope Spacecraft (TS). Figure 2-4 shows a diagram of the TS. The main physical characteristics of this spacecraft are the 3.15 m primary mirror and the 0.6 m secondary mirror, a wide-field camera for attitude determination, the 200 mm transmission optics assembly and the 10 m diameter sun shade. The TS operation is simple. The signal from the target star system is focussed by the primary/secondary mirror assembly before being relayed to the send telescope and transmitted to the BCS. Part of this signal is also used for the spacecraft attitude determination system. The modelled mass for the TS is ~900 kg. Figure 2-5 shows a diagram of the BCS. The main physical characteristics of this spacecraft are the transmission optics receive telescopes, the detector housing and the 10 m diameter sun shade. The BCS receives the three star signals from the TSs through the receive telescopes. Within the detector housing the optical paths are controlled to within 1 nm using optical delay lines. The individual beams then pass through an array of beam splitters, achromatic phase shifters and modulators before being combined within a number of single mode waveguides and passed to the detectors and spectrographs. The mass modelled for the BCS is ~1100 kg.

[1] Integration time is based on the desired signal to noise ratio (SNR) for each spectral line being observed. To accommodate the field of view duration restriction the required SNR may need to be reduced for certain star/spectral line combinations.

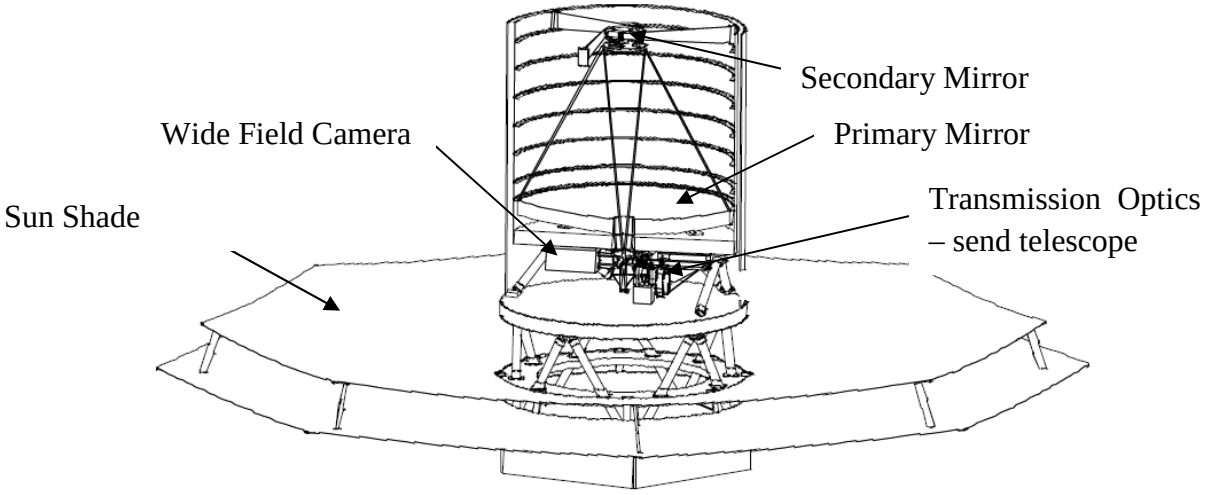

Figure 2-4 Schematic of the Telescope Spacecraft (TS) (Karlsson, et al., 2004)

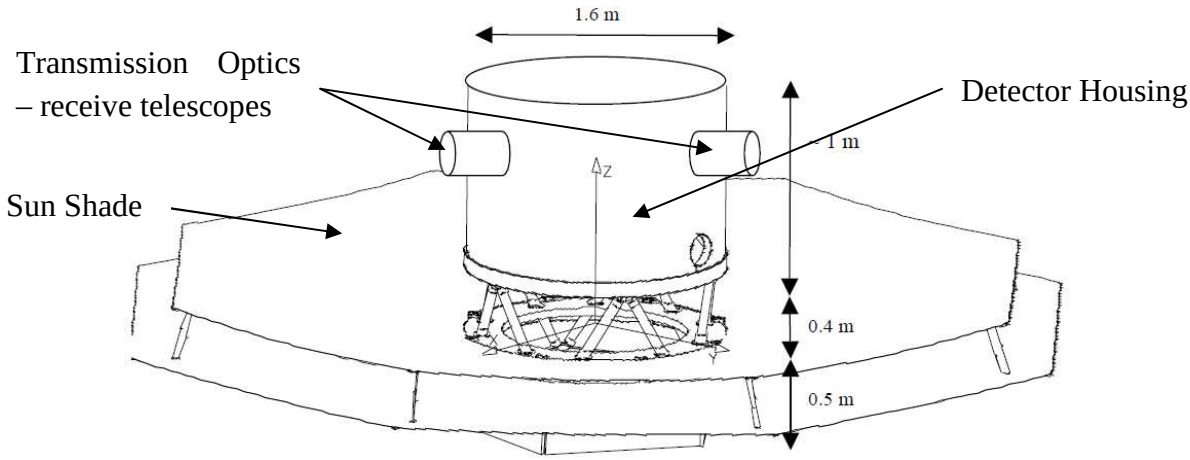

Figure 2-5 Schematic of the Beam Combiner Spacecraft (BCS) (Karlsson, et al., 2004)

For spacecraft formation flying the two most important spacecraft sub-systems to consider are the metrology and the control sub-systems. For the TS and the BCS the required pointing accuracy is <24 milli-arcseconds. This pointing accuracy is achieved using a variety of techniques. For attitude determination the TS employ a wide field camera (WFC) that uses part of the science signal to perform star tracking. This WFC will have an accuracy of 0.1 arcseconds. The BCS will carry a separate star tracker telescope capable of an accuracy of 1 arcsecond. In addition to these star trackers the formation will also employ laser attitude metrology to accurately determine pointing of the TS with respect to the BCS (spacecraft co-alignment). The BCS will use a laser to project a false star in the optical path of the TSs WFC. The WFC is capable of measuring this false star to an accuracy of ~13 milli-arcseconds. Finally, attitude determination data will also be available from the BCS fringe sensor. This device is used to measure the relative phase of the three science beams but can also measure the relative tilt of the beams with respect to an inertial reference. All spacecraft will also use coarse sun sensors and gyros for attitude determination during the non-observation phases of the mission.

The relative positions and velocities of the TS with respect to the BCS during observations are required to be within 1 cm and 250 μms^{-1} respectively. The optical path difference of the science beam needs to better than 5 nm. Like with the pointing metrology the positional metrology uses various different methods to achieve its aims. Coarse lateral and longitudinal positioning is provided by an RF system that operates similarly to GPS. This coarse positioning will provide accurate position determination to 0.2 cm and 0.05°. A coarse lateral laser metrology system will complement the RF system to provide an additional accuracy of 1 mm. For the laser metrology the BCS emits a beam towards the TS which is reflected within a corner cube on the TS back to the BCS. A CCD detector on the BCS then measures the lateral drift of the reflected beam to determine lateral displacement. The coarse metrology is used during the non-observation phases of the mission when the positional requirements of the spacecraft are less demanding. During the observation phases lateral and longitudinal position determination is provided by a fine laser metrology system. The fine longitudinal metrology system is similar to the coarse lateral system (i.e. is uses retro-reflectors on the TS) and can provide a position accuracy to within 32 µm. The fine lateral metrology system is slightly different. The BCS emits a beam towards the TS and the intensity profile is detected by a CCD. Comparing the intensity profile to a reference intensity profile provides the information necessary to achieve an accuracy of 32 µm. The final position metrology system is the BCS fringe sensor. This can measure the relative phase of the science beams to sub-nm accuracy.

In order for the stringent pointing and positioning requirements to be met DARWIN intends to employ advanced propulsion units. Coarse micro-propulsion, used for formation manoeuvres and station-keeping manoeuvres, will be provided by ion engines. These are able to provide a thrust range of 1-25 mN with a resolution of 1 µN. These engines provide the desired control for manoeuvres within the non-observation stages of the mission. Fine micro-propulsion is required for manoeuvres during the observation stages of the mission in order for the pointing and positional constraints to be maintained. This will be provided by FEEP (Field Emission Electric Propulsion) thrusters with a nominal range of 1-100 µN and a thrust resolution of <0.01 µN. For the spacecraft to be able to rotate and translate in any direction a total of 32 micro-propulsion units are required (16 mN and 16 µN units). In order to achieve the nanometre precision required for the science beam an optical delay line (ODL) is also used.

2.2.4 Observations Operational Cycle

There are many stages of the DARWIN mission that will require the spacecraft to perform manoeuvres (i.e. orbit correction manoeuvres en-route to L_2, station-keeping at L_2, formation reconfiguration manoeuvres and formation-keeping manoeuvres). Figure 2-6 shows a schematic for the GNC modes for DARWIN. As we are interested in the nominal operational modes of the mission we will only concentrate on the operational loop encompassing the baseline control mode, fringe acquisition mode, normal observation mode and reconfiguration mode. This loop describes the operation cycle for performing observations.

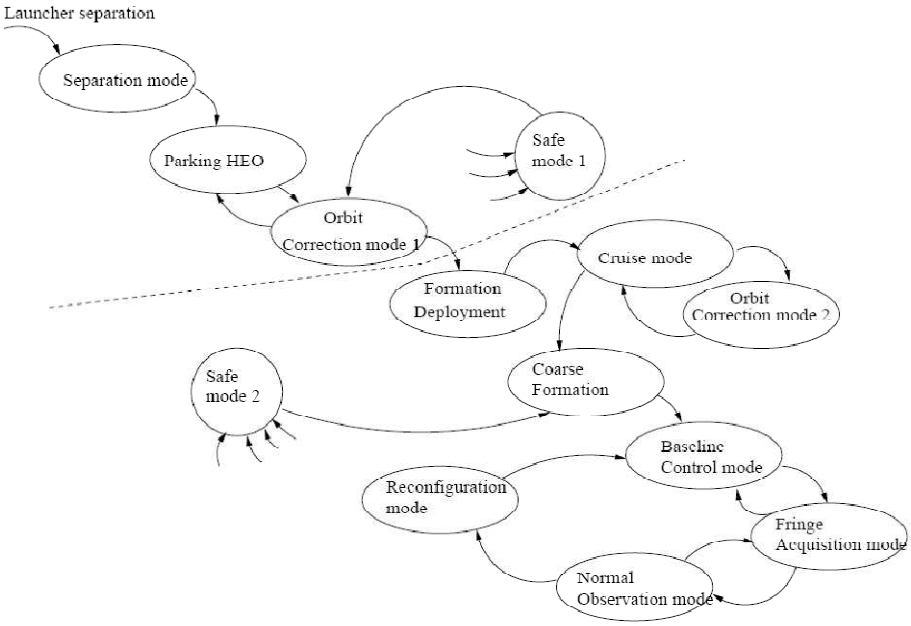

Figure 2-6 GNC Modes for the DARWIN mission (Karlsson, et al. 2004)

- Reconfiguration Mode - In this mode the spacecraft perform translational and attitude manoeuvres to achieve the required formation configuration (i.e. linTTN or triTTN) and pointing direction for the next observation. Pointing and positional accuracy in this mode are relaxed so only the RF metrology is required and the manoeuvres are performed using the milli-Newton thrusters.
- Baseline Control Mode – In this mode the relative positions of the spacecraft are stabilised to achieve better than 1 cm position accuracy. The RF and coarse lateral laser metrology systems are used along with the µN thrusters at achieve the desired baseline stability.
- Fringe Acquisition Mode – In this mode the central fringe of the interference pattern is found using the optical delay line (ODL). The full 2cm stroke of the ODL is used to find the central fringe. The ODL is then moved back to its centre position through µN thruster actuation. This allows the central fringe to be locked by the ODL when the ODL is in its central position. This gives the ODL full stroke capacity to maintain the central fringe lock during the observation. The fine laser metrology system is used in this mode along with the µN thrusters.
- Normal Observation Mode – In this mode the formation is performing science observations. The better than 5 nm optical path difference accuracy is maintained by the ODL whilst the spacecraft relative positions are controlled to better than 1cm. The fine laser metrology system and the µN thrusters are used to ensure both position and pointing accuracy. These metrology systems are also used during the required formation manoeuvre for detection observations. In order to maintain fringe lock and pointing accuracy during these manoeuvres the relative velocities of the spacecraft are constrained to a maximum of $3.7 mms^{-1}$. Once the observation is complete the reconfiguration mode is activated.

2.3 Chapter Summary

This chapter has provided a background to some of the motivational aspects for this research project: the search for extra-solar planets and the DARWIN mission. The first part of the chapter details the reasons why astronomers search for exoplanets, the types of methods used in the search, the information that can be yielded from each method and the space missions contributing to this branch of astronomy. Using a combination of detection methods can reveal a number of planetary characteristics including orbital elements, planetary mass and atmospheric composition. The requirements for the detection of Earth-like planets within a star's habitable zone have lead to the development of space missions like Kepler. For life detection studies however, direct imaging of exoplanets is required and so new and innovative techniques need to be employed for this to be realised. NASA's Terrestrial Planet Finder and ESA's DARWIN space missions represent the scale of the mission required to perform this type of astronomy.

The second part of this chapter introduced the DARWIN mission, spacecraft, science requirements and GNC requirements. This provides the background to why formation flying is required for this type of astronomy mission and introduces the aspects of the mission and GNC system that affect the design of the formation flying manoeuvre planning architecture that is the topic of this research project.

The next chapter introduces the concept of spacecraft formation flying and defines the term in relation to the GNC requirements. The chapter also reviews spacecraft formation flying applications and missions and the enabling concepts that make spacecraft formation flying missions so unique.

3. SPACECRAFT FORMATION FLYING

Distributed spacecraft systems (DSS) describe the use of multiple free-flying spacecraft working together to perform a unified mission objective. The use of DSSs to realise mission concepts is now very much a reality with examples including the Global Positioning System (GPS), Iridium, Cluster and GRACE. These DSSs are the enabling technology for an ever increasing array of spacecraft applications and offer a number of significant advantages over single spacecraft systems. The benefits of DSSs over single spacecraft include:

- mission enabler – some missions would be impossible to realise with only one spacecraft
- low cost - designing, building and launching multiple, identical spacecraft can be more cost effective than one larger spacecraft to achieve the same mission goals
- high redundancy – failure of a single spacecraft in a DSS may not be mission critical and a replacement used in its place
- easy expandability – adding further spacecraft to the formation at a later date to enhance its capabilities
- high resolution – multiple spacecraft can form very large synthetic apertures for higher resolution observations
- multiple observations – multiple spacecraft can view targets at more frequent intervals and simultaneously from various angles
- reconfiguration – ability to reconfigure the formation to perform different observations and/or be used in different missions

These benefits however are offset by an increase in operational complexity of the mission. This may involve more complex guidance, navigation and control (GNC) systems, more complex communications, an increased level of autonomy, the hardware required to realise these requirements and the increased requirements of space debris mitigation. The design of a DSS mission therefore can be a much more intensive endeavour than for a single spacecraft mission.

This chapter aims to provide an overview of spacecraft formation flying. Spacecraft formation flying will be defined to distinguish it from other DSS missions and a review of past, present and future missions and applications is presented. This is followed by a review of the enabling technology required for spacecraft formation flying missions.

3.1 Distributed Spacecraft Systems Definitions for Formation Flying

Distributed spacecraft systems (DSS) utilise multiple spacecraft to realise a common goal. There are, however, a number of different ways that the DSS can be operated and thus a number of different definitions for spacecraft formation flying have arisen. This, with the observed interchanging of terms like ‘constellation’, ‘formation’ and ‘cluster’, has made defining spacecraft formation flying a confusing endeavour. In this thesis the types of DSS are separated into two families; those that perform spacecraft formation flying and those

where the spacecraft fly in formation. The distinctions of these two DSS families are very different and help to focus the direction of this thesis. A chart showing these two DSS families, their sub-sets and examples of missions can be found in Figure 3-1.

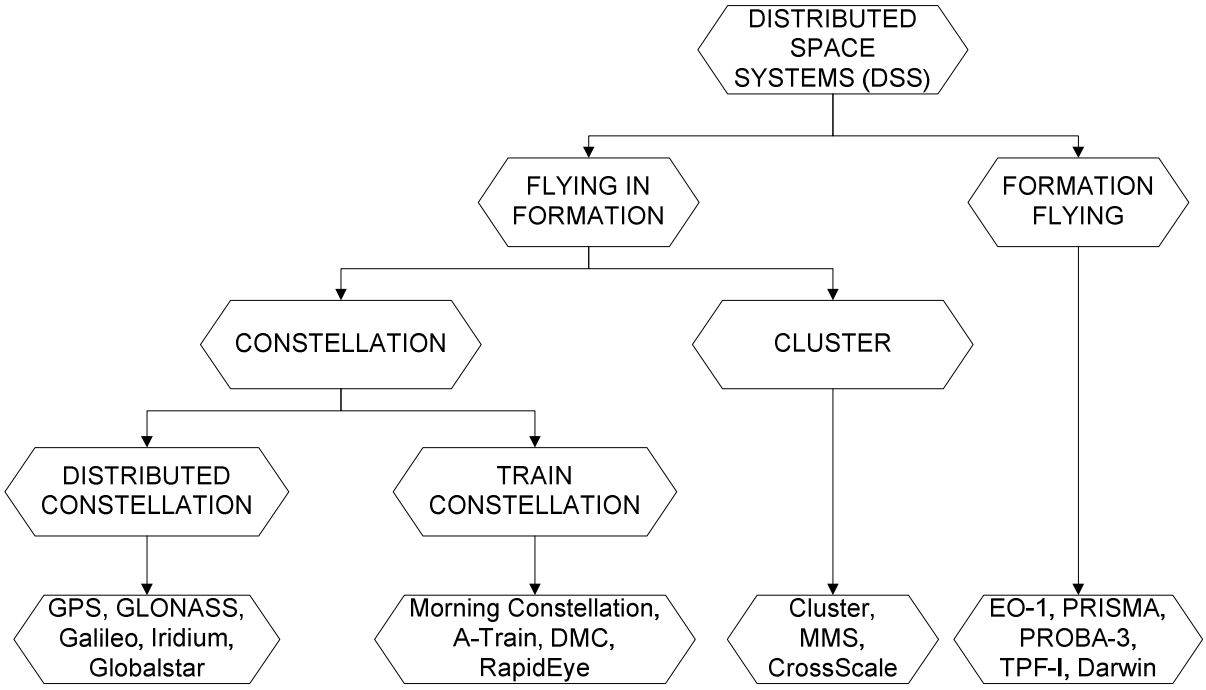

Figure 3-1 Distributed Spacecraft System (DSS) families and examples

3.1.1 Spacecraft Flying in Formation

The first type of DSS is where the spacecraft fly in formation. The formation can take any size or shape as required by the mission however the formation itself is a result of orbit/trajectory design rather than active formation-keeping control. The important aspect for this type of DSS is that each individual spacecraft is controlled independently from the ground. Therefore, although relative positioning of spacecraft within the formation may be key, that positioning is a result of individual orbit design and station-keeping. As an analogy consider the airspace around any modern airport. Aircraft operating within this airspace are separated by defined distances to ensure the safety of the aircraft in the air and their timely access to the airport itself. Although each aircraft has its own pilots flying the plane they are being controlled through air traffic controllers on the ground. Thus the pilots themselves are not making any decisions regarding aircraft separation distances and relative positioning.

Within this first DSS family there are a number of subsets that can describe the configuration of the DSS. The first is the constellation, where numerous spacecraft form a ‘fixed’ formation configuration to achieve the mission goals (i.e. although the spacecraft are orbiting the Earth their relative temporal or spatial separations remain fixed). The constellation-type of DSS typically positions the spacecraft in two ways, either distributed or as a ‘train’ constellation. A distributed constellation typically involves many spacecraft distributed over multiple orbital planes with multiple spacecraft per orbital plane separated evenly in true anomaly to provide global coverage of a service. Examples of distributed

constellation-type DSSs include Global Navigation Satellite Systems (GNSS) like NAVSTAR-GPS (GPS website, 2009), GLONASS (GLONASS website, 2009) and Galileo (Galileo website, 2009) and telecommunication systems like Iridium (Iridium website, 2009) and Globalstar (Globalstar website, 2009). An example of a distributed constellation-type DSS is shown in Figure 3-2 - left.

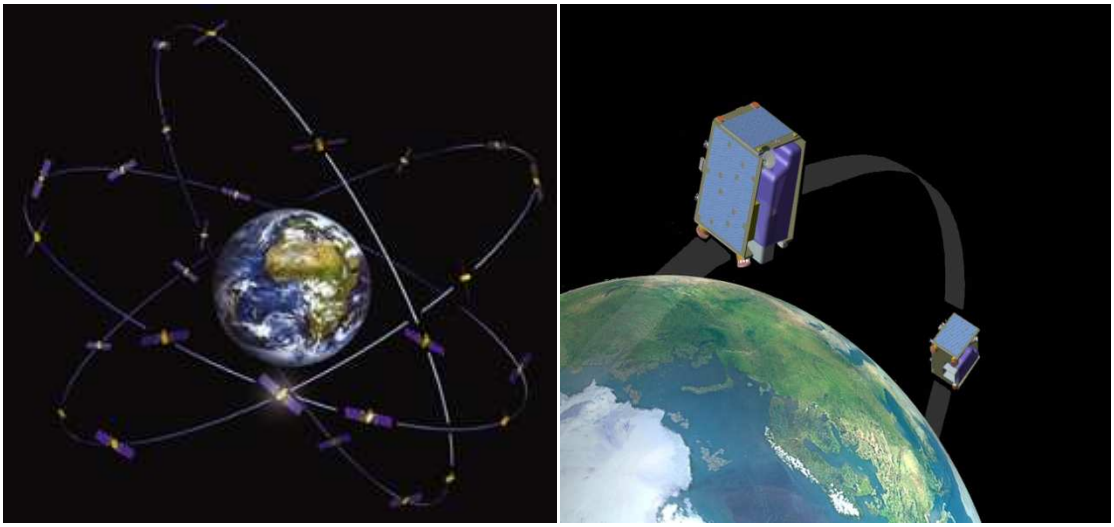

Figure 3-2 Galileo (ESA/Galileo Constellation) – left
and RapidEye (SSTL) – right

Train constellations typically involve much fewer spacecraft operating (roughly) in one orbital plane with a much smaller separation in true anomaly between the spacecraft to obtain paired scene ground observations using multiple different instruments. Examples of train constellation-type DSSs include the Morning Constellation (Colomb and Varotto, 2003), certain spacecraft in the A-Train (Vane, 2008), the Disaster Monitoring Constellation (DMC) (da Silva Curiel, et al., 2005) and RapidEye (de Selding, 2008). An example of a train constellation-type DSS is shown in Figure 3-2 – right.

As well as constellation-type DSSs, this first type of DSS family (where the spacecraft fly in formation) also contains the cluster-type DSS. The cluster-type of DSS typically involves a small number of spacecraft placed on similar orbits and similar true anomalies so that they can maintain a desired configuration at certain times within the orbit. For example, ESAs Cluster mission (ESA Cluster website, 2008) involves four identical spacecraft flying in formation to study the Earth's magnetosphere. The spacecraft are all placed in distinct highly elliptical orbits designed so that a tetrahedral configuration is formed during the apogee part of the orbit. During perigee, the configuration is broken but naturally recombines again towards apogee. The design of these orbits allows the tetrahedral configuration to be formed for a significant period of each orbit. Individual ground-controlled manoeuvres are performed to allow orbital changes that result in the change of the tetrahedron size as required by the science objectives of the mission. Such has been the success of the Cluster mission that both NASA and ESA are preparing new missions to perform an even more in-depth study of the magnetosphere, the Magnetospheric Multiscale (MMS) mission (Hughes, 2008) and Cross-Scale (ESA Cross-Scale website, 2008) respectively. Another cluster-type

DSS definition involves small number of spacecraft placed on similar orbits and similar true anomalies so that their relative motion along the orbit becomes the mission enabler. This is used for paired scene ground observations as required in the TanDEM-X dual spacecraft SAR interferometer (Zink, Krieger, Fiedler and Moreira, 2007) and the Aqua spacecraft in the A-Train (Kidder, Kankiewicz and Voder Haar, 2007). Cluster-type DSSs are also envisioned for new systems architectures involving fractionated spacecraft (Brown and Eremenko, 2006) where the individual spacecraft subsystems free-fly in a cluster providing the capability of a single monolithic spacecraft.

3.1.2 Spacecraft Formation Flying

The second family of DSSs are those where spacecraft formation flying occurs. Formation flying as defined by Scharf, Hadaegh and Ploen (2003) is "a set of more than one spacecraft whose *dynamic states are coupled through a common control law*". This explanation is expanded to stipulate that

- formation flying requires at least one spacecraft to track a desired state relative to another spacecraft in the formation
- the tracking control law must depend on the state of the tracked spacecraft.

To continue the analogy introduced in sub-section 3.1.1, this type of DSS involves control similar to that required by an air force display team like the Red Arrows. The position of each aircraft in a display formation is entirely governed by the pilot who is following a pre-determined plan and a set of rules relative to another member within the formation. Responsibility for achieving and maintaining the correct formation safely is the solely the pilot's.

This type of control was used for the Earth Observation-1 (EO-1) technology demonstration mission in 2000 and is to be used in many up-coming missions. In order to facilitate spacecraft formation flying there are a number of essential technologies required that may not be found on single spacecraft missions and DSSs where the spacecraft fly in formation. These include metrology sensors so the spacecraft know where the other spacecraft are in the fleet, inter-spacecraft communication systems, precision thrusting capability, on-board planning software, relative control laws and a high level of autonomy. Other technologies the possibility of spacecraft formation flying introduces include systems level architectures for control, communications and data processing and autonomous complex mission planning.

For the remainder of this chapter and this thesis the definition of formation flying adopted is that defined by Scharf, Hadaegh and Ploen (2003)

3.2 Spacecraft Formation Flying Applications and Missions

This section describes a number of spacecraft formation flying missions that have flown, are due to be flown or are in an advanced design phase. Only missions that utilise formation flying as defined by Scharf, Hadaegh and Ploen (2003) are introduced (i.e. autonomous

relative control of fleet members) to help focus the direction of the research being presented in this thesis.

3.2.1 Technology Demonstration

Technology demonstration of formation flying capabilities remains an important step in realising potential of formation flying for future missions. Technology demonstration is also important for future formation flying missions as they will benefit from using formation flying sensors, actuators and control architectures that have been flight-tested in the real space environment. Missions that have (or will be) directed to this application are:

- Earth Observing-1 (EO-1) - first formation flying demonstration mission (Folta and Hawkins, 2002)
- DART (Demonstration of Autonomous Rendezvous Technology) – failed spacecraft rendezvous demonstration (NASA Report, 2006)
- PRISMA – under development to test (with a high level of autonomy) guidance, navigation and control manoeuvring, evaluate a GPS-based navigation system, evaluate a vision-based sensor (VBS) and demonstrate an RF metrology package (Persson, Jacobsson and Gill, 2005)
- PROBA-3 (Project for On-Board Autonomy – 3) – under development to test a number of systems including GPS and RF-based relative position determination systems, coarse and fine optical metrology systems, validate formation flying manoeuvre algorithms and examine a range of formation flying architectures and scenarios (Borde, Teston, Santandrea and Boulande, 2004), (ESA Proba-3 Website, 2008)

3.2.2 In-flight Rendezvous and Docking

Automated rendezvous and docking of two spacecraft in-orbit was first achieved by ESA's Automated Transfer Vehicle (ATV) Jules Verne (ESA ATV Information Kit, 2008) when it docked with the International Space Station (ISS) in 2008. The ATV is an automated resupply vessel for the ISS and when docked is also capable of raising the ISS's orbit. The ATV remained with the ISS for several months before undocking and completing a destructive re-entry in September 2008. The use of autonomous formation flying for ISS cargo delivery missions greatly increases the accuracy and safety of resupply missions whilst reducing ground support requirements and the cost of similar (non-autonomous) missions. ESA plans to build five more ATVs to service the ISS in the coming years.

3.2.3 Earth Observation

The Earth Atmosphere Observatory formation (Mettle, et al., 2004) is a NASA advanced concept mission designed to view the Earth in continuous occultation from the Sun giving unprecedented observational data of the Earth's atmosphere over the 10-year mission lifetime. To realise continuous occultation of the Sun the formation needs to be placed at L_2 and remain within 200km of the Sun-Earth line. As this is an advanced concept the

hypothetical launch date is set at 2025-2030 to allow for the required technology development to occur (JPL Case Studies website, 2008) and as of writing this thesis no further information is available.

3.2.4 Astronomy

Astronomy is the scientific field most benefitting from the possibilities of spacecraft formation flying with a large number of mission proposals covering every aspect of the discipline some of these include:

- XEUS (X-ray Evolving Universe Spectroscopy) - investigate how supermassive black holes form and their effect on galaxy creation (Chabot and Udrea, 2006)
- MAXIM (Micro-Arcsecond X-ray Imaging Mission) - to use x-ray interferometry to image distant celestial objects like black holes and quasars
- NASA's Terrestrial Planet Finder Interferometer (TPF-I) - to perform infra-red interferometry with goals of finding Earth-like planets around other stars, characterising their atmospheres, studying gas giants and icy planets around other stars, performing comparative planetology and provide a platform for general astrophysics (Lawson, et al., 2007), (JPL TPF-I Website, 2008)
- DARWIN

3.3 Enabling Concepts for Spacecraft Formation Flying

Formation flying of a distributed spacecraft system as defined by Scharf, Hadaegh and Ploen (2003) requires a number of complex conditions that need to be met and involves the necessary adoption of a number of new concepts hitherto not utilised by traditional monolithic spacecraft. This sub-section introduces a number of these concepts, their variations and their comparative advantages and disadvantages.

3.3.1 Autonomous Manoeuvre Planning Organisations

Autonomy in relative position determination and control of at least one the spacecraft in a formation is the corner-stone of the formation flying definition (Scharf, Hadaegh and Ploen, 2003). However, autonomy can be extended throughout the many systems that make up the formation from individual spacecraft operations to formation manoeuvre planning and mission planning. Campbell and Schetter (2002) describe a number of multiple agent-based organisation types for spacecraft formations that describe the different levels of autonomy that can be used for the software architecture in a formation flying mission. Figure 3-3 is an adaptation of their work and describes four organisations depending on different levels of spacecraft intelligence. Campbell and Schetter (2002) compare these four organisations with four different mission scenarios for a formation of eight Earth-orbiting spacecraft acting as sparse-aperture radar. The levels of spacecraft intelligence are

- S_{NP} – spacecraft has no planning capability
- S_{IP} – spacecraft only has planning capability for itself

- S_{FP} – spacecraft can plan for all spacecraft within the formation

and the organisation descriptions with their merits are given in Table 3-1.

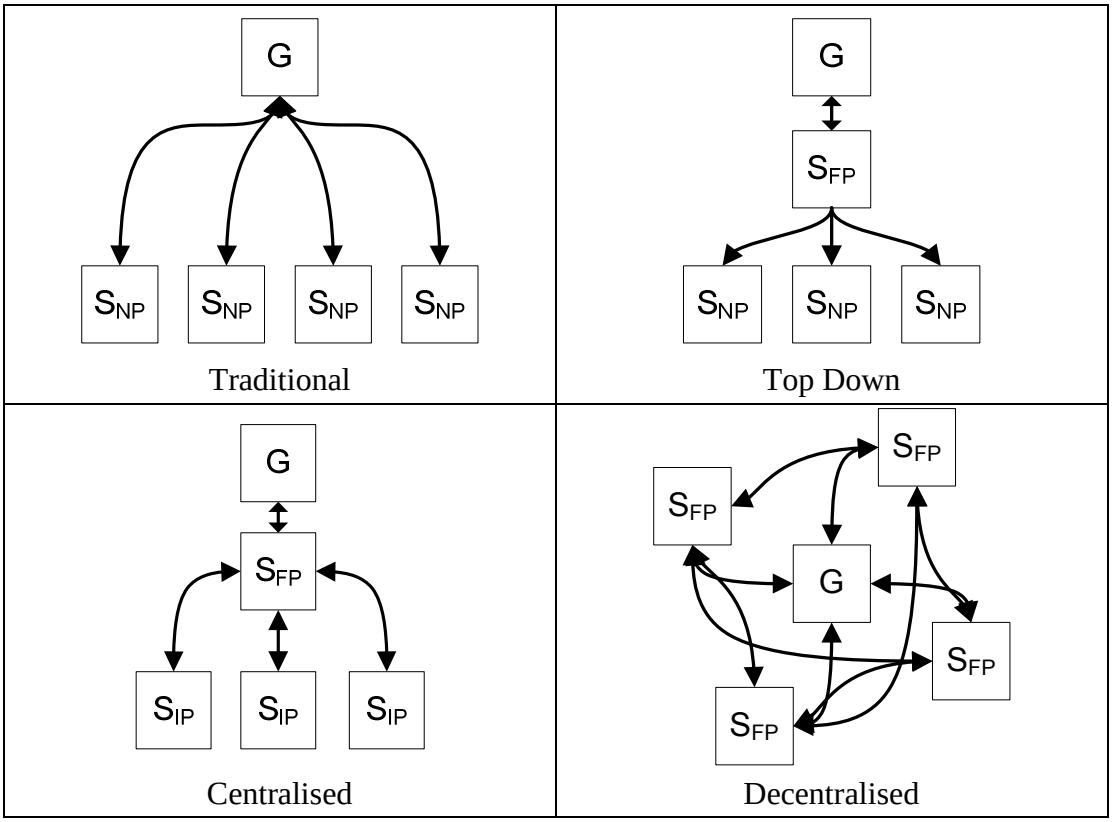

Figure 3-3 Formation organisation architectures with the ground station (G), spacecraft with no planning capability (S_{NP}), spacecraft with individual planning capability (S_{IP}) and spacecraft with formation planning capability (S_{FP}) and their links. Adapted from Campbell and Schetter (2002)

3.3.2 Control Co-ordination for Formation Flying

Formation flying control co-ordination approaches describe the control relationships between the different spacecraft in the formation and the level of autonomy required in the planning architecture. There are five main approaches to describing control co-ordination: the absolute approach, the leader-follower approach, the behavioural approach, the virtual structure approach and the virtual centre approach. This sub-section introduces these control co-ordination approaches and discusses their relative merits and weaknesses.

In the absolute control co-ordination approach each spacecraft moves independently and relative positions are not determined or controlled by the spacecraft. The formation is maintained through the careful design of reference trajectories and the spacecraft perform station-keeping manoeuvres to maintain position. This approach is used for those DSS missions where spacecraft fly in formation and so does not constitute formation flying control as per the definition in this thesis. Therefore no further reference to this type of control co-ordination approach is necessary.

Table 3-1 List of autonomous manoeuvre planning organisations with their advantages and disadvantages (adapted from Campbell and Schetter, 2002)

Organisation	Description	Advantages	Disadvantages
Traditional	entire manoeuvre planning for each spacecraft is performed on the ground	low spacecraft computational requirements, no inter-spacecraft communications and high reliability	very high ground operations costs and a large downlink capability
Top Down	S_{FP} plans for the entire formation	less ground operations support, downlink requirements and an increased working performance of the formation	significant computation load on the S_{FP} and a large crosslink capability
Centralised	S_{FP} acts as the formation planner and co-ordinator of the S_{IP}s individual plans	peak computational loads are decreased, low level for ground operations support, low downlink requirements and an increased working performance	average amount of processing increases, more complex software and spacecraft and an increased level of inter-spacecraft communications
Distributed	S_{FP}s share high bandwidth links with each other planning and co-ordinating the entire mission	peak computational loads reduced, best working performance, high reliability and low ground operations costs	average computational loads are greater complex software, spacecraft and inter-spacecraft communications capability

3.3.2.1 Leader-Follower Control Co-ordination Approach

The leader-follower control co-ordination approach is by the far the simplest and subsequently the most popular to use as a baseline when designing manoeuvre planning algorithms and control strategies. Within this approach one of the spacecraft is designated the leader whilst the others are designated followers. Some papers describe these as the chief and deputy spacecraft or the target and chaser spacecraft. The leader tracks a predefined trajectory and performs station-keeping manoeuvres to maintain this trajectory over time. The follower spacecraft tracks the position of the leader spacecraft and maintains a relative position. This relative position can either be static (describing a formation-keeping scenario) or dynamic (describing a formation manoeuvring scenario). Wang and Hadaegh (1996) describe different techniques for the leader-follower approach including nearest-neighbour tracking, barycentric tracking and leader tracking in a top-down tree-style formation organisation.

The leader-follower strategy is a very simple control co-ordination approach to implement because it requires very little information flow between the spacecraft. In addition, formation translational manoeuvres can be executed easily by simply moving the leader and dynamic disturbances experienced by the formation are captured by the leader and acted on

by the followers with no requirement for an external disturbance model. The leader-follower co-ordination approach however has its limitations. The leader serves as a single point of failure for the formation. With no explicit feedback to the followers, tracking may be difficult if the leader 'suddenly' executes a manoeuvre. Finally the leader will always use less fuel than the followers causing fuel imbalance amongst the fleet (though this can be mitigated by careful manoeuvre planning or periodic reassignment of the leader position to another spacecraft in the formation (Tillerson, Breger and How, 2003)).

The leader-follower control co-ordination approach is a very popular approach to use when designing formation flying missions and analysing control strategies. Catlin and McLaughlin (2007) use the leader-follower co-ordination approach to analyse the relative motion of spacecraft formations at the Earth-Moon triangular libration points whilst control of a spacecraft formation around the Sun-Earth L_2 libration point using the leader-follower co-ordination approach has also been examined (Hamilton, Folta and Carpenter, 2002). A study in using impulsive control for formation establishment and reconfiguration for Earth-orbiting formations adopts the leader-follower strategy (Vaddi, Alfriend, Vadali and Sengupta, 2005) and Vignal and Pernicka (2006) use the leader-follower approach to develop low thrust formation keeping controllers for Earth-orbiting formations.

3.3.2.2 Behavioural Control Co-ordination Approach

Within the behavioural control co-ordination approach each spacecraft's control effort is guided by its weighted desire to follow a modelled behaviour. Potential behaviours that can be modelled include formation-keeping, spacecraft goal-seeking, collision avoidance and obstacle avoidance (Beard, Lawton and Hadaegh, 2001). One way to approach behavioural control co-ordination is through the potential function method (McQuade, Ward and McInnes, 2003). In this method an analytical potential function describing the 'potential energy' of the formation is constructed. Desired behaviours (collision avoidance, goal-seeking etc…) are mathematically modelled using the current formation state and the desired formation state so that the behaviour potential is equal to zero when the current and desired states are identical. The potential function is simply the sum of the behaviour potentials and has a value of zero when the formation is in its desired state. A control feedback law is applied to each individual spacecraft in the formation to ensure the rate of change of the potential function remains negative, thus driving the potential to zero and the formation to its desired state. This method is shown by McQuade, Ward and McInnes (2003) to safely reconfigure a DARWIN-like formation at L_2 from random initial positions and is also used by Bennet and McInnes (2008) to study formation reconfiguration using bifurcation theory.

Another behavioural approach, similar to the potential function method, is equilibrium shaping (Pettazzi, Izzo and Theil, 2006). In this method the desired behaviours are represented as velocity vectors. The desired velocity for any spacecraft is simply the weighted sum of the individual behaviours and a velocity tracking feedback control law is implemented to manoeuvre the spacecraft to their desired positions. Equilibrium shaping is used by Pettazzi, Izzo and Theil (2006) in their analysis of formation flying reconfiguration manoeuvres using a swarm of coulomb satellites (King, Parker, Deshmukh and Chong, 2003). As well as being used to simulate translational formation reconfiguration manoeuvres

the behavioural control co-ordination approach has also been used in the design of control strategies to perform synchronised multiple spacecraft attitude manoeuvres using nearest-neighbour tracking (Lawton and Beard, 2002).

One of the advantages of the behavioural approach is its decentralised nature allowing it to be more flexible, reliable and robust. The behavioural approach allows a simple control strategy even with multiple competing objectives and the explicit feedback provides fast reaction of individual spacecraft to disturbances and the manoeuvres of other spacecraft within the formation. The behavioural approach does have its drawbacks however. It is very difficult to analyse mathematically and so certain formation characteristics like stability cannot be guaranteed (Beard, Lawton and Hadaegh, 2001). Also since the group behaviour is an emergent property, this can lead individual spacecraft in the formation to periodically act 'abnormally' preventing any optimisation of manoeuvres.

3.3.2.3 Virtual Structure Control Co-ordination Approach

The fourth formation flying control co-ordination approach appearing in the literature is called the virtual structure approach. In the virtual structure approach the entire formation is treated as a rigid-body with individual spacecraft co-ordinates fixed relative to a formation frame. During manoeuvres the formation frame translates and rotates in inertial space as if it were a rigid structure thus the formation configuration is maintained throughout. The virtual structure approach is examined for simple formation reorientation manoeuvres by Hammer, Piper, Thorp and Watkins (2004) whilst it is also used by Xin, Balakrishnan and Pernicka (2005) in the development of a deep-space spacecraft formation position and attitude control strategy. Beard, McLain and Hadaegh (2000) use a virtual structure-like within an optimal manoeuvre planning strategy that aims to generate an optimal control law for constrained formation reorientation manoeuvres in free space through optimising the position of the formation frame.

The virtual structure control co-ordination approach has its advantages in that it is easy to describe a coordinated behaviour for the formation and a tight formation can be maintain during reconfiguration manoeuvres. However, the maintenance of this tight formation may be detrimental to the fuel/time optimality of the manoeuvre for large reorientation manoeuvres (Hammer, Piper, Throp and Watkins, 2004) and the rigidity required by the virtual structure limits the potential applications of any formation adopting this approach.

3.3.2.4 Virtual Centre Control Co-ordination Approach

The virtual centre control co-ordination approach is similar to the virtual structure control co-ordination approach except the reference point of the formation frame is dynamic (relative to the spacecraft in the formation) throughout the manoeuvre. This idea was first used by Beard, McLain and Hadaegh (1998) in developing an optimal manoeuvre planning strategy for separated spacecraft interferometry missions (though they do not use the term 'virtual centre'). The position of the virtual centre in the formation represents the weighted average motion of the formation and includes the average disturbances on the formation. The spacecraft relative states are used to determine the state of the virtual centre (in position and velocity) and that in turn is used to calculate the desired relative states of the spacecraft. The

virtual centre therefore moves in inertial space during the reconfiguration manoeuvre as each individual spacecraft's updated state influences the virtual centre's state. Since the virtual centre is a weighted average it can be used to encourage fuel minimisation or fuel balancing during the manoeuvre.

The virtual centre concept has its weaknesses however. While Tillerson, Breger and How (2003) find the virtual centre approach more fuel efficient for Earth-orbiting formation reconfiguration manoeuvres than the leader-follower approach, Beard, McLain and Hadaegh (1998) find that, while more robust, the virtual centre concept increases overall fuel consumption since additional thrust is required to move relative to a dynamic point than a static point. Another drawback of the virtual centre concept is its requirement for a centralised computational architecture and therefore a greater need of inter-spacecraft communication to collate and distribute the required information amongst the fleet. Further issues arise due to the execution of error correcting control inputs. These are planned assuming a fixed virtual centre but the position of the virtual centre will move as the spacecraft move resulting in poorly planned manoeuvres. Tillerson, Breger and How (2003) suggest a method to solve this issue. In their solution each spacecraft broadcasts its planned control inputs to the other spacecraft. These plans are then used to calculate the expected motion of the virtual centre in the near future. This method however involves significantly more inter-spacecraft communication and computational effort.

As well as the literature introduced above the virtual centre concept is used to design optimal reconfiguration manoeuvres for a Darwin-like interferometry mission at L_2 (Penin, Araujo and Avila, 2005). Campbell, Zanon and Kulkarni (2004) also use the virtual centre control co-ordination approach for formation keeping control of a circular formation of spacecraft tracking a Halo trajectory at L_2.

3.3.3 Autonomous Control Strategies

Formation flying control allows spacecraft to maintain relative position and/or perform translational manoeuvres relative to each other. Formation-keeping requires the spacecraft to maintain a desired relative position with respect to another spacecraft in the formation in the face of external disturbances. Formation manoeuvring is when the spacecraft are required to follow trajectories relative to one another. The control strategies developed to tackle the formation-keeping and formation manoeuvring aspects of formation flying missions are introduced in this sub-section:

- Proportional/Integral/Derivative (PID) – for formation-keeping within the leader-follower control co-ordination approach (Açikmese, et al., 2004)
- Proportional/Derivative (PD) - to demonstrate the virtual structure control co-ordination approach (Penin, Araujo and Avila, 2005), (Wei and Beard, 2004)
- Linear Quadratic Regulator (LQR) - for formation-keeping within the leader follower control co-ordination approach (Gurfil and Kasdin, 2004), Smith and Hadaegh (2005), Roberts (2005)
- Linear Quadratic Gaussian (LQG) – for formation flying with a centralised planning organisation (Lagadec, Lebas and Ankersen, 2003), (Davidson, et al., 2006)

- Adaptive neural control - for formation-keeping within the leader-follower control co-ordination approach (Gurfil, Idan and Kasdin, 2002 and 2003)
- Input Feedback Linearisation (IFL) - for formation-keeping within the leader-follower control co-ordination approach (Folta et al., 2004), (Howell and Marchand, 2003), (Marchand and Howell, 2005)
- Differential corrections - for formation-keeping within the leader-follower control co-ordination approach (Carlson, Pernicka and Balakrishnan, 2004)
- Θ-D - for formation-keeping within the leader-follower control co-ordination approach (Xin, Balakrishnan and Pernicka, 2004) and for formation manoeuvres using the virtual structure control co-ordination approach (Xin, Balakrishnan and Pernicka, 2005)
- H_2 – for formation-keeping using the leader-follower control co-ordination approach (Chabot and Udrea, 2006)

3.4 Chapter Summary

This chapter introduced the concept of spacecraft formation flying and gave a definition of that concept to avoid confusion with other distributed spacecraft systems (DSS). For a DSS to be formation flying two requirements need to be met: the spacecraft must track a state relative to another spacecraft in the formation and the tracking control law must depend on the state of the tracked spacecraft. This definition helped to remove the ambiguity present in the literature and press that use the term ‘formation flying’ to describe many different types of DSS mission.

Following the definition of the concept a review of past, present and future spacecraft formation flying missions were given. This review highlighted the fact that spacecraft formation flying as a usable concept is still very much in its infancy with the few missions flown (and soon to be flown) performing formation flying between only two spacecraft in a technology demonstrating capacity (with the exception of the ATV). This section also showed how future missions, especially astronomy missions, aim to use spacecraft formation flying of more than two spacecraft. This allows astronomers to simulate telescope apertures much larger than would be available on a monolithic spacecraft and perform complex imaging, like interferometry, with widely separated telescopes.

The final section in this chapter introduced a number of related spacecraft formation flying concepts that are enabled by, and unique to, this technology for space missions. These included the use of autonomous control architectures for the potentially distributed planning for formation flying manoeuvres, control co-ordination for the formation GNC system and the control strategies designed to execute the manoeuvres.

In the next chapter the concept of mathematical optimisation is introduced and examples given to demonstrate the different types of optimisation problem commonly found. The next chapter also reviews a number of optimisation algorithms that can be applied to find solutions to different types of optimisation problem.

4. OPTIMISATION

Optimisation is a field of research that spans across many different industries and can be applied in a large number of situations. Optimisation techniques are used to design aircraft, manage machine jobs in factories, schedule network data in IT systems and solve complex problems in microeconomics. For space missions there are a large number of factors that undergo optimisation processes during design (e.g. the trajectory, power requirements, mass requirements, fuel requirements, etc.) but there are equally important optimisation processes adopted during the missions (e.g. planning attitude manoeuvres, scheduling payload operations, etc.). Operations for spacecraft formation flying missions are much more complex due to multiple spacecraft and their requirement to operate as one unit. These missions require the use of additional optimisation for planning manoeuvres. In this chapter the concept of optimisation, in a mathematical sense, is introduced, techniques are presented that address the optimisation problem and references are made relating optimisation within the Manoeuvre Planning Architecture (MPA).

4.1 The Optimisation Problem

In mathematics the optimisation problem can typically be described by the following equation:

$$\arg\min_{X} \begin{pmatrix} J_1(X) \\ J_2(X) \\ \vdots \\ J_n(X) \end{pmatrix} \tag{4.1}$$

where

$$\left[J_1, J_{2,} \ldots, J_n \right] = f\left(x_1, x_2, \ldots, x_m\right) \equiv f\left(X\right) \tag{4.2}$$

Equation (4.1) describes an optimisation problem where the goal is to find a value for X that minimises the objective J, Equation(4.2). This is called the objective function. X is an independent variable with m terms. These terms (x_1, x_2, etc.) are called variables. J is the global objective of the optimisation problem and has n terms. These terms (J_1, J_2, etc.) are called objectives. The generalised objective function given in Equation (4.1) can easily be adapted to encompass different or more complex problems, for example

$$\arg\max_{X} \begin{pmatrix} J_1(X) \\ J_2(X) \\ \vdots \\ J_n(X) \end{pmatrix} \tag{4.3}$$

or

$$\begin{aligned} &\arg\min_{X} J_1(X) \\ &\arg\max_{X} J_2(X) \\ &\arg\min_{X} J_3(X) \end{aligned} \tag{4.4}$$

Equation (4.3) describes the problem where the global objective is to maximise the values of the individual objectives. For Equation (4.4) the global objective is to minimise J_1 and J_3 whilst maximising J_2. It should be noted that achieving optimisation of the global objective does not necessarily mean that the individual objectives will be optimised. The global optimum may have to represent a trade-off between conflicting objectives (e.g. fuel minimal vs. time minimal manoeuvre). Objectives can, however, be prioritised to ensure that more important objectives are least affected by the optimisation goals of less important objectives.

In addition to managing the objectives, constraints can also be introduced into the optimisation problem to limit the range of the elements of the independent variable or enforce other relationships within the cost function. These can take the form of equality and inequality constraints, e.g.

$$\begin{aligned} x_1 &\geq 0 \\ x_2 + x_3^2 &= 0 \end{aligned} \tag{4.5}$$

The scope of the optimisation problem can be constrained by bounding the number of variables and objectives giving three optimisation types in order of increasing complexity:

- when $m=1$ and $n=1$ - this is single-variable optimisation
- when $m=1$ and $n>1$ - this is multi-variable optimisation
- when $m>1$ and $n\geq 1$ - this is multi-objective optimisation

Each different type of optimisation can be solved a number of ways however in general the more complex types require more complex solution methods. Solution methods where the optimisation problem can be described by a cost function are detailed in sub-section 4.4. Multi-variable optimisation is used in Chapters 8 & 9 to find solutions to some of the formation flying manoeuvre planning problems introduced in Chapter 1.

Sometimes the optimisation problem is so complex that it cannot be defined within a cost function. Without a cost function traditional optimisation solution methods cannot be used. Instead, bespoke algorithms must be created for the specific type of problem to be solved. This is encountered in Chapter 7 where the manoeuvre scheduling problem is indefinable within a cost function. Other optimisation problems may be definable using a cost function but analytical methods can be used to find the optimal solution. This is the case in Chapter 10 where the station-keeping manoeuvres are optimised.

4.2 The Solution Space

The solution space of an *n* variable optimisation problem is an *n+1* dimensional space that contains all the possible solutions to the objective function for all possible combinations of the variables. When trying to optimise an objective function, numerical solvers calculate solutions to the objective function, analyse the solutions found and then generate new solutions based upon the analysis. In this way the solvers move through the solution space until the desired solution has been found.

Figure 4-1 shows a graphical example of an *n=2* solution space generated using the MATLAB® function '*peaks(20)*' and illustrates some of the properties found in solution spaces. The main observation is the existence of multiple maxima and minima in the solution space. The true maximum or minimum of any optimisation problem is called the global solution whilst the other observed maxima or minima are called local solutions. For the solution space represented in Figure 4-1 $\arg\max_{X,Y} J = (25, 38)$ and $\arg\min_{X,Y} J = (25, 10)$.

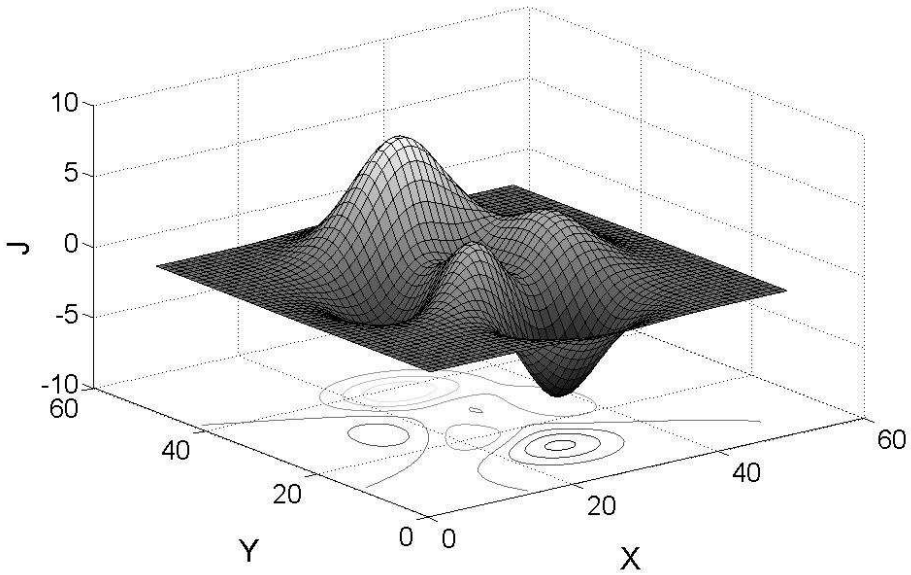

Figure 4-1 Solution space example with *n=2*. This solution space demonstrates multiple local maxima and minima and the global maximum and minimum.

Solution spaces with many local solutions can be difficult to solve because some solvers get 'trapped' at these local solutions and are unable to find the true global solution to the problem. Another issue with trying to find the global solution involves proving that the global solution has indeed been found. Complex optimisation problems require the use of numerical

solvers because analytical methods cannot be used. In these cases it is very difficult to prove the global optimum of the solution space. Confidence that the best solution has been found can be built through repeated applications of the problem, through the use of multiple different optimisation techniques and more detailed analysis of the solution space.

Another feature that some solvers find difficult is the discontinuities within the objective function that some problems present. Those solvers that use the gradient of the objective function to move through the solution space cannot operate effectively when faced with discontinuities within the solution space. A final feature of solution spaces that affect the efficiency of optimisation solvers is the size of the space. As well as being (potentially) multi-dimensional the variables that make up these dimensions can have a range $-\infty \le X \le \infty$. With such a vast solution space to navigate through it is important to constraint the sizes of the dimensions (where appropriate). For example, when optimising an angle the dimension range should be at the most $-\pi \le X \le \pi$.

The complexity of a solution space is a function of the number of variables, *n*, that input to the cost function. Many of the solution space problems detailed above are experienced by the solution space of the cost function defined in Chapter 9. This cost function aims to find collision free and thruster plume free trajectories for all the spacecraft in the formation whilst managing fuel consumption and has 16 input variables. The resulting topography of the solution space excludes the use of many of the simpler optimisation algorithms and requires the use of a much more complex solver that requires greater processing capacity. This requirement for greater processing capacity has implications on the algorithms suitability of on-board operation.

4.3 Optimisation Problem Examples

Single variable optimisation can be viewed simply by trying to find the minimum value of a single variable equation

$$\arg\min_{x} y(x) \text{ where } y(x) = x^4 - x^3 - x^2 + x - 1 = 0 \tag{4.6}$$

For this type of problem the simple analytical technique of finding and evaluating the roots of the derivative of the function shows the solution to be $x = -0.64$. The solution space of this function can be found in Figure 4-2 (left).

In the multi-variable optimisation problem there is still only one objective but more than one variable that affects the value of the objective function. This can be seen with the example in Figure 4-2 (right) of the problem

$$\arg\max_{x,y} z(x,y) \text{ where } z(x,y) = x.e^{\left(-x^2-y^2\right)} \tag{4.7}$$

Figure 4-2 (right) shows that the answer is $(x, y) = (0.5, 0)$ but using analytical or graphical methods to find the exact answer is difficult, especially when the number of variables is

greater than two. An optimisation algorithm like '*fminunc*' from MATLAB's Optimisation Toolbox, however, is capable of finding a more accurate answer, $(x, y) = (0.71, 0.00)$.

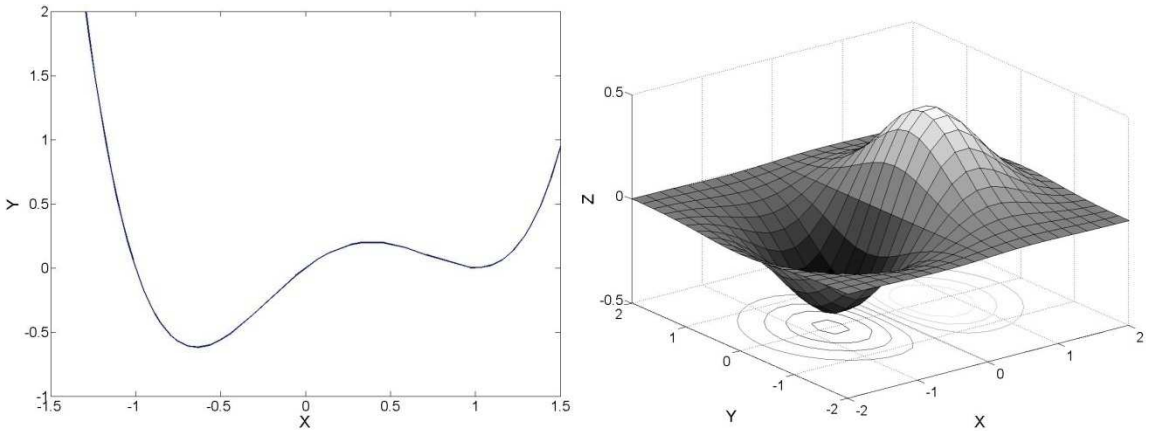

Figure 4-2 Graphs representing the single variable (left) and the multi-variable (right) optimisation examples

Moving to the multi-objective example increases the problem complexity with the addition of multiple variables and objectives. Often in multi-objective cases the objectives are in conflict so that lowering the cost of one objective raises the cost of another. Concise graphical and analytical solutions to such problems are all but impossible. As an example, consider the design of an aircraft wing. Simple objectives would be to maximise lift, minimise drag and minimise weight, whilst the design variables could be wing area, thickness, flap position, dihedral angle and construction material. Examining all the perturbations for even this simplified problem would take a great deal of time due to the size of the solution space. One method employed to solve multi-objective optimisation problems is Pareto-optimisation (Liu, Yang and Whidborne, 2003). For any *m* dimensional multi-objective optimisation problem the Pareto-optimal solution is defined as a solution where the cost of one objective cannot be reduced without increasing the cost of at least one other objective. This can be explained further by referring to Figure 4-3. Here the optimisation problem is to find the maximum of two objectives, J_1 and J_2. Any solution within the attainable set, e.g. point A, is sub-optimal since the cost of both objectives can be increased. Point B represents solutions outside the attainable set that cannot be achieved. The boundary between the attainable and unattainable sets is called the Pareto-optimal set. The Pareto-optimal set is a set of attainable solutions where the cost of one objective cannot be increased without reducing the cost of another (e.g. the cost for J_2 for solutions C and D cannot be increased without reducing the cost of J_1). Cleary for Pareto-optimisation there is no global optimum and the required solution must be selected from the Pareto-optimal set that represents the best trade-off between the competing objectives.

Other methods employed to solve multi-objective optimisation problems include the weighed sum method, the goal attainment method, the method of inequalities and multi-objective genetic algorithms (GAs) to name but a few (Liu, Yang and Whidbourne, 2003).

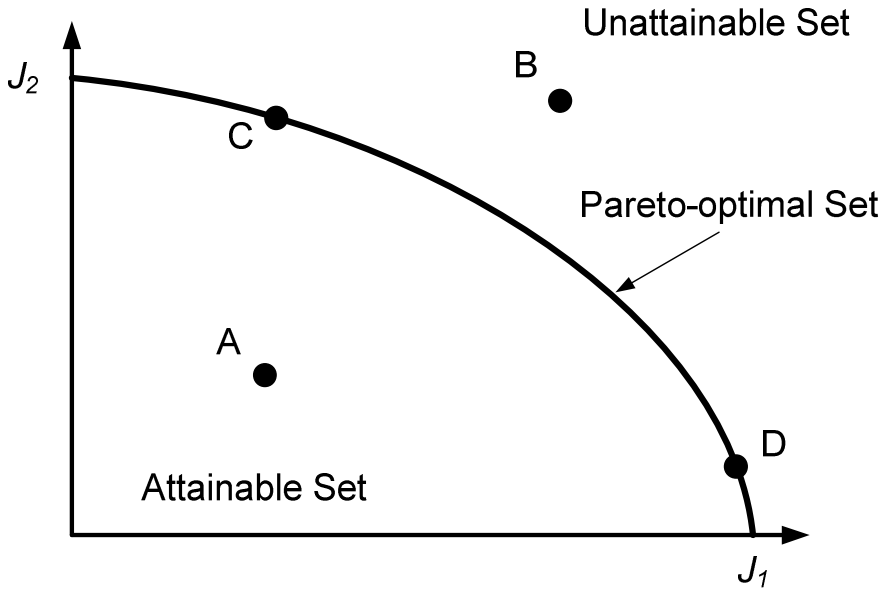

Figure 4-3 Attainable and Pareto-optimal sets in multi-objective optimisation problems. Point A represents a sub-optimal solution, point B represents an unattainable solution and point C and D represent Pareto-optimal solutions (where decreasing one cost increases the other).

4.4 Optimisation Algorithms

The complexity of multi-variable and multi-objective optimisation problems has required researchers to employ more complex and innovative techniques to solve them. These optimisation algorithms fall, roughly, into four categories; gradient methods, direct search methods, stochastic methods and evolutionary methods. In this section these four types will be introduced and examples given of the innovative optimisation algorithms that are used to solve optimisation problems today.

4.4.1 Gradient Methods

Gradient methods are the simplest and most commonly used types of optimisation algorithms and in general work by requiring the use of the derivative of the objective function. The simplest of these is the Gradient Descent algorithm. Put simply, the algorithm iterates towards the local minimum of the solution space by taking steps towards that solution in proportion to the size of the gradient of the objective function. Given an initial guess x_0, the iteration

$$x_{n+1} = x_n - \gamma \nabla F(x_n), n > 0 \tag{4.8}$$

should converge at the local minimum of the function *F(x)*. The value, γ, is a step size that can be changed during iteration. Due to its simplicity the gradient descent method can be very time consuming for complex objective functions. It is also very sensitive to initial conditions and cannot find the global minimum of a function with many local minima unless the initial guess is close to the global solution.

The gradient descent method is the basis for a number of more complex gradient optimisation techniques including Newton's Method (Press et al. Section 9.4, 1992), the Conjugate Gradient Method (Press et al. Section 10.6, 1992), the Gauss-Newton Method (Wikipedia, 2009a), the Levenberg-Marquardt algorithm (Wikipedia, 2009b) and the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method (Press et al. Section 10.7, 1992). To avoid falling into local minima, gradient methods can be adapted to allow the algorithm to jump out of the local minimum and continue searching the solution space for the global solution. Additionally, gradient methods can be combined with non-gradient methods, e.g. simulated annealing (Press et al. Section 10.9, 1992), to overcome the local minima problem (Bailey, McLain and Beard, 2001).

Gradient methods also lend themselves to be easily derived for specific purposes. Singh and Hadaegh (2001) use cost function gradients to calculate energy minimum reconfiguration manoeuvres for formation flying incorporating collision avoidance whilst Sultan, Seereeram and Mehra (2004a, 2004c) derive a sequential gradient-based algorithm to provide fuel optimised reconfiguration manoeuvres for formation flying incorporating fuel equalisation and collision avoidance.

Gradient methods are very useful as secondary optimisation tools that use the solution of a more complex optimisation algorithm as their initial input. This suits the types of optimisation algorithm that, due to stochastic processes, are unable to directly find the minimum and can only find solutions within the vicinity of the minimum. This technique is employed for the optimisation problem tackled in Chapter 9.

4.4.2 Direct Search Methods

Direct Search methods are a type of optimisation that does not require knowledge of the derivative of the objective function to find solutions. These methods therefore lend themselves better to problems where the objective function cannot be differentiated, is discontinuous or involves stochastic processes. Direct Search methods are small population-based calculating a small number of simultaneous solutions to the objective function during iteration. This allows them to search quickly through a solution space for the optimal result.

4.4.2.1 Nelder-Mead (Simplex) Method

The Nelder-Mead (Simplex) algorithm is a direct search optimisation method capable of finding the global minimum of an objective function of *n*-variables. First conceived in 1965 (Nelder and Mead, 1965) it falls under the class of an unconstrained non-linear optimisation algorithm.

A simplex is a geometrical construct in multi-dimensional space. For example a 2-simplex is a triangle; a 3-simplex is a tetrahedron. In the Nelder-Mead (NM) algorithm a simplex is formed with $n+1$ vertices for an objective function of n variables. The algorithm performs the following operations:

- The solutions of each of the simplex's vertices are calculated.
- These solutions are sorted into ascending order.

- The position and solution of the centroid of the simplex is calculated from the vertex solutions.
- If the centroid solution falls within the algorithm's stopping conditions then the point is accepted and the algorithm terminates.
- If the centroid solution is outside the algorithm's stopping conditions it is compared to the *n+1* vertex solutions. From this comparison a set of rules are followed that ultimately constructs a new simplex. The process is repeated until the algorithm terminates with a solution.

The comparison rules cause the simplex to undergo a number of transformations which can include reflection, expansion and contraction. After these transformations a new simplex is formed and the algorithm iterates. Over successive iterations the simplex moves around the solution space and eventually shrinks around a minimum solution.

The NM algorithm is an extremely popular and effective optimisation algorithm and can be found in the MATLAB® Optimisation Toolbox (Coleman and Zang, 2005) embedded into the '*fminsearch*', '*linprog*' and '*quadprog*' functions. However the solution the algorithm converges to cannot be guaranteed to be global as the initial simplex is governed by the initial conditions input to the algorithm. Therefore in a solution space with many local minima, the proximity of the initial simplex to one of these minima will likely prevent the algorithm from converging to other minima (one of which may be the global solution).

The NM algorithm, within the context of an embedded MATLAB® function, is used by Beard, McLain and Hadaegh (1998), Beard and Hadaegh (1999) and Beard and McLain (2000) to find the minimum for a cost function to achieve fuel optimisation for unconstrained and constrained retargeting manoeuvres of spacecraft formations in free-space. Bailey, McLain and Beard (2001) also use the NM algorithm within an algorithm for calculating fuel optimisation via an optimal tour of retargeting and imaging of multiple stellar sources.

4.4.2.2 PatternSearch

The PatternSearch (PS) algorithm is a direct search method implemented in the MATLAB® Genetic Algorithm and Direct Search (GADS) toolbox (Abramson, 2006). The algorithm is capable of finding the minimum of unconstrained non-linear multi-variable objective functions and operates in a similar fashion to the Nelder-Mead algorithm.

To iterate, the PS algorithm forms a mesh about the current solution in the solution space. The mesh is constructed by adding to the current solution a number of fixed vectors (forming a pattern about the current solution). The solutions of each point in the mesh are calculated and the best solution is used to form the next iteration. If the best mesh solution improves upon the current solution then a new mesh of increased size is constructed around the best mesh solution for the next iteration of the algorithm. If the current solution remains the better option a new, smaller, mesh is constructed around the current solution for the next iteration. The size changes in the mesh help the algorithm to escape local minima since the mesh increases in size with successful searches. As the algorithm progresses the mesh moves through the solution space, increasing and decreasing in size until it shrinks around the function minimum.

The PS algorithm is the chosen optimisation algorithm to solve the optimisation problem posed in Chapter 8. The cost function input has 4 elements and the algorithm consistently performs well with the resulting solution space. The limitations of the algorithm however are highlighted in Chapter 9 where the 16-dimensional solution space of the optimisation problem posed is too complex for the PS algorithm to perform consistently well in.

4.4.3 Stochastic Methods

Stochastic methods are governed by their reliance on a probability function that determines how the algorithm proceeds at each step. The stochastic nature of the method helps the algorithm to avoid getting trapped at local minima within the solution space. However, due to this stochastic nature, it is unlikely that the algorithm will follow the same route twice thus making comparative analysis of solutions more complex.

4.4.3.1 Simulated Annealing

Simulated Annealing (SA) is a probability-based method for finding the minimum of a multi-variable function (Press et al. Section 10.9, 1992). The method draws its inspiration from the process of annealing in thermodynamics.

In the SA algorithm, the value of the current state is compared to the value of a neighbouring state. The algorithm then decides whether to move to the neighbouring state or to remain at the current state until the next iteration. This decision is based on the Boltzmann probability distribution function which gives the probability of a transition as a function of the value of the current state, the value of the neighbouring state and a time varying global parameter, *T*. The function is designed so that it favours "downhill" moves (i.e. to a better state) over "uphill" moves, but that "uphill" moves are not ruled out. This helps to prevent the algorithm from remaining at a local minimum (the 'freezing' problem). An annealing schedule defines the parameter, *T*, which decreases to zero over the allotted time frame for the algorithm. The probability function is designed so that at the start of the algorithm the probability, *P*, is 0.5 (i.e. equal probability of "uphill" or "downhill" move), increasing to $P = 1$ by the end of the allotted calculating time for the function (i.e. only "downhill" moves allowed).

As with all optimisation methods SA is not without its drawbacks. If the SA algorithm finds a minimum early in the time frame the algorithm could escape from that minimum and never find it again (since the likelihood of an "uphill" move is greater at the beginning). To combat this, a 'best solution so far' variable can be included so that the algorithm can be restarted at that best solution should the initial solution prove sub-optimal. Another problem is coined the 'freezing problem'. Any local minima can trap the algorithm especially towards the end of the time frame (when the likelihood of an "uphill" move is less). However, if the algorithm is started again from that local solution, the chance of it escaping and converging at the global minimum increases. Therefore, for the SA algorithm to operate robustly it needs to be implemented a number of times.

Simulated annealing is used by Bailey, McLain and Beard (2001) as part of an algorithm to find an optimal tour (incorporating retargeting and imaging manoeuvres) between multiple stellar sources for a formation flying mission.

4.4.3.2 Other Stochastic Methods

Due to the success of simulated annealing other stochastic methods have been developed to improve on the technique. Simulated Quantum Annealing (Stella, Santoro and Tosatti, 2005) is similar to simulated annealing except the probability distribution function is based on quantum tunnelling instead of thermodynamic cooling. Stochastic Tunnelling (Wenzel and Hamacher, 1999) is a more robust version of SA able to circumvent the 'freezing' problem by allowing tunnelling to a different part of the solution space. The Cross-Entropy Method (de Boer, et al., 2005) uses the Kullback–Leibler divergence of two probability distributions to find the optimal solution within the solution space.

4.4.4 Evolutionary Methods

Evolutionary methods are inspired by the probabilistic nature of evolutionary biology which, over hundreds of generations, improves biological life. Evolutionary methods are large population-based initially using a large number of simultaneous solutions to the objective function. This means that the methods are very good at finding the best solutions of complex optimisation problems that contain many local minima. They do however require considerable computational resources to operate and, due to their stochastic nature, make comparative analysis of solutions more complex.

4.4.4.1 Genetic Algorithms

Genetic Algorithms (GAs) are an increasingly popular method of global optimisation (Coello Coello, Van Veldhuizen and Lamont, 2002). The GA begins by randomly populating the solution space of an objective function with a number of agents. Each agent is made up of the *n* variables that make up the objective function. At each time-step the solution for each agent is calculated and a stochastic selection process chooses a proportion of the population. This process is designed so that agents with better solutions will be selected (but there is a chance that agents with worst solutions will also be selected). The next generation of agents is formed by transforming the selected agents using the genetic operators, crossover and mutation. In crossover new agents are formed by swapping the individual variables from two of the selected agents. In mutation, one or more variables from an agent may randomly change value. Crossover and mutation over the whole population creates a new population with an average better solution than the old population whilst maintaining enough diversity within the population to avoid local minima. Over time the population evolves towards an agent whose solution represents the global minimum in the solution space.

The accuracy of the optimal agent depends on the time allocated for the GA to run. To minimise this time, the GA can be stopped early with a population of agents whose solutions surround the global minimum. One of these agents can then be used as the initial condition for a faster local optimisation routine that can find the exact solution.

A genetic algorithm can be found in the MATLAB® Genetic Algorithm and Direct Search toolbox (Abramson, 2006). Yang et al. (2002) use a genetic algorithm within an optimization process to find fuel optimal reconfiguration manoeuvres for multiple spacecraft formation flying. Seereeram et al. (2000) uses a genetic algorithm to optimise a multiple

spacecraft trajectory planner to include fuel minimisation and equalisation, collision avoidance, distance minimisation and time minimisation.

The GA is part of a two stage optimisation process used to solve the optimisation problem posed in Chapter 9. The 16-dimensional solution space of the cost function proves too complex for direct search methods but the GA consistently performs well in finding good solutions.

4.4.4.2 Differential Evolution

Another evolution-based optimisation method is Differential Evolution (DE) (Stron and Price, 1997). The DE algorithm begins by randomly generating a population of agents over the solution space of an objective function. Like in the GA, each agent is made up of the *n* variables that make up the objective function. At each time-step the solutions of each agent are calculated. For every agent, two other agents are randomly selected and their difference taken. This difference is weighted and added to a third randomly chosen agent. This is the 'mutation' part of the algorithm. Crossover is performed on the mutated agent and the initially selected agent to form a mutated/crossover agent whose solution is calculated. This is compared to the solution for the initially selected agent and the agent with the better solution survives to the next generation. Over time the population evolves towards an agent whose solution represents the global minimum in the solution space. Similar to the GA's optimal agent accuracy, the DE algorithm's accuracy is time-dependant. So a faster local optimisation routine can be used to achieve a global result in minimum time.

Differential Evolution is used by Pettazzi, Izzo and Theil (2006) to find energy optimal solutions to formation-keeping and reconfiguration manoeuvres for a swarm of electrostatically propelled satellites.

4.4.5 Swarm Methods

Swarm optimisation methods are also inspired by nature but this time by the observed movement of swarms, (e.g. birds flocking, fish schooling, etc.). These methods are large population-based and have very similar optimisation properties to the evolutionary methods introduced above.

4.4.5.1 Particle Swarm Optimisation

Devised by Kennedy and Eberhart (1995), Particle Swarm Optimisation (PSO) is a population-based optimisation technique inspired by observed bird-flocking and fish-schooling mechanisms.

In the PSO algorithm the solution space of an objective function is populated with a number of agents whose parts are made up of the *n* variables of the objective function. Each agent is also imparted with a random velocity vector through the solution space and two weighted 'desire' attributes. These attributes are the desire to continue travelling through the solution space (individuality) and the influence of other agents (sociality). At each time step each agent's solution is calculated and compared to the solutions of a number of neighbouring agents. These neighbours remain the same regardless of the distance in solution

space separating the agents. Depending on the agent's solution, the neighbour's solutions and the agent's individuality/sociality weights the velocity vector may be changed for the next time step. For example a social agent will be attracted to a neighbour with a slightly better solution. However an individual agent will not be attracted to a neighbour unless the neighbouring solution is vastly better than its own. This influence of neighbouring agents and 'desire' attributes help to avoid convergence to local minima. Over time the agents will all converge to the one agent that sits near the globally optimal solution. Like in GAs and DE, the accuracy of the solution is time-dependant, so to avoid long convergence times, the algorithm can be stopped early and a local optimisation routine run to find the exact solution.

4.4.5.2 Ant Colony Optimisation

Ant Colony Optimisation (ACO) is a method devised by Dorigo and Di Caro (1999) and was inspired by the ability of ant colonies to always find the shortest route between the nest and their food source.

For the ACO algorithm the solution space is populated with a number of agents made up of the *n* variables of the objective function and the solution for each agent is calculated. At each time-step the agent probabilistically moves to another solution by discretely changing one of the *n* variables. Every time a move is made a marker (called a pheromone) is deposited at the old solution so that another agent arriving at that solution will know where the previous agent went. The pheromone is a function of the difference between the values of the old solution and the new one. For example, if the agent moves to a worse solution the pheromone deposited would repulse other agents from that path. But if a better solution was found the pheromone would attract other agents to that path. When an agent arrives at a solution, its next move is probabilistically influenced by the number (strength) and type (attractive/repulsive) of pheromones at the solution. The agent is more likely to follow a path to a better solution when many other agents have followed that path before. This is called autocatalysis. Of course due to the stochastic nature of the decision where to move there is always the chance the agent will select a path that has not been visited or a path marked with a repulsive pheromone. This keeps the agents from blindly following each other to local minima. At each time-step the 'best solution found so far' for each agent is broadcast and the algorithm terminates when a certain proportion of the agents broadcast the same solution.

ACO is a complex and potentially CPU-intensive optimisation method which can be improved upon by adding extra abilities to the agents like look ahead, backtracking and local optimisation (Dorigo, Di Caro and Gambardella, 1999).

ACO is the inspiration of the bespoke optimisation algorithm designed in Chapter 7 to optimise the manoeuvre scheduling problem. Whilst the optimisation problem prevents the use of multiple agents simultaneously, the stochastic nature of the path decision process is emulated so that the choice of path is governed by a tuneable probabilistic procedure.

4.5 Chapter Summary

In this chapter the mathematical concept of optimisation was introduced and different types of optimisation problems defined. These problems are characterised by the number of

objectives and variables present in a ‘cost function’ that describes how the objectives change with respect to the variables. Single variable, multi-variable and multi-objective optimisation problems can be constructed to describe a large number of real, and imaginary, systems that can be optimised to exploit the capabilities of the system. The concept of the cost-space was also introduced and common features such as maxima, minima and local and global solutions discussed. Both these concepts were forward-referenced to chapters in the thesis where optimisation is found. A number of examples of optimisation problems were given and it was shown how analytical and graphical methods can be employed for 1-dimensional and 2-dimensional problems.

For more complex optimisation problems, using multi-variable or multi-objective cost functions, numerical methods must be used to find solutions within these systems. A large number of algorithms have been developed to cope with these more complex problems ranging from the gradient methods, which use the derivative of the cost function, to stochastic, evolutionary and swarm methods, that emulate processes found in nature. As with all numerical methods their accuracy is limited by the amount of time devoted to computing a solution and the computing resources available at the time. A list of common optimisation algorithms was provided with descriptions and forward-referencing indicated which optimisation algorithms were used in the research project.

5. FORMATION FLYING CONCEPTS ANALYSIS AND SELECTION

The previous three chapters have introduced a large number of concepts that must be discussed before further definition of the manoeuvre planning architecture can be realised. This chapter introduces these concepts with reference to manoeuvre planning for spacecraft formation flying missions and various decisions are made regarding their use within the manoeuvre planning architecture envisaged.

5.1 DARWIN Guidance, Navigation and Control Mode Analysis

The selection of DARWIN as the reference mission for this research project was made during the initial proposals for the research. DARWIN encompasses all the motivations of the stakeholders for this research project since it represents a mission devoted to finding and characterising extra-solar planets, it involves multiple spacecraft flying in formation performing separated spacecraft interferometry and various phases of the mission require formation flying manoeuvre planning and execution. There are many guidance, navigation and control (GNC) modes that require translational formation flying manoeuvres so it is necessary to select one mode to study regarding the manoeuvre planning. The GNC modes for DARWIN are shown in Figure 5-1. As mentioned in sub-section 2.2.4 only the operational loop consisting of the baseline control mode, fringe acquisition mode, normal observations mode and reconfiguration mode are of interest since this project is concentrating on the normal mission operations (i.e. no Launch and Early Operations (LEOP) or safe mode).

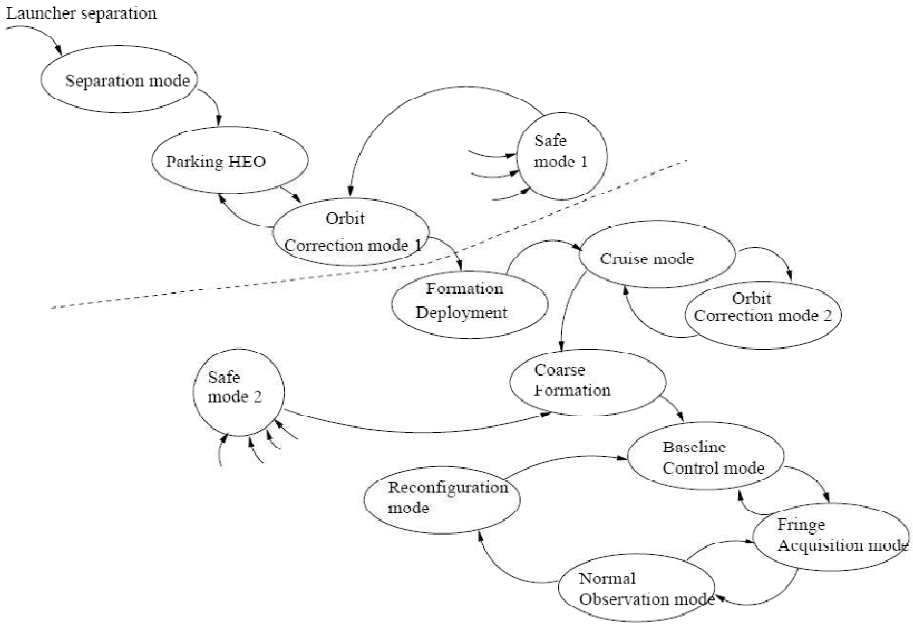

Figure 5-1 GNC Modes for the DARWIN mission (Karlsson, et al., 2004) reproduced from Figure 2-6

For each GNC mode of the operational loop the different types of potential translational manoeuvre to be performed by each spacecraft and the formation as a whole needs to be defined. These manoeuvres are (with respect to the DARWIN mission):

- Station-keeping – a 3 translational degrees of freedom (DOF) manoeuvre designed to maintain spacecraft position relative to a reference point external to the formation
- Formation-keeping – a 3 translational DOF manoeuvre designed to maintain the spacecraft position relative to a reference point within the formation
- Resize – a 1 translational DOF manoeuvre within the formation plane designed to increase or decrease the spacecraft's distance from a reference point
- Rotate – a 2 translational DOF manoeuvre within the formation plane designed to change the position of the spacecraft relative to the reference point whilst maintaining a constant distance from that point
- Retarget – a 6 DOF (in translation and attitude) manoeuvre designed to change the pointing direction of the individual spacecraft and the formation plane. Each individual spacecraft's position relative to the reference point before and after the manoeuvre remains the same
- Reconfigure – a 3 translational DOF manoeuvre designed to change the position of the spacecraft relative to the reference point whilst maintaining spacecraft pointing
- Slew – a 6 DOF manoeuvre designed to maintain individual spacecraft and formation plane pointing in the L_2 reference frame

These seven manoeuvre types can be used in each of the GNC modes of the DARWIN operational loop as in Table 5-1.

Table 5-1 Manoeuvres for the GNC Modes of DARWIN and planning requirements

GNC Mode			
Reconfiguration	Baseline Control	Fringe Acquisition	Normal Observation
Station-keeping Retarget Reconfigure	Station-keeping Formation-keeping	Station-keeping Formation-keeping Resize	Station-keeping Formation-keeping Rotate Slew
Planning Required?			
Yes	No	Possibly	Yes

In the reconfiguration mode the formation pattern and pointing direction is changed so that a new star can be analysed. This requires a combination of both reconfigure manoeuvres, to change the formation patterns, and retarget manoeuvres, to change the pointing direction of the spacecraft and formation. The relative positions of the spacecraft in the formation do not need to remain fixed during this reconfiguration manoeuvre. Manoeuvre planning is required for this mode since all the spacecraft may be required to perform large translational manoeuvres that will require safe trajectories that can be optimised for fuel management and manoeuvre duration.

In the baseline control mode the relative position constraints of the formation are tightened so that the relative position of the spacecraft can be brought to within 1cm. The type of manoeuvre required to perform this is simply the formation-keeping manoeuvre. The formation-keeping manoeuvre cannot be planned in advance since its requirements depend on the real-time perturbation environment. The formation-keeping manoeuvre can be optimised however through the use of an optimal control law.

The fringe acquisition mode involves two steps. In the first step formation-keeping manoeuvres are used to hold the formation steady whilst the optical delay line (ODL) on the Beam Combiner Spacecraft (BCS) uses its full stroke to acquire the interferometry fringe. Only formation-keeping is required for this part and so manoeuvre planning is not required. In the second step resize manoeuvres are performed by the Telescope Spacecraft (TS) in order to move the ODL back to the centre of its stroke. These resize manoeuvres could utilise manoeuvre planning to aid fuel management however since the entire manoeuvre is likely to involve a translation of less than 1 cm is seems unlikely than any real fuel management gain could be achieved with respect to the complexity of the optimisation process. For the purposes of this research therefore the fringe acquisition mode is considered to require no manoeuvre planning

The final GNC mode from the operational loop is the normal observation mode. This is the mode in which the scientific observations are made and a number of translational manoeuvres are required. Firstly, formation-keeping manoeuvres are required to ensure the ODL does not have to compensate for more than 1cm of relative spacecraft drift during the observation. Again no manoeuvre planning can be applied to the formation-keeping manoeuvres. For planet detection observations the entire formation is required to rotate within the formation plane 180° (see sub-section 2.2.2). Rotate manoeuvres are used to achieve this. These manoeuvres can be planned for fuel management since, although the relative positions of the spacecraft remain fixed with respect to each other (i.e. no risk of collisions), the reference point for the rotation does not have to be fixed. This type of manoeuvre planning can be found in Beard, McLain and Hadaegh (2000). Finally within the normal observation mode is the requirement for slew manoeuvres. For long duration spectroscopy observations the formation plane is likely to require slewing to maintain its required perpendicularity to the target's position vector. This slew manoeuvre requires both translational and attitude manoeuvres by the individual spacecraft and is similar to the retarget manoeuvre (except the slew rate will be significantly slower than the retarget rate). Again, these manoeuvres can be planned for fuel management in a similar way to that presented by Beard, McLain and Hadaegh (2000).

For all the GNC modes presented in Table 5-1 it is also possible to perform station-keeping manoeuvres. The station-keeping manoeuvre is designed to ensure the spacecraft remain close to a pre-determined trajectory that ensures the formation stays within the vicinity of the L_2 point. The station-keeping manoeuvre involves the effective simultaneous translation of the entire formation. The timing of the station-keeping manoeuvre can be planned to minimise the frequency and associated fuel costs. However, due to the stringent position accuracy required for three of the GNC modes (baseline control, fringe acquisition and normal observation mode) it is sensible to avoid any unnecessary stress on the control system. One way to assist this is to prevent the station-keeping manoeuvres from being

executed in any GNC mode other than the reconfiguration mode. Although this may limit the optimality of station-keeping manoeuvres (due to the addition of a temporal constraint), planning can be performed to mitigate this.

Based on the GNC modes analysis in the previous paragraphs the author decided to focus the research on manoeuvre planning for the reconfiguration mode for DARWIN. This mode allows the greatest flexibility in spacecraft movement during the manoeuvre and so presents a greater challenge to optimise for fuel management and manoeuvre duration. Although the normal observation mode does involve manoeuvres that could be optimised the author believes that the method presented in Beard, McLain and Hadaegh (2000) is sufficient to perform this task and little improvement could be made. In selecting the reconfiguration mode for this research the focus can be drawn to planning for manoeuvres that do not require the maintenance of fixed relative positions between the spacecraft. This condition allows the research to be generalised to include other types of formation flying mission where unconstrained reconfiguration of the formation is required.

5.2 Formation Flying Concepts Analysis

Chapter 3 introduced a number of concepts that are enabled with formation flying missions and not present in single-spacecraft missions. Before the manoeuvre planning architecture can be defined it is necessary to analyse and select which variations of these concepts to pursue.

5.2.1 Autonomous Manoeuvre Planning Organisation

In sub-section 3.3.1 the concept of autonomous manoeuvre planning organisation was introduced. Examples defined included traditional, top-down, centralised and de-centralised organisations. Each of these organisations describe how autonomous manoeuvre planning can be organised amongst formation members with various levels of planning capabilities defined, i.e. no planning (S_{NP}), individual planning (S_{IP}) and full planning (S_{FP}). The trade-off in selecting a manoeuvre planning organisation relates to ground operations cost and complexity, ground communications bandwidth requirements, on-board software complexity, on-board hardware capability and inter-spacecraft communication bandwidth capability and requirements.

For the purposes of this research project the author has chosen to adopt the top-down autonomous manoeuvre planning organisation as shown in Figure 5-2. The top-down organisation involves the ground station maintaining a low bandwidth link to a spacecraft capable of formation planning (S_{FP}). The S_{FP} maintains full knowledge of the other spacecraft in the fleet and can plan for the entire formation. Once the plan is complete the S_{FP} cross-links the required elements of the plan to each of the S_{NP}s who execute it accordingly. This organisation was chosen because:

- it provides the required level of autonomous manoeuvre planning as laid down in the research project's problem statement and aims without being too complex

- the brief communication sub-system description in Karlsson, et al. (2004) indicates only the BCS will communicate with the ground station whilst the TSs use the RF metrology signal for inter-spacecraft communications, therefore discounting the de-centralised approach
- the author wanted to concentrate on the structure of the manoeuvre planning algorithms themselves not how the computational resources would be managed
- simulating the centralised or de-centralised organisations would involve computing resources/expertise unavailable to the author

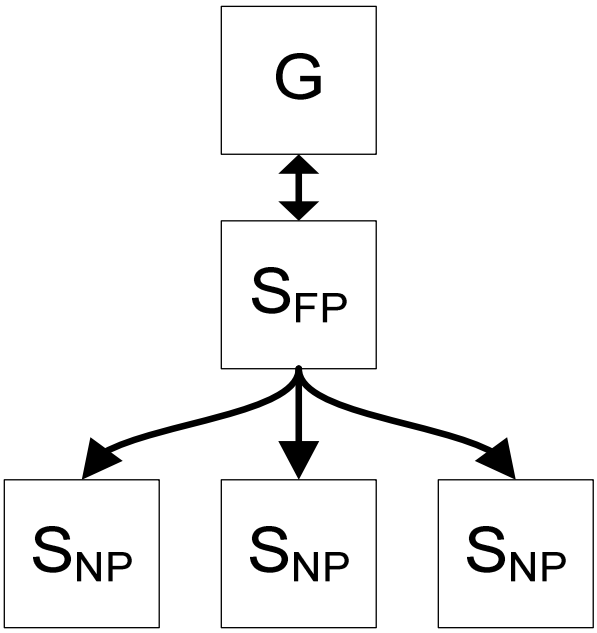

Figure 5-2 Top-down autonomous manoeuvre planning organisation. G represents the ground station, S_{FP} is a spacecraft with formation planning capability and S_{NP} is a spacecraft with no planning capability.

For DARWIN it is natural to assume the S_{FP} spacecraft will be the BCS since it maintains constant RF links with the TS spacecraft and is the only spacecraft with a ground station link. As long as the inter-spacecraft and ground station links remain available however it is feasible to use one of the TSs as a backup S_{FP} should the BCS lose its ability to plan manoeuvres (but remain fully functional in every other aspect).

5.2.2 Control Co-ordination

The concept of control co-ordination for spacecraft formation flying was introduced in sub-section 3.3.2 describing the absolute method, leader-follower, behavioural approach, virtual structure and virtual centre. The trade-off in selecting the type of control co-ordination required includes complexity of the approach, the ability to optimise manoeuvres, inter-spacecraft communication capability and requirements and the size of the manoeuvre error margins.

For the purposes of this research the author chose to adopt the leader-follower control co-ordination approach. Within this approach one of the spacecraft is designated the leader

whilst the others are designated followers. The leader tracks a predefined trajectory and performs station-keeping manoeuvres to maintain this trajectory over time. The follower spacecraft tracks the position of the leader spacecraft and maintains a relative position. This control co-ordination approach was chosen because

- its simplicity for use within the reconfiguration GNC mode
- the behavioural approach is difficult to optimise manoeuvres and it requires de-centralised organisation
- the virtual structure is too restricted and more suited to the normal observation GNC mode
- in the virtual centre it is harder to optimise manoeuvres and it requires significant inter-spacecraft communication and centralised organisation

The greatest issue with the leader-follower control co-ordination approach is its reliance on the leader and the single-point-of-failure characteristic that this implies. For DARWIN however, the BCS already represents a single-point-of-failure for the formation as it is the only spacecraft in the formation that can combine the science beams and disseminate the resulting images to the ground station. Placing the BCS as the leader of the formation for GNC purposes therefore is immaterial since the loss of the BCS of any reason would constitute the loss of the mission.

5.2.3 Control Strategy

Sub-section 3.3.3 details a number of autonomous control strategies that appear in the literature designed to execute the manoeuvres output from a planning algorithm. While this research project will not implement a formation flying control strategy to simulate the planned manoeuvres it is important for the definition of the manoeuvre planning architecture to have a control strategy in mind. This research focuses on the GNC reconfiguration mode for DARWIN. This involves unconstrained translational manoeuvres between the spacecraft using the coarse metrology and thruster configurations introduced in sub-section 2.2.3. The specifications of the proposed RF metrology and ion thrusters to be deployed during the reconfiguration mode should be able to maintain the positional and velocity errors to 1 cm and 250 μms^{-1} respectively. In reality however these are tight constraints for the reconfiguration mode.

Based on the proposed manoeuvre planning architecture, the proposed formation flying control co-ordination method and a control strategies literature review the author decided to adopt a dual control strategy approach to the problem. In this approach the BCS will perform its manoeuvres through an open-loop thruster timing command structure with no explicit feedback from the other spacecraft in the formation for nominal manoeuvres. The TS spacecraft will perform their manoeuvres through a closed-loop command structure with relative position and velocity feedback from the BCS for nominal manoeuvres.

The choice for the BCS control strategy to be open-loop stems from the lack of a reference point to base the BCS manoeuvre on. Closed-loop manoeuvres require some reference to provide feedback. In deep-space at L_2 there will be no external references to

measure the BCSs position except for the other spacecraft in the formation. These spacecraft can only provide information relative to the formation, not to inertial space. To be able to perform BCS manoeuvres relative to inertial space requires ranging data from the ground station on Earth. As this will not always be available during a manoeuvre a combination of previous ranging data and dead-reckoning must be used for the spacecraft to autonomously calculate its position relative to inertial space. The BCS control strategy uses open-loop thruster timing schedules to ensure the spacecraft follows the trajectory planned by the manoeuvre planner. This is the same for both station-keeping and formation reconfiguration manoeuvres. Using this simple method for the BCS manoeuvres reduces the risk of errors which is important since the BCS acts as a leader to the other spacecraft in the formation.

With the BCS forming the inertial reference for the formation the TSs can adopt autonomous closed-loop control strategies to perform their manoeuvres. For reconfiguration manoeuvres the output from the manoeuvre planner will contain a desired trajectory to follow relative to the BCS. The combined translation of the BCS and TS will create a trajectory for the TS in inertial space that will be safe and fuel managed. For station-keeping (i.e. controlling the position of the formation relative to a reference trajectory around L_2) the TSs simply have to formation-keep with the BCS (that will actually be performing the planned station-keeping manoeuvre). This control strategy for the TSs can be simply designed, tuned and implemented. The major flaw with this approach however is its reliance on the BCS. Should the BCS manoeuvre not follow the planned thruster schedule (for any reason) then the relative trajectories of the TSs may not lead to safe and fuel managed trajectories in inertial space. Mitigation for this will be covered in the next sub-section. As detailed in sub-section 3.3.3 there are a whole host of closed-loop control strategies that could be adopted for this problem. The LQG controller introduced by Lagadec, Lebas and Ankersen (2003) and Davidson, et al. (2006) is designed specifically for the DARWIN mission and performs to the requirements of the reconfiguration mode.

5.2.4 Manoeuvre Error Mitigation

Manoeuvre planning for spacecraft formation flying ensures that reconfiguration manoeuvres can be performed safely and be optimised for fuel management. These nominal manoeuvres satisfy the requirements for the reconfiguration manoeuvres but depend heavily on the flawless operation of every spacecraft in the formation and a perturbation environment similar to that employed by the planner. Should one spacecraft malfunction during the execution of a manoeuvre it may pose a collision risk or it may move on to an unrecoverable trajectory. In addition, with the trajectories planned to a certain level of accuracy to gain fuel optimisation any deviation from that trajectory can pose a serious risk to the other spacecraft. Though manoeuvre error mitigation cannot be planned for at the manoeuvre planning stage it is important to have an understanding of the type of error mitigation system that can be adopted. This sub-section introduces such a system and completes this discussion of the GNC for the reconfiguration mode for DARWIN.

There are two scenarios that can be envisioned; a malfunction of the BCS and a malfunction in one or more of the TSs. If a problem should arise with the BCS and its ability to execute the planned manoeuvre then no collision risk is imposed. Each TS follows a

trajectory relative to the BCS. There is no risk of collision because any error in the BCSs inertial space manoeuvre will be echoed in the inertial space manoeuvres of the TSs. An error in the BCS manoeuvre will, in effect, force the TSs to perform formation-keeping manoeuvres in addition to the reconfiguration manoeuvre. Whilst this does not aid the fuel management of the manoeuvre at least the formation remains safe. It does however require that the TSs are able to remain tracking the errant BCS. The second scenario involves the malfunction of one or more of the TSs. Should this happen then the situation is much more serious as the collision risk is much higher.

There are two required elements for a manoeuvre error mitigation system. The first is spacecraft tracking whilst the second is mitigation actions. During the execution of the manoeuvre the TSs can be easily tracked by the BCS using the RF metrology system (since the laser metrology will not be usable due to the misalignment of the spacecraft). The BCS is therefore able to compare the planned TS trajectory against the actual TS trajectory. Should the two trajectories diverge by a significant amount then an anomaly on the TS can be assumed. This anomaly can either be due to a malfunction of the TS or the TSs inability to track the BCS due to a malfunction by the BCS. For a BCS malfunction then the course of action is simply to stop the manoeuvre on all the spacecraft and allow the formation to enter a safe mode where the TSs formation-keep with the BCS in a pre-defined ‘safe’ configuration (like the triTTN). This ensures that the formation does not drift too far apart to reinitialise the mission should the BCS anomaly be fixed. For a malfunction on one of the TSs the situation is more complex. The same safe mode process as before could be used but since the TS is likely to be unable to formation-keep due to whatever caused the manoeuvre error the collision risk remains and if the spacecraft drifts away from the formation it may not be able to return once the anomaly has been fixed. In this situation the solution is for the formation to enter a different safe mode where the formation-keeping is switched relative to the anomalous TS. This ensures the collision risk is mitigated since the BCS and remaining TS will maintain a fixed relative position with respect to the anomalous TS and ensures the formation does not drift apart. Whilst both these mitigation actions will result in poor fuel management the risk of losing the entire mission is much more pertinent.

5.3 Optimisation Techniques Analysis

The final aspect to explore before defining the Manoeuvre Planning Architecture (MPA) is the use of optimisation and the types of algorithms to be employed. Chapter 4 introduced the concept of optimisation and listed a number of useful optimisation techniques that are employed to solve modern optimisation problems. Whilst optimisation is an important aspect of this research project, and the author found this part of the review particularly interesting, it was noted that the research project was not about developing optimisation routines for spacecraft formation flying manoeuvre planning but developing a manoeuvre planning architecture that optimised formation flying manoeuvres. There are two aspects of the optimisation for this research project that need to be examined; the type of optimisation attempted and the optimisation techniques used. This sub-section discusses these aspects.

Chapter 4 introduced three optimisation types differentiated by the number of objectives and the number of variables that affect those objectives. The problem identified for this

research project is clearly multi-objective with numerous, often conflicting, objectives to optimise. These include manoeuvre schedule optimisation, time minimal reconfiguration manoeuvres, fuel minimal reconfiguration manoeuvres, fuel balancing reconfiguration manoeuvres and optimal station-keeping manoeuvres. Though this is identified as a multi-objective problem the additional desire to solve this on-board a computer-limited single spacecraft restricts the computational complexity that can be deployed. Complex multi-objective problems typically require large computational resources over a long period of time to get satisfactory results. This approach however is impractical for an optimal manoeuvre planner that is required to find good solutions in real-time.

For this research project the author decided to examine the manoeuvre planning problem as a series of individual single-objective (but multi-variable) problems that are incorporated within a Manoeuvre Planning Architecture (MPA). Whilst this approach is unlikely to be able to find the global optimum for the manoeuvre planning problem a combination of the individual solutions should lead to a good solution within the computation complexity constraints. Another advantage of this approach is the ability to analyse each individual component of the optimisation problem easily and identify problems and limitations quickly.

With the optimisation types detailed the second aspect to investigate is which optimisation technique to employ. This was mainly driven by the stakeholder requirement that the coding be done within the MATLAB® software environment (The Mathworks Inc., 2006). In this software environment there are a number of built-in optimisation routines within the Optimisation Toolbox (Coleman and Zang, 2005) and the Genetic Algorithm and Direct Search (GADS) Toolbox (Abramson, 2006).

Table 5-2 Optimisation routines used in this research project

MPA Optimisation Module	Optimisation Method	Technique	MATLAB toolbox
SOM (Chapter 7)	bespoke (ACO inspired)	quasi-swarm	n/a
PAM (Chapter 8)	PatternSearch (PS)	direct search	GADS
TDM (Chapter 9)	PatternSearch (PS)	direct search	GADS
	Genetic Algorithm (GA)	evolutionary	GADS
	'*fmincon*'[2]	gradient based	optimisation
SKM (Chapter 10)	bespoke (literature)	bespoke	n/a

Whilst designing the MPA the author was confronted with three different optimisation forms. In the first form the author was able to describe the optimisation problem as a mathematical cost function, as detailed in sub-section 4. For this form of optimisation problem it was possible to use the embedded MATLAB® optimisation routines from the Optimisation and GADS toolboxes. The optimisation routines used and where in the MPA are shown in Table 5-2. The second form of optimisation problem encountered was where the author was unable to describe the optimisation problem as a mathematical cost function. As the MATLAB optimisation routines cannot be implemented without this format the author was compelled to write a bespoke algorithm to solve the optimisation problem. This was used

[2] '*fmincon*' is a non-linear constrained gradient-based optimisation technique found in MATLAB's Optimisation toolbox

to solve one optimisation problem within the MPA and was inspired by the Ant Colony Optimisation method detailed in sub-section 4.4.5.2. The final optimisation problem form encountered was when the problem and a bespoke solution method were found in the literature. No design therefore was required of the author. The use of all the optimisation routines is detailed further in Chapters 7-10.

5.4 Chapter Summary

Using the background information provided in Chapters 2-4 analysis has been performed on the mission aspects that affect the design of a manoeuvre planning architecture and decisions made. The work in this research project has been restricted to the reconfiguration mode of the envisaged DARWIN GNC system since it offers the greatest potential for manoeuvre optimisation and spacecraft proximity violations. The other modes within the operation GNC are much more restrictive in terms of spacecraft flexibility during manoeuvres and the author believes that existing planning methods are suitable for these modes. For the GNC analysis it was also decided to restrict the timing of station-keeping manoeuvres so that they only occur during the reconfiguration mode. Whilst this may make the manoeuvres sub-optimal is it essential to reduce load of the spacecraft control systems during the delicate observation phases.

The top-down autonomous manoeuvre planning organisation was chosen as it represented the simplest structure whilst providing the autonomy required in the problem statement. In addition the top-down organisation fitted best with the envisaged DARWIN communications sub-system from Karlsson, et al., (2004). For control co-ordination the chosen method was the leader-follower approach for its simplicity in comparison to the other methods introduced. Though this approach suffers from a single-point-of-failure limitation, the Beam Combiner Spacecraft (BCS) already represents this limitation for the entire mission. Using the BCS as the leader for the manoeuvres removes this issue from the control co-ordination trade-off.

The envisaged control strategy differs between the BCS and Telescope Spacecraft (TS). With no external position reference the BCS uses open-loop thruster timing schedules to execute the planned reconfiguration and station-keeping manoeuvres. Inertial navigation for the BCS is provided by previous ranging data and autonomously calculated dead-reckoning. With the BCS as the leader of the formation, the TSs can use it as a reference for their manoeuvres. The TSs adopt a closed-loop LQG controller to track and follow the movement of the BCS. During station-keeping manoeuvres the TSs formation-keep with the BCS. During reconfiguration manoeuvres the TSs follow a planned trajectory relative to the BCS. This chapter also introduced a manoeuvre error mitigation strategy to address the problem spacecraft malfunction during a manoeuvre. Whilst not part of manoeuvre planning, the error mitigation strategy allows the Manoeuvre Planning Architecture (MPA) to ignore such concerns and plan manoeuvres as optimally as possible.

The final section of this chapter dealt with the issue of optimisation within the MPA. A combination of stakeholder requirements and problem statement assessment lead the author to adopt an optimisation approach that views the MPA a series of single objective multi-variable problems. This approach improves transparency for analysis purposes and reduces

the computational burden of the MPA. The MATLAB® software environment and the associated Optimisation and Genetic Algorithm and Direct Search (GADS) toolboxes were chosen as the platform to implement and analyse the MPA.

This chapter concludes the first section of this research thesis. In Chapters 6-0 the MPA is developed and each of the associated ‘optimisation modules’ are introduced, designed and analysed.

6. MANOEUVRE PLANNING ARCHITECTURE DESIGN

As identified in the problem statement (sub-section 1.3) there are a number of manoeuvre planning issues that can be applied to the manoeuvre planning problem presented by a mission like DARWIN. These planning constraints need to be addressed within a Manoeuvre Planning Architecture (MPA) that can not only plan manoeuvres to effectively include these constraints but also plan manoeuvres efficiently in terms of the planning time required. This chapter introduces a number of optimisation modules that deal with the above constraints and presents various MPAs in which these modules can operate.

6.1 Manoeuvre Planning Optimisations and Constraints

The manoeuvre planning optimisations that can be realised by an effective manoeuvre planning strategy are as follows:

- Maximising mission science returns – essential for an astronomy observation mission such as DARWIN to justify the cost of developing and operating the system.

- Manoeuvre duration minimisation - the time taken to perform manoeuvres detracts from the available time for useful science observations which, in a time limited mission, can reduce the possible scientific returns.

- Formation fuel consumption minimisation - the amount of propellant used during the mission is an important factor affecting the cost and mass budget

- Formation fuel balancing - although the spacecraft work together within the formation, during a formation reconfiguration manoeuvre they act independently to each other. Therefore each spacecraft may use different amounts of fuel to achieve the required formation configuration. Over time, it is possible for some spacecraft in the formation to experience fuel-starvation whilst others may have plenty of fuel. This would reduce the effectiveness of the formation and its ability to achieve the science goals and possibly end the mission prematurely. The manoeuvres must therefore be planned to prevent single-spacecraft fuel starvation and promote fuel-balancing throughout the formation.

- Libration-point station-keeping – the mission requires the formation to remain within the vicinity of L_2. This cannot be achieved without active corrective manoeuvres that influence the time, fuel and fuel balancing optimisations described previously. The requirement for station-keeping therefore must be included in any manoeuvre planning strategy.

The manoeuvre planning constraints that need to be addresses by any manoeuvre planning strategy are as follows:

- Collision avoidance - spacecraft manoeuvring in close proximity to each other pose a collision risk. Manoeuvre planning must therefore ensure the calculated trajectories do not violate collision avoidance criteria.

- Thruster plume avoidance - thruster exhaust plumes from manoeuvring spacecraft can pose a risk to spacecraft surfaces. Manoeuvre planning must therefore incorporate a thruster plume avoidance strategy.

6.2 Manoeuvre Planning Systems Model

The design of the Manoeuvre Planning Architecture (MPA) requires a detailed analysis of the manoeuvre planning system and a trade-off analysis of the various solutions presented. Figure 6-1 illustrates a systems overview of the manoeuvre planning problem. At the time of manoeuvre planning the formation is in normal operations mode and has a fixed pointing direction and configuration. Through a combination of ranging and dead-reckoning through the gravitational model the absolute spacecraft positions can be estimated at the time the next reconfiguration manoeuvre is to commence. At this point the reconfiguration manoeuvre is executed to achieve the planned final absolute spacecraft positions (that satisfy the optimisation goals). The final absolute spacecraft positions must also satisfy the observation requirements (i.e. that the final relative spacecraft positions position the spacecraft correctly for the formation plane to point at the chosen target star and the formation configuration be correct for the chosen observation on that star). This model indicates two distinct (but interlinked) decision processes that must be followed to successfully plan an optimised manoeuvre; the selection of the target star/observation combination and the selection of the individual spacecraft manoeuvres.

6.2.1 Target Star/Observation Selection Model

The first decision process involves the selection of the target star/observation. Figure 6-2 illustrates the systems model for the target star/observation selection process. The model shows the available input information, calculations data flow and constraints. The target star/observation selection addresses the planning goal to maximise the mission science returns. This can be achieved by maximising the number of individual observations performed by the formation during the lifetime of the mission and by extension, minimising the total time required to perform each observation. This total time is the sum of the reconfiguration manoeuvre duration, the formation calibration time and the actual observation duration. The reconfiguration manoeuvre duration in-turn depends partly upon the selection of the target star and the target observation to be performed on that star (i.e. detection or spectroscopy). The latter also determines the observation duration.

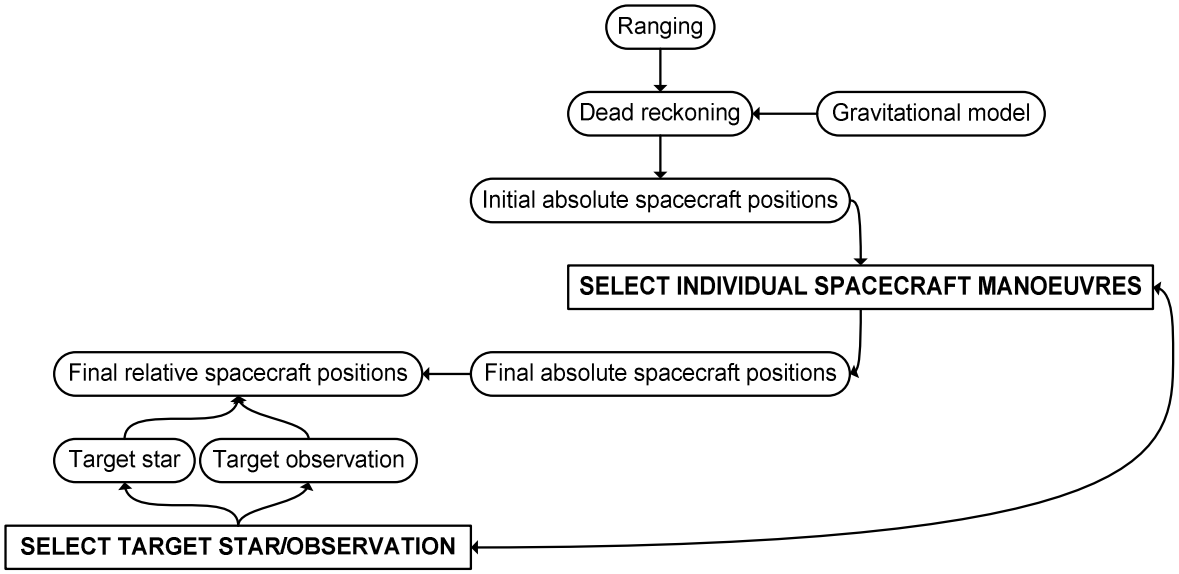

Figure 6-1 Manoeuvre planning systems overview

In addition there are a number of constraints that contribute to make the target star/observation selection a complex problem. These factors are as follows:

- Calibration time – After the manoeuvre is completed there will be a calibration time associated with star acquisition and stabilising the formation. This time will depend on the type of star being acquired and the formation configuration being adopted.
- Formation pointing constraints - The formation must remain within ±45° of the anti-Sun vector at all times. As the formation will remain in the vicinity of the L_2 point in the Sun/Earth-Moon system and the L_2 point orbits the Sun with the same period as the Earth this implies that the field of view (FOV) for the mission rotates about the Sun at a rate of 0.986 °/day. Stars move in and out of the FOV over the course of the mission and remain within the FOV for approximately 91 days per year. This complicates the choice of science task as a task cannot be selected if the star moves out of the FOV before the observation can be completed. This is particularly pertinent to some of the spectroscopy tasks that may take up to 85 days to complete.
- Detection task scheduling – For a planet to be found and its orbital characteristics calculated three detection observations are required to be performed on each star. These three observations need to be separated in time so that any potential planet can be observed at different points around its orbit. The duration of this separation depends on the spectral type of the parent star and so will vary from star to star.
- Planet detection analysis time – Once the three detection observations have been performed there will be a period of time required to confirm planet detection and calculate the orbital parameters. No further observations can be scheduled on the parent star until this has been completed.
- Planet orbital characteristics – should a confirmed planetary orbit involve occultation of and by the parent star it becomes desirable to schedule the spectroscopy observations to avoid these occultation periods. For the longer spectroscopy

observations and certain star/planet/orbit configurations this may be impossible so the observation may need to be stopped whilst the occultation is occurring. During this time it may be possible to re-task the formation to another star and return to the original star to complete the long duration spectroscopy at a later date.

- It is unclear from the literature whether planet detection is more important than planet spectroscopy in terms of science task selection however during the course of the mission it is reasonable to assume that mission scientists will want to influence the task schedule to ensure certain observations get completed. This can be achieved by adding a weight to the desired observation to allow the SOM to optimally ensure the task is included in the schedule.

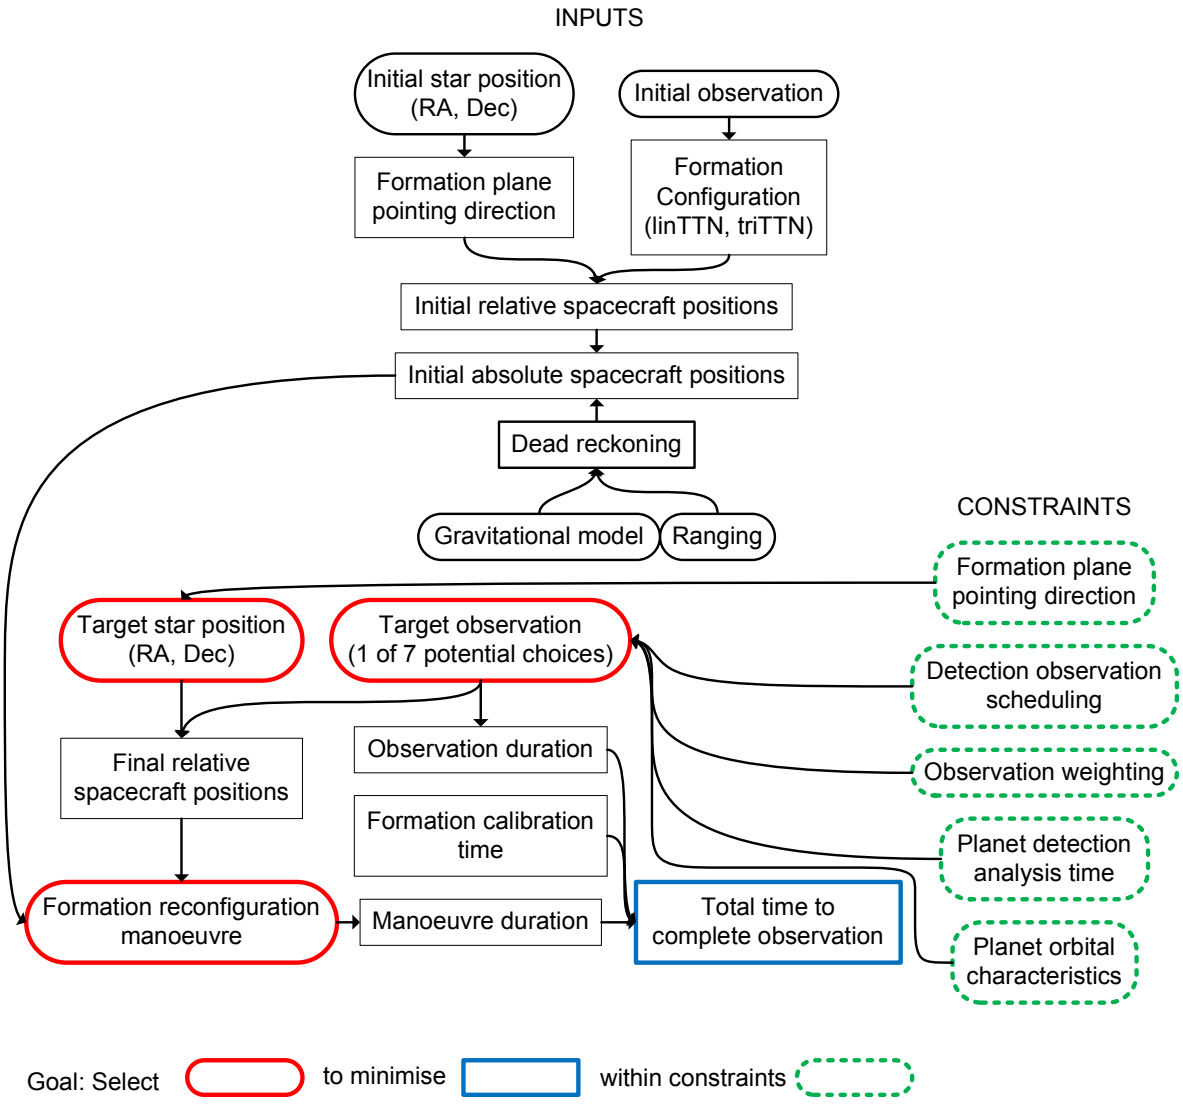

Figure 6-2 Target star/observation selection model

As can be seen from Figure 6-2 the selection of the target star/observation is a complex problem made even more so by the complexity of the individual spacecraft manoeuvres selection.

6.2.2 Individual Spacecraft Manoeuvres Selection Model

The second decision process involves the selection of the individual spacecraft manoeuvres. Figure 6-3 shows the selection model for the individual spacecraft manoeuvres. The model shows the available input information, calculations data flow and constraints. The individual spacecraft manoeuvres selection addresses the planning goals to optimise the formation manoeuvre duration, manoeuvre fuel consumption and formation fuel balancing whilst ensuring spacecraft collisions and thruster plume impingements are avoided. This can be achieved through the selection of the final absolute spacecraft positions and the thrust duration, timings and magnitudes required attaining those positions.

The relative positions the spacecraft are required to take in order to perform an observation depend on the target star/observation chosen. First, the spacecraft must form a formation plane that is perpendicular to the direction vector of the star. Within that plane the spacecraft must be in the correct configuration (linTTN or triTTN, Figure 2-2 and Figure 2-3) and at the correct relative distances from each other (the formation baseline). The selection of the absolute spacecraft positions must conform to the requirements of the relative positions. In inertial space there are an infinite number of possible absolute spacecraft positions that could satisfy these relative spacecraft positions. In addition therefore the absolute positions must be chosen to reflect the optimisation goals of the manoeuvre and the manoeuvre constraints.

The individual spacecraft thrust duration, magnitude and timing schedule determines the fuel consumption for that manoeuvre and in turn the overall fuel balancing for the formation. These thrust parameters also determine the trajectory that each spacecraft follows from the initials to the selected final positions. The trajectory will also be affected by the gravitational model used within the planning environment. These trajectories must in turn comply with the manoeuvre avoidance constraints and the absolute trajectory maintenance requirements.
As can be seen from Figure 6-3 the selection of the individual spacecraft manoeuvres is a complex problem made even more so by the complexity of the target star/observation selection.

6.3 Manoeuvre Planning Systems Trades

The design of the manoeuvre planning architecture requires the system models given in Figure 6-1, Figure 6-2 and Figure 6-3 to be analysed and trades made to maximise the system efficiency. The systems trades begin by analysing the selection of the target star/observation.

6.3.1 Target Star/Observation Selection Trades

For this selection the goal is to minimise the total time to complete the observation. This time is the sum of the observation time, the formation calibration time and the formation reconfiguration manoeuvre time. Table 6-1 shows the order of the time component affecting the target star/observation selection process. The observation time has the order of days as introduced in Table 2-2. The formation calibration time assumes a manoeuvre requiring the full 1 cm stroke for fringe acquisition using 50 μN thrust on a 1000 kg spacecraft and a bang-bang thrust profile (see Equation (8.11) and it's derivation in Chapter 8) and has the order of

a few minutes. The manoeuvre time is likewise calculated assuming a 100m manoeuvre using 6 mN thrust on a 1000kg spacecraft and bang-bang thrust profile. This manoeuvre has the order of a few hours and represents the maximum duration manoeuvre any one spacecraft is likely to make (since the maximum baseline for any formation configuration is ~100 m). All these values are taken from Karlsson, et al. (2004) and are detailed further in Chapter 2.

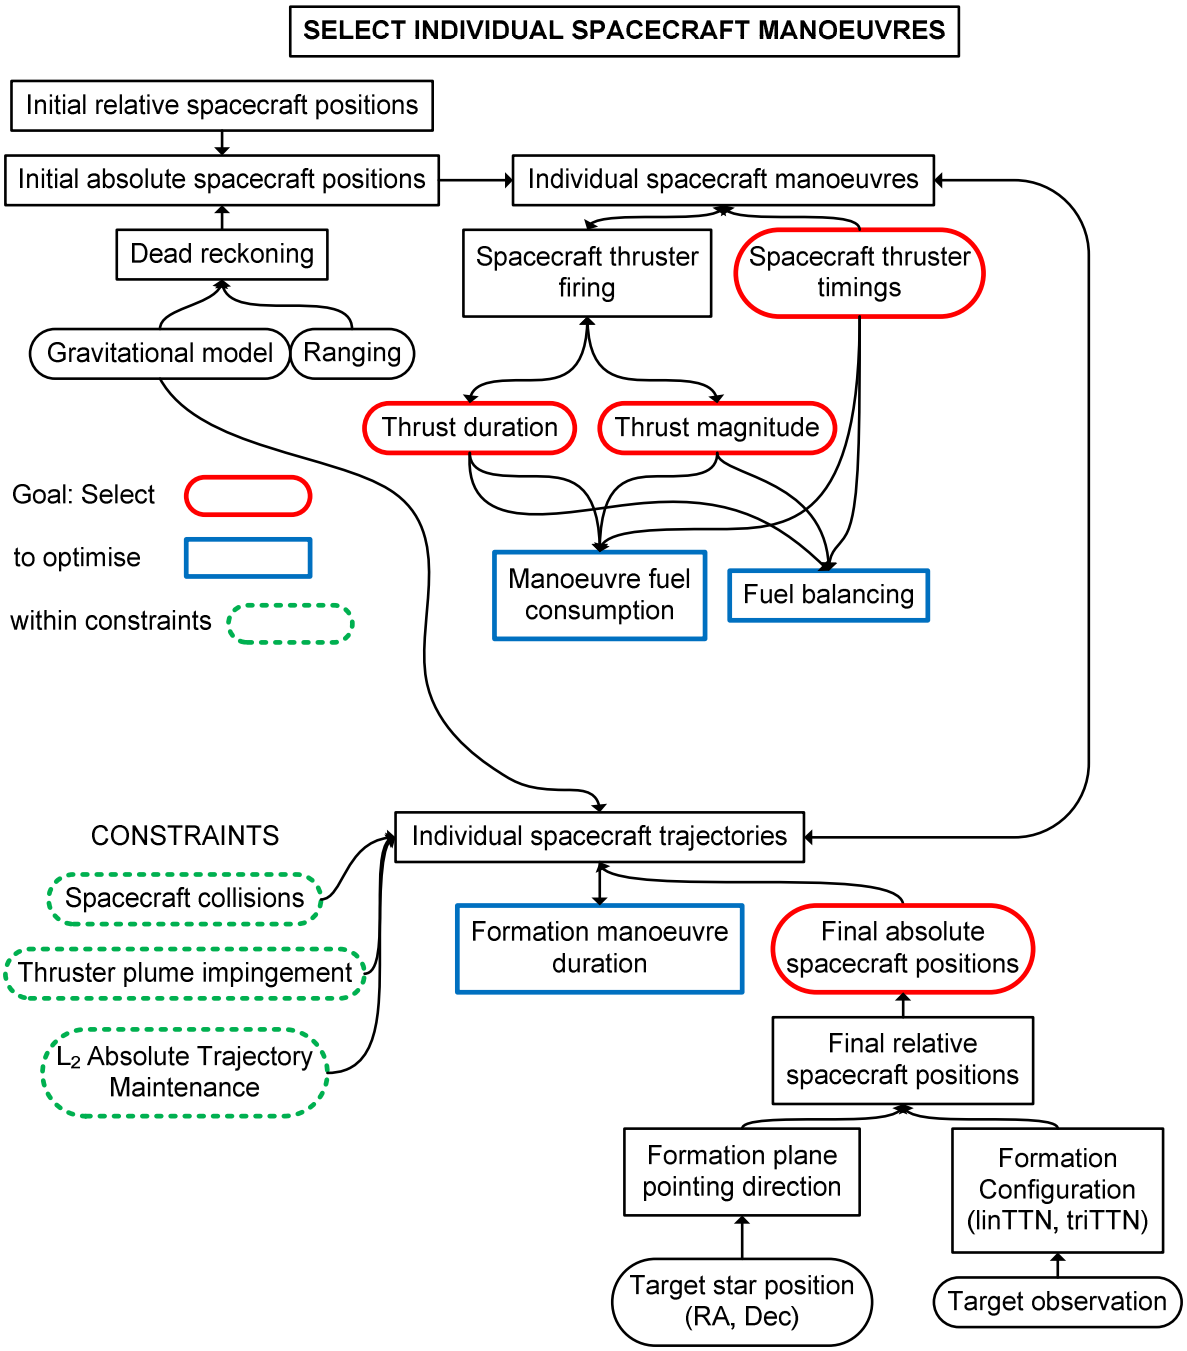

Figure 6-3 Individual spacecraft manoeuvres selection model

Table 6-1 The order of the time components influencing the target star/observation selection

	Order	Example
Observation time	days	observation times range from 8 hours to ~91 days
Formation calibration time	mins	~15 mins for 1 cm fringe acquisition
Manoeuvre time	hrs	max manoeuvre time ~2¼ hrs

Table 6-1 shows that the order of the observation time far out-weighs the order of the other two components affecting the total time to complete the observation. Thus any efforts to minimise the total time to complete the observation must clearly be focussed on minimising the observation duration which is a function of the target star/observation only. As the selection of the individual manoeuvres adds an extra layer of complexity to the selection of the target star/observation the former selection can be de-coupled from the latter by assuming a reconfiguration manoeuvre duration that does not depend upon the selection of the individual spacecraft manoeuvres. This is shown in Figure 6-4.

The simplified target star/observation selection model in Figure 6-4 shows how the manoeuvre duration can be calculated solely as a function of the target star selection. This selection model assumes that the duration of a reconfiguration manoeuvre depends on the scale of the formation manoeuvre required to rotate the formation plane towards the target star. Since no formation plane rotation will be greater than 90° (due to formation pointing constraints) this is a reasonable assumption to make even if the relative spacecraft positions during the manoeuvre are flexible.

The simplified model in Figure 6-4 also omits the formation calibration time as an influencing factor on the total time to complete the observation. Calculation of the calibration time depends upon the errors in the final relative positions of the spacecraft and, as demonstrated in Table 6-1, only affects the order of the total time to complete the observation by a few minutes. It is reasonable to assume that omission of this contribution will have little effect on the results of the selection but a greater (positive) effect on the time required to achieve that solution.

6.3.2 Individual Spacecraft Manoeuvres Selection Trades

As seen in the previous sub-section, de-coupling some of the selection processes can decrease the complexity of other selection processes without a significant loss of accuracy. This type of de-coupling can also be employed in the selection of the individual spacecraft manoeuvres.

The trajectories that the spacecraft follow depend upon the selection of the final absolute positions, the thrusters parameters used and the local gravitational model. For DARWIN the formation will be subject to the gravitational environment about the L_2 point in the Sun/Earth system. Consider the n-body problem defined in Equation(6.1), $m_i, m_j, \ldots, m_n$ describe the masses of *n* point masses with position vectors $r_i, r_j, \ldots, r_n$.

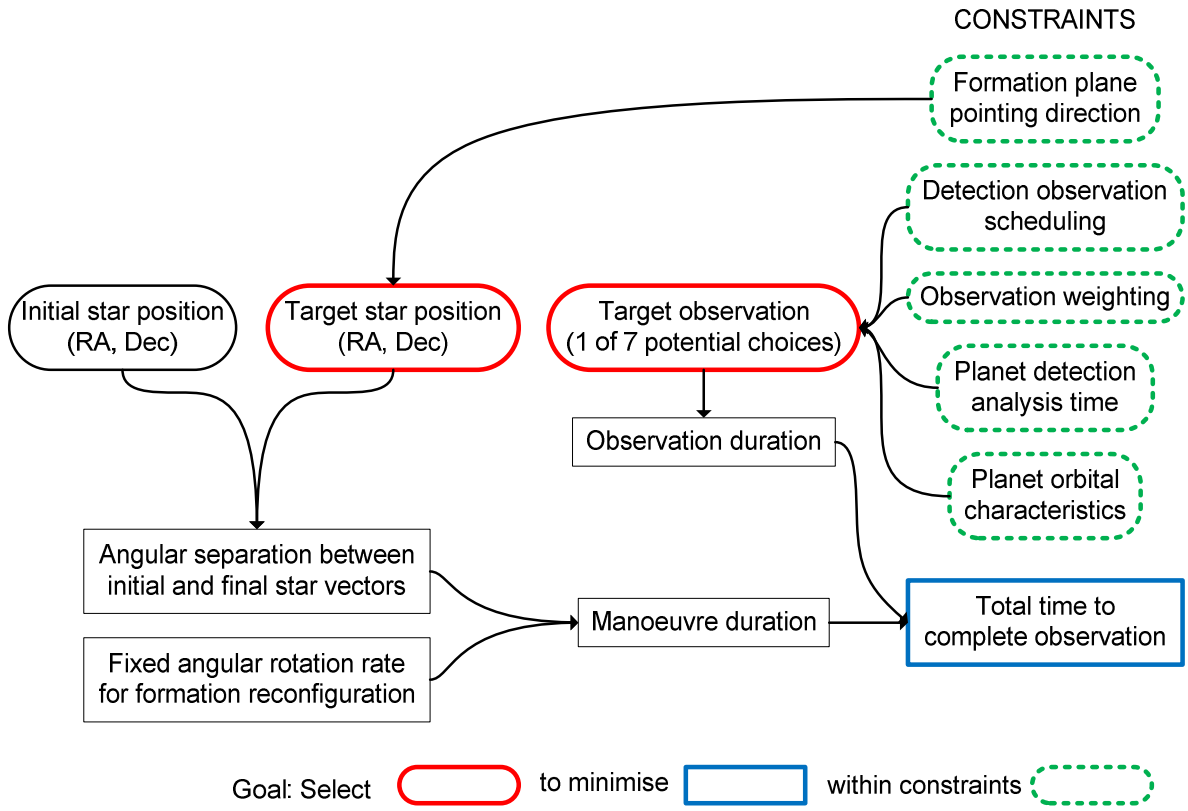

Figure 6-4 Simplified target star/observation selection model

$$m_i \ddot{\mathrm{r}}_i = G \sum_{i \neq j} \frac{m_i m_j \left(\mathrm{r}_j - \mathrm{r}_i \right)}{\left| \mathrm{r}_j - \mathrm{r}_i \right|^3}, \quad i = 1, \ldots, n \tag{6.1}$$

In the three-body system of the Sun, Earth and spacecraft the force exerted on the spacecraft by the Sun and the Earth can be calculated. Using the values found in this force is found to be ~6N (putting the spacecraft 30,000 km 'above' L_2 in the Z-direction). Using the same 100m reconfiguration manoeuvre calculated in sub-section 6.3.1, this acceleration equates to a perturbed distance of ~196 km. This is a significant perturbation that accentuates the importance of the gravitational model. However, in essence, it indicates how far the local gravitational environment affects the entire formation, not the relative distances between the spacecraft in the formation. Placing a second spacecraft 100 m from the first in the positive Z-direction gives a distance perturbation of only ~0.2 mm difference to the first spacecraft's perturbation. Thus, over a 100 m manoeuvre, a pair of spacecraft separated by 100 m only drifts apart by ~0.2 mm. This illustrates how, over short timescales of a few hours, the local gravitational environment has very little affect on the relative distances of the spacecraft in the DARWIN formation.

Table 6-2 Values used to calculate perturbing acceleration on a spacecraft around L_2 in the Sun/Earth system

Property	Value
Mass of the Sun m_S	1.9891×10^{30} kg
Mass of the Earth m_E	5.9736×10^{24} kg
Mass of the spacecraft m_{sp}	1000 kg
Position of the Sun $\mathbf{r_S}$	(0, 0, 0) km
Position of the Earth $\mathbf{r_E}$	$(1.496\times10^{11}, 0, 0)$ m
Position of the spacecraft $\mathbf{r_{sp}}$	$(1.511\times10^{11}, 0, 3.0\times10^{8})$ km
Gravitational constant G	6.67428×10^{-11} $m^3kg^{-1}s^{-2}$

The analysis above has shown how important the local gravitational environment of L_2 is when trying to select the absolute final spacecraft positions, but how little affect it has on the relative positions of the spacecraft within the formation. Therefore we can ignore the local gravitation environment for manoeuvre planning between the initial and final relative positions. The spacecraft trajectories therefore become a function of the natural trajectory through the local gravitational environment and the controlled trajectory of the reconfiguration manoeuvre to change the relative positioning of the spacecraft. As they are independent of each other they can be de-coupled for manoeuvre planning purposes. This de-coupling creates two separate selection processes (derived from the individual spacecraft manoeuvre selection process, Figure 6-3) given by Figure 6-5 and Figure 6-6.

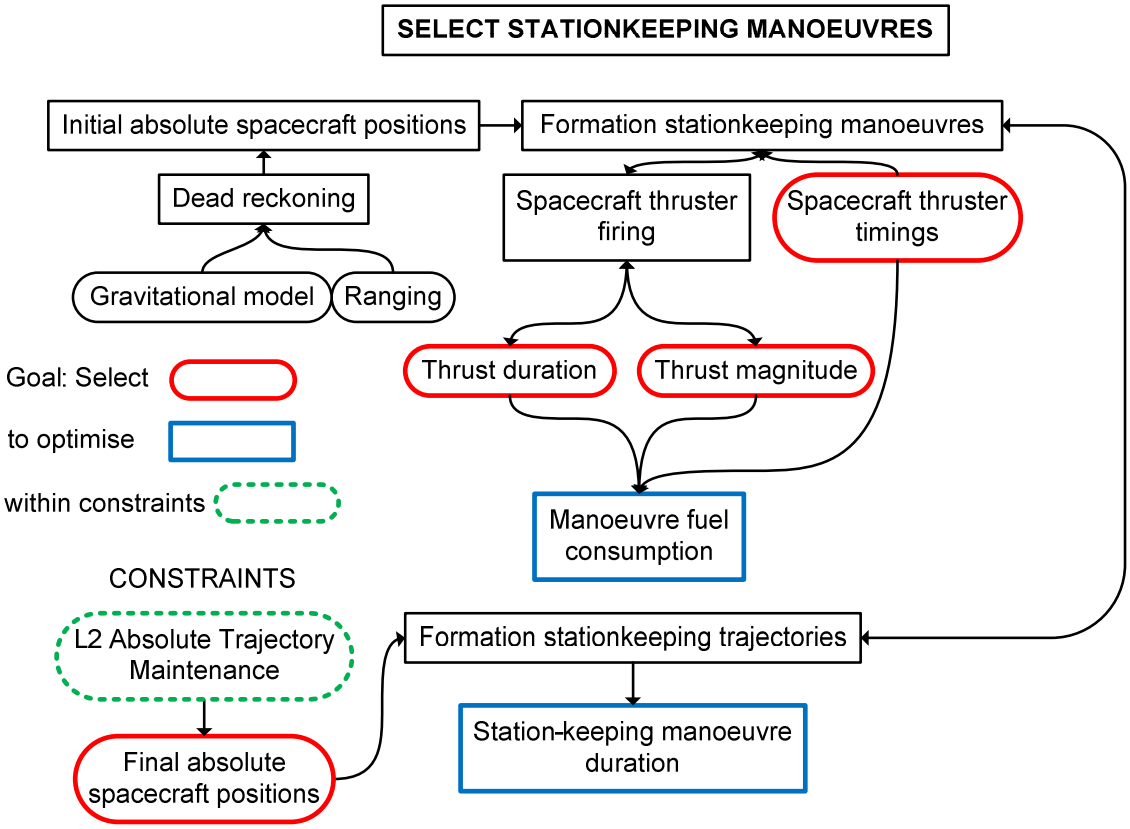

Figure 6-5 Select stationkeeping manoeuvres model

The stationkeeping manoeuvre selection model is shown in Figure 6-5. Stationkeeping manoeuvres are performed by the formation (i.e. spacecraft relative positions are maintained during the manoeuvre) with the goal to maintain the formation on a trajectory that keeps it within the vicinity of the L_2 point. The selection of the final absolute spacecraft positions addresses this goal. In addition the fuel consumption and manoeuvre duration should be optimised by the selection of appropriate thruster firing parameters. With similar sized spacecraft, there is no fuel balancing issue since all the spacecraft will consume the same amount of fuel during the stationkeeping manoeuvre. Maintenance of the spacecraft relative position during the manoeuvre also negates the need to collision avoidance and plume impingement constraints.

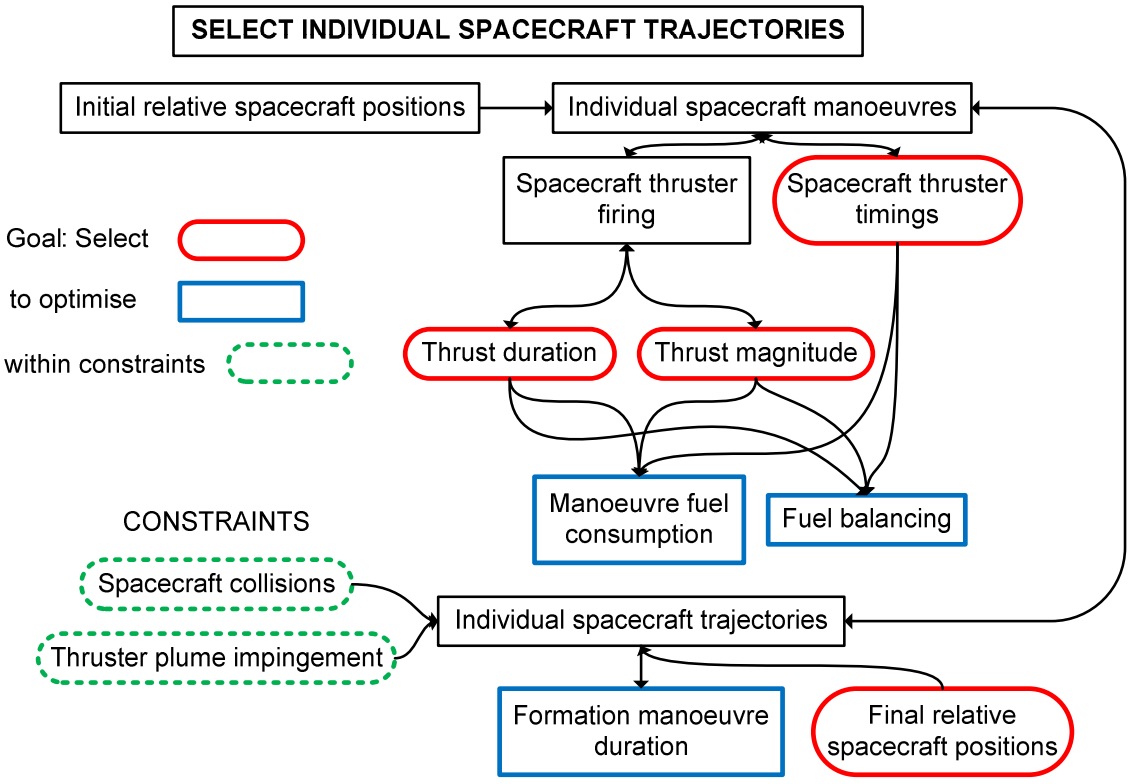

Figure 6-6 Selection individual spacecraft trajectories model

The individual spacecraft trajectory selection model is show in Figure 6-6. Here the local gravitational model has been completely de-coupled and so a free-space gravitational model is used in its place. Though the relative spacecraft positions to perform the observation would be provided by the target star/observation selection, there is additional scope for relative positioning and so it remains a selection parameter. Position selection can be seen as a two stage process. First the positions are selected then potential trajectories are examined. This position/trajectory process is then iterated until a combination is found that suits the optimisation goals. The most optimal trajectory between any two points in free-space is a straight line. Analysis in section 9.3.4.1 shows that at least 56 % of all relative position combinations for the DARWIN mission conform to the avoidance criteria using straight-line

trajectories. For these relative positions, much calculation time would be wasted searching for a more optimal trajectory using the optimisation model described above. There are ways however that the position/trajectory optimisation process can be enhanced.

There are many ways in which the position/trajectory selection could be performed. Figure 6-7 shows three ways this could be achieved and details the computational costs involved. A position calculation (i.e. calculating fuel cost, balancing and manoeuvre time for a set of positions a straight-line trajectory generated by a bang-bang thrust profile) is assumed to take 1 computational unit (CU)[3]. A trajectory calculation (i.e. calculating fuel cost, balancing and manoeuvre duration for the same set of positions using a non-straight trajectory generated by two perpendicular bang-bang thrust profiles) is assumed to take 2CU and a trajectory check (i.e. assessing all trajectories for collisions and plume impingement), 5CU. Finally is it assumed that a position optimisation takes 10 position calculations and a trajectory optimisation takes 10 trajectory calculations/checks. In Figure 6-7, Process1 describes the position/trajectory selection process described above. For Process2, after every position check, the straight-line trajectory is checked and if it passes the avoidance criteria then the next position calculation is performed. If it fails the avoidance criteria a trajectory optimisation is performed. For Process3 a complete position optimisation is performed, the straight-line trajectory checked and if it fails the avoidance criteria a trajectory optimisation is performed.

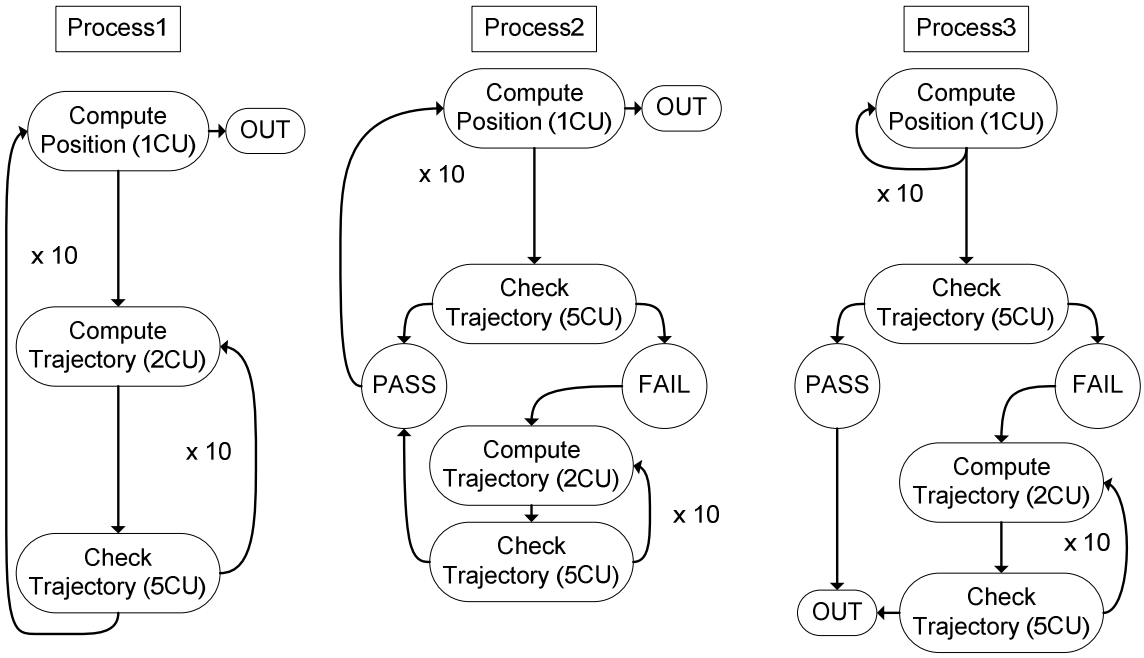

Figure 6-7 Example position/trajectory optimisation processes

Table 6-3 shows the computational load examples for the three processes given in Figure 6-7. Loads are calculated for one position optimisation and for 100 position optimisations. The pass rate for straight-line trajectories conforming to the avoidance criteria is 60%. Process1 is the most computationally intensive mainly so because it doesn't perform a

[3] An arbitrary unit of CPU time used to compare relative the relative performance of differing algorithms

trajectory check on the straight-line trajectory from the position calculation. When this is done, as in Process2, the computational load drops to 45 % of Process1. Process3 performs the best because no trajectory checking is performed within the compute position loop. For a 1 position optimisation the computational load drops to 2% if the straight-line trajectory passes and 12% if it fails. Over a 100 position optimisation the computation load is just 6% of the Process1 and 13% of the Process2 computational loads. In terms of the computational load, Process3 is clearly the better option.

Table 6-3 Computational load examples for different position/trajectory optimisation processes

Process	Computation load for 1 position optimisation (CU)		Computation load for 100 position optimisations (CU)
1	710		71000
2	320		32000
3	Pass 15	Fail 85	4300

The analysis of the select individual spacecraft trajectories model shows that optimisation Process3 is the most efficient. This allows for de-coupling between the relative position selection and the thruster firing parameter selection elements in this model. The position selection model is shown in Figure 6-8. Here only the selection of the relative spacecraft positions important and no constraints are added. The manoeuvres are governed by straight-line trajectories with fixed thrusters firing parameters, thus the optimisation goals depend wholly on the selection of the relative spacecraft positions. The avoidance trajectory selection model is given in Figure 6-9. This selection is only required if the straight-line trajectories form the position selection model fail the avoidance criteria. Here the final relative spacecraft positions are fixed, however fuel optimisation can still be achieved through selection of the thrusters firing parameters. These also ensure the resulting trajectory conforms to the avoidance criteria.

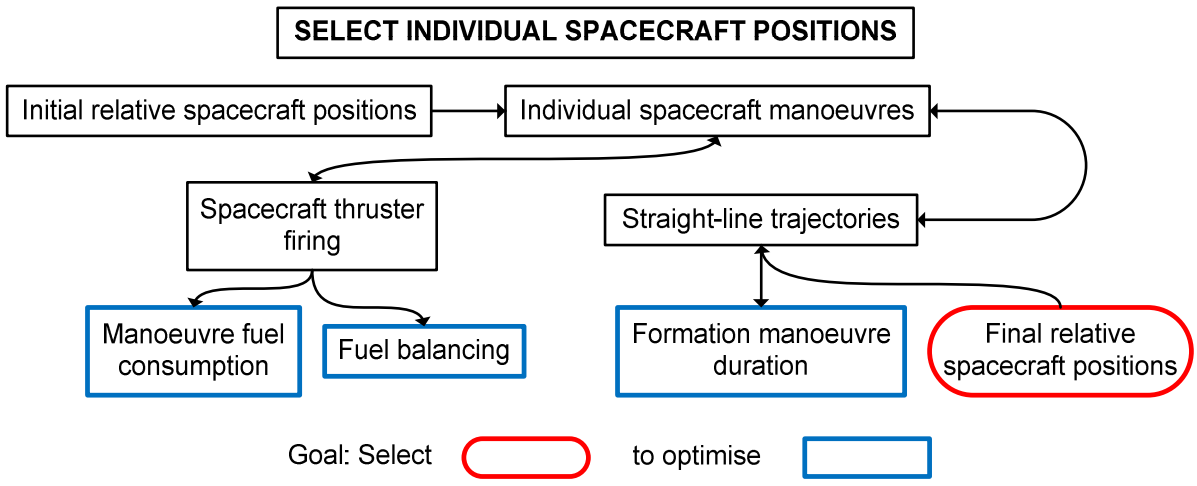

Figure 6-8 Individual spacecraft positions selection model

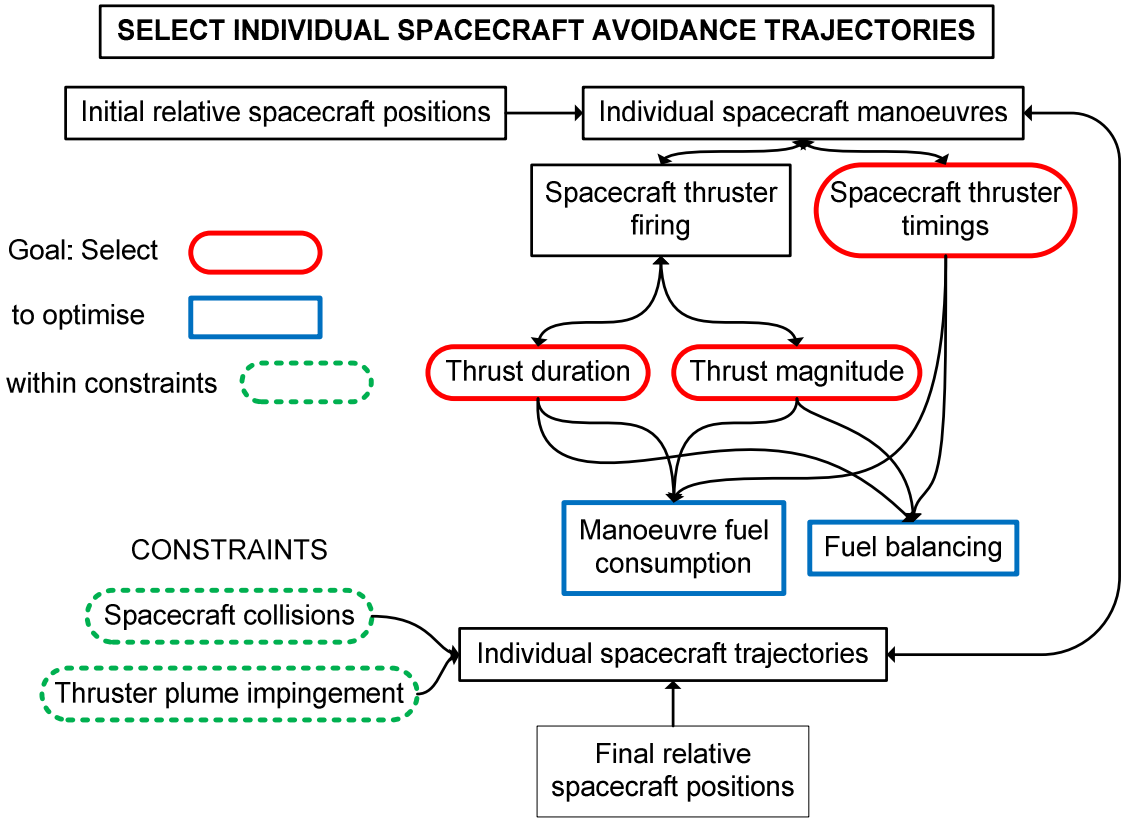

Figure 6-9 Individual spacecraft avoidance trajectories selection model

The manoeuvre planning systems trades have help to reduce a very complex planning problem defined by Figure 6-1, Figure 6-2 and Figure 6-3 into a significantly less complex form (Figure 6-4, Figure 6-5, Figure 6-8 and Figure 6-9) without any apparent loss of flexibility, functionality, efficiency and optimality. It must be remembered however that these simplifications may produce unforeseen problems and so any further analysis must take into consideration the system trades made.

6.4 Manoeuvre Planning Architecture

The core of this research project is the design of a Manoeuvre Planning Architecture (MPA) for the optimisation of spacecraft formation flying reconfiguration manoeuvres. The envisaged design, the Separate Modular Manoeuvre Planning Architecture (SepM-MPA), is introduced in the next sub-section and is heavily influenced by the trade-off processes described in Chapter 5 and sub-section 6.3.

6.4.1 Optimisation Modules

The system trades performed have naturally led to a modularised approach for the Manoeuvre Planning Architecture. Each of the independent selection process models can be characterised by a distinct task, optimisation goals, constraints and input/output values. These tasks are defined as follows and optimisation module summaries are given in Table 6-4.

Table 6-4 Optimisation module summaries

Module	Optimisation Goals	Constraints	Inputs	Outputs
Science Operations Module	Minimise total time to complete observation	formation pointing, detection task scheduling, planet detection analysis time, planet orbital characteristics, task weighting	initial formation pointing direction, direction of anti-sun vector	final pointing direction, final formation configuration
Stationkeeping Module	Minimise fuel consumption and manoeuvre duration	maintain formation within limits of reference trajectory	initial absolute spacecraft positions	thrust direction, duration and timing for stationkeeping manoeuvre
Position Assignment Module	Minimise fuel consumption and manoeuvre duration and control fuel balancing	none	initial absolute spacecraft positions, final pointing direction, final formation configuration	final relative spacecraft positions, manoeuvre duration
Trajectory Design Module	Minimise fuel consumption and manoeuvre duration and control fuel balancing	spacecraft collisions, thruster plume impingement	initial and final relative spacecraft positions, manoeuvre duration	thrust direction, duration and timings for the reconfiguration manoeuvre

- Science Operations Module (SOM) [Figure 6-4] – this module decides what the next science task (i.e. target star/observation) will be after the current science task has been completed
- Station-Keeping Module (SKM) [Figure 6-5] – this module controls the planning of the station-keeping manoeuvre to ensure the formation remains within the vicinity of its reference trajectory
- Position Assignment Module (PAM) [Figure 6-8] – this module optimises the relative final spacecraft positions that, at the end of the manoeuvre, satisfy the science task formation configuration requirements
- Trajectory Design Module (TDM) [Figure 6-9] – this module optimises the spacecraft trajectories and ensures there are no collisions or thruster plume impingement issues during the manoeuvre

These four modules are described in greater depth in Chapters 7, 8, 9 and 10.

6.4.2 Manoeuvre Planning Architecture Design

Within the SepM-MPA each of the optimisation modules act as individual entities providing only feed-forward information to the next module. A flow diagram of the SepM-MPA can be found in Figure 6-10 and is highly linear in nature. The architecture uses an on-board mission catalogue that contains details of the co-ordinates of all the target stars for the mission and baselines and observation duration times of every science task for each star. The catalogue, coupled with the current formation configuration state, forms the inputs to the SepM-MPA.

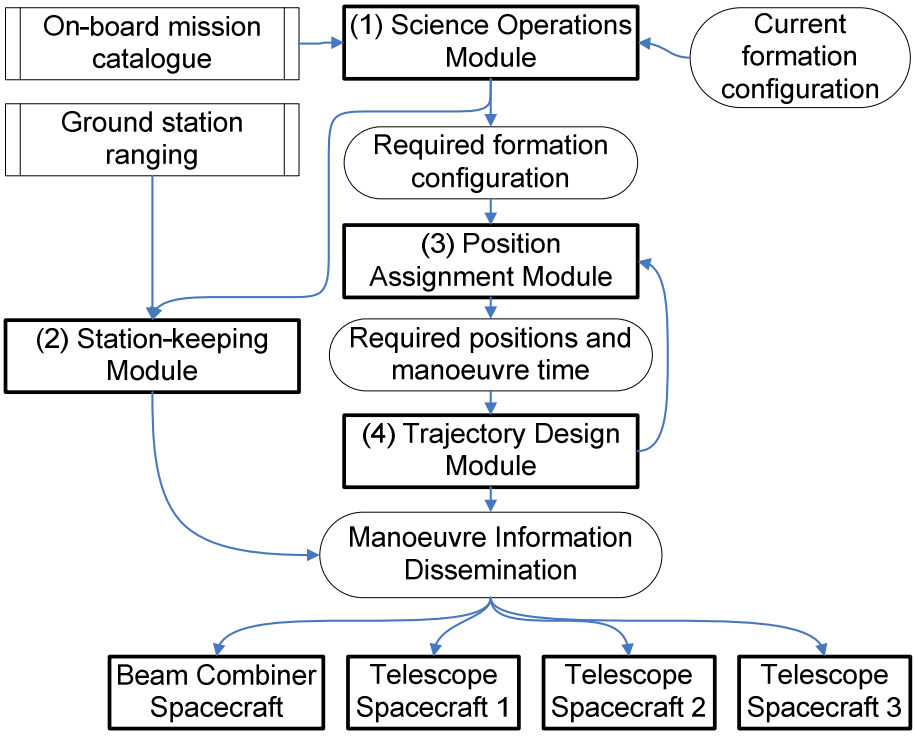

Figure 6-10 Separate modular manoeuvre planning architecture (SepM-MPA)

The Science Operations Module (SOM) decides what the next science task should be based on the availability of science tasks once the current task has completed. The SOM uses an optimisation algorithm to select the task minimises the total time to complete the observation. A more in-depth description of the SOM and an analysis of its operation can be found in Chapter 7. The task configuration details for the selected task (e.g. pointing direction, baseline and formation configuration), along with the current configuration details, are passed to the Position Assignment Module (PAM).

The PAM uses the current and desired formation configuration information to find the post-manoeuvre spacecraft positions that satisfy the science requirements whilst optimising the manoeuvre duration, fuel consumption and the fuel balancing across the fleet. As these three optimisation goals are generally mutually-exclusive a decision algorithm is employed to select the levels to which these goals are to be achieved. This decision is based on the current fuel levels for the formation and prior analysis of how the PAM optimisation operates. Further information relating to the PAM, its operation and its performance analysis can be

found in Chapter 8. The current positions of the spacecraft relative to the formation centre and the optimised relative positions, along with the manoeuvre duration, are passed to the Trajectory Design Module (TDM).

Using the output from the PAM, the TDM first checks to see whether any of the spacecraft proximity or thruster plume proximity constraints have been violated in the PAM optimised manoeuvre. If there are no proximity violations then no trajectory modification is required and so the manoeuvre thruster firing parameters are passed for manoeuvre information dissemination. If a proximity violation has occurred however, the TDM modifies the spacecraft trajectories to ensure the proximity constraints are adhered to and the resulting modified trajectory is optimised for manoeuvre duration, fuel consumption and fuel balancing across the fleet. As for the PAM, a decision algorithm is required to select appropriate levels to which these optimisation goals are achieved. This decision is based on the current fuel levels for the formation and prior analysis of how the TDM optimisation operates. Further information relating to the TDM, its operation and its performance analysis can be found in Chapter 9. For the modified trajectories the TDM passes the revised manoeuvre thruster firing parameters for manoeuvre information dissemination.

The manoeuvre planning performed by the PAM and the TDM does not take into account the dynamics of spacecraft motion around L_2 or the necessity to remain in the vicinity of the reference trajectory. The SKM uses ground station ranging data and trajectory propagation calculations to estimate the position of the BCS relative to the reference trajectory. The SKM then calculates the size and duration of the station-keeping manoeuvre required to return the formation to the reference trajectory over a pre-defined time period. If the upcoming manoeuvre phase is identified as an optimal time to perform a station-keeping manoeuvre then the station-keeping manoeuvre thruster firing parameters are passed for manoeuvre information dissemination. If no station-keeping manoeuvre is required then there is no output from the SKM. A more detailed introduction to the SKM can be found in Chapter 9.

The manoeuvre is executed during the manoeuvre information dissemination stage. The manoeuvre information is passed in two stages. First the station-keeping manoeuvre data is sent to all the spacecraft in the formation. For the BCS the data used is open-loop as there is no feedback available. For the TSs, the data can be used in the estimation section of the controllers. For the station-keeping manoeuvres the TSs will simply formation-keep with the BCS, however the SK manoeuvre data (i.e. data showing when the BCS will perform manoeuvres and how large they will be) can aid the TS closed-loop controllers to avoid over-compensation when the BCS thruster firings occur. After the station-keeping manoeuvre has been completed the reconfiguration manoeuvre data is passed to each spacecraft. For the BCS this is simply open-loop thruster switch times. For the TSs the optimised trajectory computed by the algorithm is passed as a trajectory relative to the BCS. The TS closed-loop controllers therefore formation-keep with the BCS and change the spacecraft's relative position to the BCS following the optimised trajectory data. This results in a formation reconfiguration manoeuvre that is safe and optimal in inertial space.

6.4.3 Manoeuvre Planning Architecture Limitations

The SepM-MPA describes the fastest method for implementing a complete manoeuvre planning strategy that encompasses everything from science task assignment to reconfiguration manoeuvre optimisation and station-keeping manoeuvre optimisation. The relative simplicity of the SepM-MPA to reduce computational complexity however does come at the price of reduced accuracy in the calculations due to a number of assumptions and omissions made during the manoeuvre planning system trades process:

- Manoeuvre duration in the SOM is a function of the angular separation of the initial and target star vectors. This will not be the same as the manoeuvre duration calculated by the PAM.
- The optimal position found by the PAM may become sub-optimal if TDM trajectory modification is required.
- Planning of manoeuvres in free-space will produce a different outcome to executing those manoeuvres in the L_2 dynamic environment.
- Separate stationkeeping and reconfiguration manoeuvres will take longer to execute than combining them.
- Reconfiguration manoeuvres may have an effect on the size and timing of future stationkeeping manoeuvres.

Addressing these accuracy issues within the SepM-MPA will invariably require the use of more computing power but also a change in the manoeuvre planning architecture.

6.5 Manoeuvre Planning Architecture Execution Options

There are two approaches that can be employed in executing the Manoeuvre Planning Architecture; a global approach and a local approach. For the local approach the planning is performed exclusively to optimise the upcoming manoeuvre and no fore-thought is given to the optimisation of future manoeuvres. In the global approach the planning is performed over a mission sub-set (either a set duration or a set number of manoeuvres) and each individual manoeuvre within the sub-set is planned to achieve optimisation over the entire mission sub-set. A comparison of these two approaches follows:

- manoeuvre planning architecture complexity – the local approach to manoeuvre planning is relatively simple compared to the global sub-set since the latter involves the planning of multiple manoeuvres whilst the former plans only one
- manoeuvre planning speed – the local approach will plan much quicker that the global sub-set for the same reasons as above however over the mission sub-set the global approach may be quicker than the sum of the individual local manoeuvre plans over the same sub-set
- changing mission environment – over a mission sub-set the local approach can adapt to changes within the mission environment on a manoeuvre-by-manoeuvre basis

whilst the global sub-set approach needs to incorporate a fixed mission environment over the same planning sub-set

- planning optimality – over a mission sub-set a combination of manoeuvres planned using the local approach is unlikely to be more optimal than the combination planned using the global sub-set approach

The modular nature of the SepM-MPA allows for these two execution approaches to be analysed in two separate domains: over the entire MPA and over the individual optimisation modules.

6.5.1 Executing the Manoeuvre Planning Architecture

The greatest thing to affect the accuracy and optimality of a manoeuvre plan is the difference between the planning environment used to make the plan and the actual mission environment during the execution of the plan. For DARWIN there are many potential differences including:

- changing availability of observation tasks: the DARWIN science schedule will be very dynamic due to a number of factor that cannot be planned for e.g.
    - detecting a planet allows 6 more observations that can be scheduled – though they cannot be scheduled before the planet has been detected
    - special tasking directed by the science team
- changes in spacecraft fuel amounts and balancing due to
    - formation-keeping manoeuvres
    - attitude manoeuvres
    - calibration manoeuvres
- changes in the entire mission environment due to safe-mode operations or other malfunction

Using a global execution approach to optimise a sub-set of manoeuvres allow the changes between the planned and the actual mission environment to increase over the timescale of the plan. Reducing the size of the sub-set will reduce these differences but selection an appropriate sub-set size that balances optimisation gains against environment accuracy is not possible to achieve until an architecture that optimises locally has been developed. Therefore the SepM-MPA will be executed locally, on a manoeuvre-by-manoeuvre basis.

6.5.2 Executing the Individual Optimisation Modules

Though the architecture is to be executed locally on a manoeuvre-by-manoeuvre basis the individual optimisation modules within the architecture can be executed globally. This potentially offers the gains achievable with global optimisation without the losses of using a static planning environment over an extended mission sub-set. This is because the global plan for each individual model is revised on a manoeuvre-by-manoeuvre basis and thus the most

up-to-date planning environment is always employed. The cost for adopting this execution method is an increase in the planning time required and so some optimisation modules may not benefit from this type of execution approach.

For the SOM the differences between potential 'total time to complete the observation' could be in the order of days to tens of days in the worst case scenario (see Table 6-1). Selection of which observation task to perform does have far-reaching effects on the availability of tasks due to the constraints listed in sub-section 6.2.1. The loss of tens of days of potential observation time due to a poor selection of locally optimal tasks is clearly unacceptable and so the SOM would definitely benefit from taking a global approach to its execution. This can be implemented by having the SOM find an optimal tour of observations that maximise the mission science returns over a mission sub-set. The first observation within that tour then becomes the selection observation for that manoeuvre phase.

For stationkeeping at L_2 previous missions to L_1 (Farquhar et al, 1980) and analysis of missions to L_2 (Dunham and Roberts, 2001; Rohrbaugh and Schiff, 2002 and Williams et al, 2000) show that the required frequency of stationkeeping manoeuvre to maintain loose trajectories around collinear libration points is one manoeuvre every 3-6 months. The frequency is far smaller than the reconfiguration manoeuvre frequency for the DARWIN mission (on average one every few days – see Table 2-2). Although the execution of stationkeeping manoeuvres will be in the order of one every few months it will still be necessary to monitor the formation's position around L_2 in order to optimise the timings and fuel consumption of those manoeuvres. It is prudent therefore for the stationkeeping module to plan a manoeuvre for every reconfiguration phase but only execute the manoeuvre if it falls within clearly defined execution parameters.

For the manoeuvre planning execution the PAM and the TDM need to be considered together. If planned locally then they will operate the same as Process1 defined in sub-section 6.3.2 and Figure 6.7. For global planning, the observation tour calculated by the SOM would be used to generate a tour of manoeuvres. This tour would then be optimised by modifying each individual manoeuvre within the tour until a global solution is found. These are several problems with this approach. First, it could be very computationally intensive. A ten manoeuvre tour would take 430 CU so optimising this takes the same amount of CU for every tour calculated (which could be tens to hundreds). Secondly, the tour does not take into account the fuel differences caused by formation-keeping or attitude manoeuvres or potential stationkeeping manoeuvres (since the SKM is de-coupled from the PAM/TDM), thus the longer the tour the larger the error that will be induced in the optimisation. For these reason a local implementation of the PAM/TDM will be performed for the SepM-MPA.

6.6 Hardware Requirements

For the SepM-MPA to operate autonomously on-board one of the DARWIN spacecraft it needs to be capable of performing within the limits of the processing power available. Figure 6-11 illustrates how the speed of spacecraft CPUs has increased over the last two decades. With an estimated launch date for DARWIN circa 2030, it will likely have on-board a processor from the early 2020s.

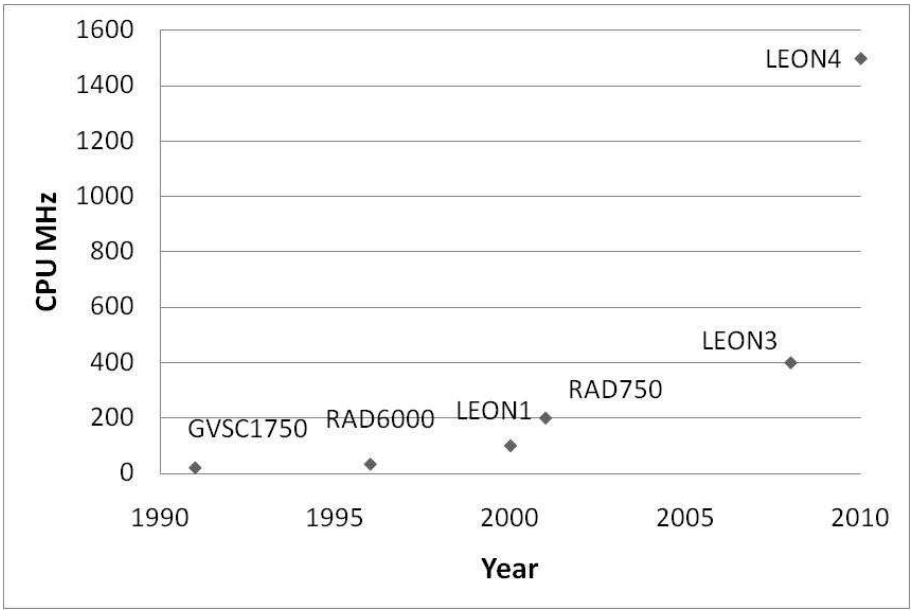

Figure 6-11 MHz increase of popular spacecraft CPUs since the 1990s – GVSC1750, RAD6000 and RAD750 (BAE Systems, 2002), LEON1 (Gaisler, 2000), LEON3 (Aeroflex Gaisler, 2008), LEON4 (Aeroflex Gaisler, 2010)

The analysis of the SepM-MPA was performed within the MATLAB 2006a software environment running in Windows XP Professional on a 3.06 GHz Intel Pentium 4 PC with 1 GB of RAM. Extrapolating the data from Figure 6-11 it is reasonable to assume that DARWIN will have at least this processing capability by launch.

6.7 Chapter Summary

The planning of safe and optimised formation flying reconfiguration manoeuvres requires a software architecture that balances an accurate representation of the problem to be solved with the computation requirements to achieve a viable plan. In this chapter, the manoeuvre planning systems model for DARWIN is presented and a trade-off analysis is performed to reduce its complexity without the loss of functionality. This trade-off resulted in the definition of four independent optimisation modules that aim to address the manoeuvre planning issues of target selection, fuel consumption and manoeuvre duration optimisation, fuel balancing, collision and thruster plume avoidance and optimal station-keeping around L_2. These modules were then implemented within the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) providing a fast method for realising a complete manoeuvre planning strategy that satisfies the issues presented down in the problem statement.

7. SCIENCE OPERATIONS MODULE

From the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) the goal of the Science Operations Module (SOM) is to find an optimal science operations schedule that maximises the mission science returns in the scheduled time. For this design this translates to maximising the number of completed observations within the scheduled time. There are almost 450 stars in the DARWIN star catalogue leading to over 4000 potential science observation tasks ranging in duration from ⅓ day to ~85 days (Karlsson, et al., 2004). The DARWIN mission is only being planned for 5 years. It is imperative that the available time is not wasted by unnecessary and lengthy manoeuvres. Careful planning of the science operations schedule can help to ensure this. *A priori* mission planning however cannot be performed as there are no indications as to which stars have planets. The science operations planning must therefore be performed in real-time in parallel with the mission. This allows for the planned schedule to be adapted when updated information is available (i.e. the existence of a planet has been confirmed) and be receptive to special tasking by the science team (i.e. the science team may want to image a particular planetary system).

In deciding the choice of the next observation to perform the most logical choice is to choose the observation that takes the least time to complete (if the optimisation goal is to maximise the number of observations completed). Selecting the shortest task from a selection of all available tasks is the locally optimal solution. However, as discussed in sub-section 6.5.2, a global solution is much more desirable. To achieve this global optimisation it may be more efficient to choose stars further away (in angular separation) but with shorter observation durations than those closest but with longer observation durations. This also implies it may be better to remain at the current star and perform a different (but longer) observation than to move to a different star but perform a shorter observation. This is called the Science Task Assignment (STA) problem for the duration of this chapter. Solving the STA problem is the goal of the SOM as the solution will provide the next manoeuvre goal for the PAM and an optimal schedule for analysis by the SKM.

7.1 Previous Contributions

Planning and scheduling for autonomous spacecraft is a relatively new concept that has been put into practice in various forms on a number of space missions. These include Cassini (Muscettola et al., 1995), Rosetta (Ferri and Sorensen, 1998) and Deep Space One (Smith, Rajan and Muscettola, 1997) to name but a few. Chein, et al. (1998) provides a useful introduction to the planning systems used at NASA and specifically for Deep Space One and Earth Orbiter One from the New Millennium Program. They describe the benefits of ground-based and on-board operations planners as “reduced costs, increased responsiveness, increased interactivity, increased productivity and simplified self-monitoring”. The planners use “symbolic AI” routines to review any given plan, identify the flaws and iteratively remove the flaws to create a flawless plan. These papers however describe operations planners that take into account all of the subsystems and payloads on a spacecraft and find a schedule that allows the most efficient use of spacecraft resources.

More recently an autonomous mission planning system was introduced by Rui, et al. (2003) and Rui, Ping-yuan and Xiao-fei (2005) called the Multi-Agent Planning System (MAPS). It is designed for a fully autonomous deep-space mission and combines the results from a number of planning agents (representing the spacecraft subsystems) with a planning manager agent to create mission operations schedules for the complete spacecraft. This is implemented in C++ and Java respectively but the authors only provide the architecture for the model, not the details of how the planner actually achieves its goals. The papers mentioned above address the problem of automating spacecraft operations; however they do not address optimisation within the planning environment. The tasks are scheduled within the time and resource allocation constraints defined but no attempt is made to improve a plan once an achievable schedule has been found. This is different to the situation of the STA problem. Since there will always be many more tasks available than can be completed in any defined time period the SOM must optimise the schedule to increase the number of observations completed.

Bailey, McLain and Beard (2001) have addressed a similar mission planning problem to that posed by the SOM. The authors use NASA's Starlight dual-spacecraft interferometry mission (Blackwood et al, 2003) as a baseline for examining interferometry mission schedules that reduce fuel expenditure and encourage fuel balancing between the spacecraft. Fuel consumption dynamics for each manoeuvre are based on a scheme by Beard and Hadaegh (1999) and manoeuvres are separated into three types: retarget, resize and reorientation. Each star, separated by retarget manoeuvres, is imaged using a number of resize and reorientation manoeuvres. The optimisation of the manoeuvre schedule is likened to a travelling salesman problem (TSP). A "Chained Local Optimisation (CLO)" algorithm, combining simulated annealing and local search optimisation techniques, is developed to solve the TSP. The authors use a benchmark tour as a comparison that involves completing all the resize and reorientation manoeuvres for each star before retargeting to the next star. The results show that, for a number of different scenarios, the optimized tour is capable of reducing the fuel consumption of the spacecraft performing all manoeuvres by a significant amount compared to the benchmark tour. This is achieved by combining resize, reorientation and retarget manoeuvres instead of performing all the resize and reorientation manoeuvres on one target before moving to the next. These results are noteworthy in understanding the necessity for the SOM as an analogy can be made. In the paper, each star has a number of resize and reorientation manoeuvres each with their associated fuel costs. Likewise, the stars in the DARWIN catalogue each have a number of observation tasks with their associated time costs. By extension it may be inferred that combining the science tasks for DARWIN in an appropriate schedule is more time optimal than performing the observations in a systematic fashion. There are a number of factors however that Bailey, McLain and Beard (2001) do not address. The whole problem is time invariant and there are no pointing constraints. This means that a star's availability for observation is not included and the duration of each manoeuvre is omitted (since fuel consumption is the only factor being investigated). Furthermore, a fixed number of manoeuvres are investigated and each solution finds a tour that encompasses them all. As previously discussed, the SOM will have to find a tour with temporal constraints from a large number of possible tasks that will not all be able to be completed within the given time. Therefore, although Bailey, McLain and Beard (2001)

reveal some interesting trends with respect to interferometry mission scheduling, the relative simplicity of their method prevents its use to emulate the SOM.

7.2 Science Task Assignment Analogies in Graph Theory

Graph theory is the mathematical study of graphs. The graphs in this theory are made up of a number of ‘nodes’ and each pair of nodes in the graph are connected by ‘edges’. The STA problem can be described as a graph where the science tasks are the nodes and the time to complete to task (from any starting task) defines the weight of the edge connecting the two tasks. This is illustrated in Figure 7-1 where the circles represent nodes and the lines represent the edges. For the STA problem not all the nodes are connected to each other due to the various constraints found within the problem. Finding a path through the graph from one node to another whilst minimising the sum of the weights of the edges is called a shortest path problem and there are many analogies to the STA that can be found in these problems.

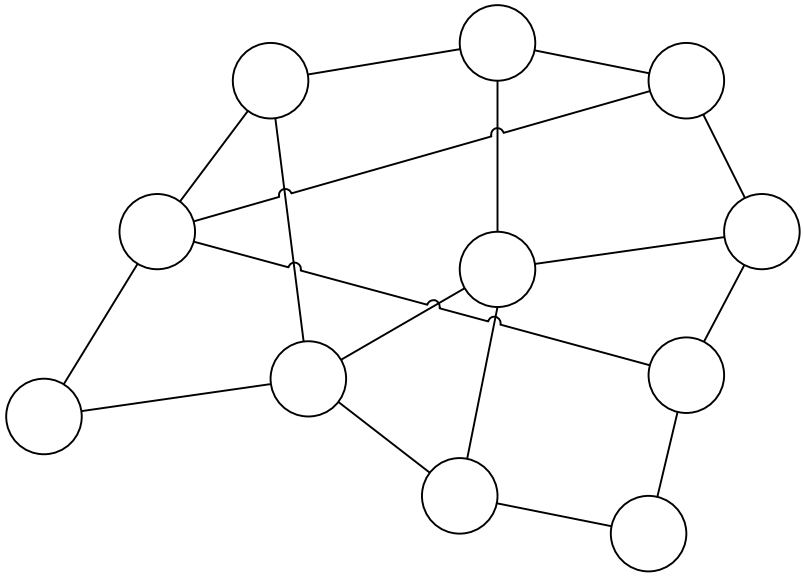

Figure 7-1 An example graph where the nodes are represented by circles and the edges connecting them by lines.

The STA problem can be likened to a travelling salesman problem (TSP). In the basic TSP the salesman has to visit a number of cities and the tour is optimal if the shortest route is chosen. The TSP and a variant, the Vehicle Routing Problem (VRP), are cornerstone challenges within the field of Operations Research. Similar problems can also be found in computer network routing literature. There are many different kinds of constraint that can be added to the TSP including time windows (visiting a point within a specified time period) (Bansal et al, 2004), path constraints (minimising the number of times a separate point is crossed, e.g. a river) (Press et al. Section 10.9, 1992) or sequential constraints (visiting a point before another) (Hernádvölgyi, 2003). These constraints can all be made applicable to the STA problem. There are also a vast number of different methods that can be employed to solve these problems. Exact algorithms, like branch and bound, cutting plane and linear programming are good for finding solutions with a relatively small number of points to visit.

For larger problems, heuristics have been developed that quickly find good solutions (>97% optimal). These include nearest neighbour (Hurkens and Woeginger, 2004), simulated annealing (Press et al. Section 10.9, 1992), genetic algorithms (Merz and Freisleben, 1997) and ant algorithms (Dorigo, Di Caro and Gambardella, 1999). The problem with the TSP in relation to the STA problem is that the TSP aims to find a path connecting a fixed number of cities but for the STA the number of total tasks is dynamic and will decrease throughout the mission as planets are not found around target stars.

Another close analogy to the STA problem is the single machine scheduling problem (SMSP) found in manufacturing. In this problem there is a list of *n* jobs (science tasks). Each job has a release time (the time the task enters the FOV), a duration (the observation time of the task), a deadline (the time the star leaves the FOV) and a weight (the importance of completing the task). The goal of any algorithm to solve this problem is to minimise the total number of late jobs (i.e. those completed after their deadline). Again a number of conventional and innovative methods have been employed to solve this problem including branch and bound (Brucker, Hilbig and Hurink, 1999), local search (Crauwels, Potts and Van Wassenhove, 1998) and memetic algorithms (Franca, Mendes and Moscato, 2001). This analogy however still fails to emulate the STA problem as it assumes a fixed number of tasks and no task can be selected if it over-runs its deadline (i.e. the star leaving the FOV).

There are graph theory analogies similar to the STA problem that deal with the situation where the graph changes after each step has been made. The Canadian Traveller Problem (Papadimitriou and Yannakakis, 1989) describes the a graph where at the beginning of the path the entire graph and edge costs are know but at each step some of the edges fail and cannot be used. In this problem a ‘policy’ is defined that can be used to determine edge selection at each node. The optimal ‘policy’ is then calculated using dynamic programming. The Canadian Traveller Problem (CTP) (and its variants) however only deal with paths between known nodes. For the STA problem, the goal is not to reach a chosen node in the minimal time but to maximise the number of nodes over a given time, thus the end node could be any node that is available in the FOV at the chosen time.

None of the shortest paths graph theory problems investigated appear to be exact analogies to the STA problem. They all assume a target node (or visiting a fixed number of nodes) instead of finding an open-ended path that maximises the number of node visits in a given time. This means that conventional combinatorial optimisation methods employed for solving the TSP, VRP, SMSP and CTP cannot be used for solving the STA problem. Even if the optimisation occurs over a fixed number of tasks (rather than over time), the pool of tasks that can be selected is far larger than the size of the set to be optimised. This also renders conventional combinatorial optimisation techniques incompatible with the STA problem.

7.3 Science Operations Module

Due to the apparent incompatibilities of standard combinatorial optimisation algorithms with the STA problem the Science Operations Module (SOM) is designed using variants of these popular algorithms. There are two versions of this planning algorithm; one optimised version (the SOM) and one benchmark version (the BSOM) to act as a comparison tool. These two versions select the tour members following different selection procedures however

they share the same core algorithm. The constraints introduced earlier in the chapter are implemented within the planning algorithms as per
Table 7-1.

There are three factors that have been omitted mainly due to lack of data and three factors only partially implemented. The partial implementation of the manoeuvre time is a result of the trade-off in designing the SepM-MPA. The angular rate that governs the formation's manoeuvre duration is based on data from the Position Assignment Module (PAM) where a retarget manoeuvre of 90° typically takes 4500s to complete. Only this fixed formation angular rate (0.02 $°s^{-1}$) is used to calculate the manoeuvre duration. The planning algorithms do not take into consideration reconfiguration or resize manoeuvre influences on the manoeuvre duration. The other partially implemented factors add complexity to the problem but are not fully representative of a DARWIN-like mission. The only factor fully implemented is the 0.986 °/day angular rotation of the FOV.

Table 7-1 Science Operations Module constraints and implementation

Constraint	Inclusion	Implementation
Manoeuvre time	Partially	Fixed at 0.02 $°s^{-1}$
Calibration time	No	N/A
Observation Time	Yes	As per Table 7-2
Formation pointing constraints	Yes	FOV moves at a fixed rate of 0.986 °/day
Detection task scheduling	Partially	At least one other task must be scheduled before a detection task can be repeated[4]
Planet detection analysis time	No	Spectroscopy can be scheduled immediately after third detection task
Planet orbital characteristics	No	N/A
Science task weighting	Partially	Tasks are given equal weighting

7.3.1 Problem Definition

Consider a set N of nodes, representing science tasks, and a set S of star positions, representing the ecliptic longitude of each star in the catalogue. Each pair of nodes, $n_i \in N$ and $n_j \in N$, is connected by a edge $v_{i,j} \in V_i$, where the set V_i represents all possible edges $v_{i,j}$ connecting n_i to all other nodes. Relative weighting of the nodes with respect to the other nodes is given by the flag f ($f \in \mathbb{N}: \{1,\ldots,\infty\}$). If the node contains a flag (i.e. n_i^f) then the node (task) can gain priority scheduling with f=*1* the highest priority, f=*2* the next highest and so on. If the node does not contain a flag (i.e. n_i) then there is no weighting for that task. Flags and their priorities can be added to any task throughout the mission by the science team

[4] With this definition the time between repeated detection tasks will up to ~2.5 days (Table 7-2) however in reality the required separation time is likely to be much larger (>90 days) depending on the orbital characteristics of the planet

to prioritize tasks within the schedule. M_i is a set of nodes ($\subset N$) representing all achievable tasks from node n_i connected by edges in the set U_i ($\subset V_i$). A diagram of the nodes and sets defined thus far is given in Figure 7-2 and the construction of M_i is detailed in the next section.

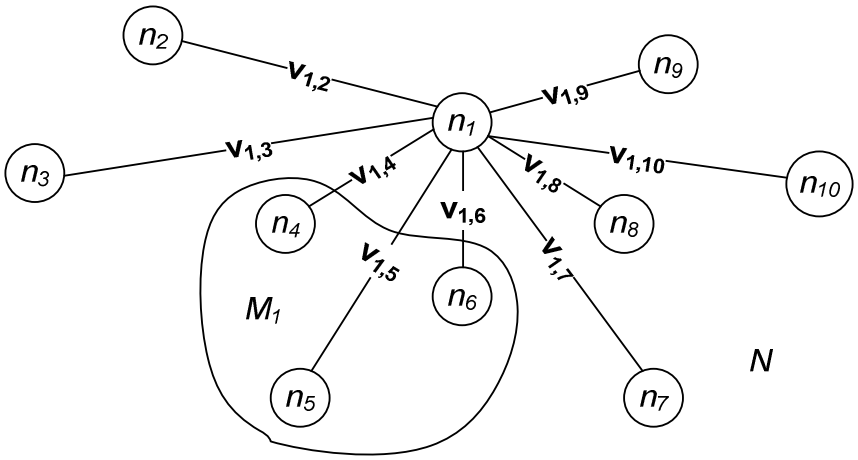

Figure 7-2 Diagram of the node and edge sets used in the Science Operations Module

Let $l_{i,j}$ be the length (in days) of the edge $v_{i,j} \in V_i$ connecting n_i to $n_j \in M_i$,

$$l_{i,j} = m_{i,j} + c_j \tag{7.1}$$

where $m_{i,j}$ is the manoeuvre duration and c_j the task completion time. For this implementation, $m_{i,j}$ is calculated from the angular separation of the initial star, $s(n_i) \in S$, and the target star, $s(n_j) \in S$ and the fixed formation angular rate, $\dot{\alpha} = 0.02\ ^\circ s^{-1}$,

$$m_{i,j} = \dot{\alpha} \left| s(n_i) - s(n_j) \right| \tag{7.2}$$

If $s(n_i) = s(n_j)$ then $m_{i,j} = 0$.

The tour is a set T, of p edges, $v_{i,j} \in T$, of duration L where

$$L = \sum_{k=1}^{p} l_{i,j}^{k} \tag{7.3}$$

and

$$L \geq \mathrm{T_{target}} \tag{7.4}$$

with $\mathrm{T_{target}}$ as a user defined minimum tour duration. The goal of the SOM is to find a tour that maximizes the number of edges within the minimum tour duration $\mathrm{T_{target}}$. $\mathrm{T_{target}}$ can have a range from 0.116 days (the shortest task duration) to the entire mission lifetime of 5 years however the size of $\mathrm{T_{target}}$ will directly affect the length of time it takes to calculate a tour and

the relative performance of the tour. For example, a combination of five T_{target} = 4 days tours will be quicker to calculate than one T_{target} = 20 days tour but the latter will likely be a more optimal tour than the former. T_{target} is a minimum tour duration (as opposed to a maximum tour duration) so that the tours generated have a duration of at least T_{target}. If T_{target} was a maximum tour duration the user would have less control over the length of the tour desired.

7.3.2 Core Algorithm

At each node n_i, the set of nodes M_i is generated from which the set of edges U_i can be calculated. This allows for the scheduling constraints detailed in sub-section **Error! Reference source not found.** to be applied to V_i. M_i is found as follows:

1. Generate set $O_i \subset N$, representing all possible nodes with star positions $s(n_i) \in S$ such that:

$$\beta(t_i) - \frac{\pi}{4} \leq s(n_i) \leq \beta(t_i) + \frac{\pi}{4} \tag{7.5}$$

   where $\beta(t_i)$ is the ecliptic longitude of the anti-sun vector and t_i is the time when the node (task) n_i completes. O_i contains all the tasks for stars within the FOV bounded by ±45° from the anti-sun vector.
2. Calculate set W_i, representing all the edges connecting n_i to $n_j \in O_i$.
3. To obtain the set of edges U_i (and hence M_i) remove from the set of edges, W_i, all the edges where:

$$n_i = n_j \tag{7.6}$$

   and

$$s(n_j) \leq \beta(t_j) - \frac{\pi}{4} \tag{7.7}$$

   where

$$\beta(t_j) = \beta(t_i) + l_{i,j}\dot{\alpha} \tag{7.8}$$

   $s(n_j)$ is the star position for node n_j, $\beta(t_j)$ is the ecliptic longitude of the anti-sun vector and t_j is the time when the node (task) n_j completes. The first condition, shown in Equation (7.6), ensures that the task the formation is originally assigned to, n_i, is not included in the set M_i, thus preventing task duplication. The second condition, shown in Equation (7.7), eliminates any tasks that cannot be scheduled as their star would leave the field of view before the task could be completed. After applying the constraints, U_i contains all the achievable edges from n_i. The set of nodes M_i can be obtained by examining U_i and extracting all the nodes n_j.

4. If during 3. $U_i \Rightarrow \varnothing$ then no suitable edges exist and a wait edge of 24hr duration is imposed on the tour to allow $\beta(t_i)$ to move across the sky. The 24 hr wait duration is chosen as it provides a ~72 % chance that a new star will enter the field of view during the wait. If no new star enters the field of view during the wait period then another wait period is imposed on the tour. If two 'waits' are required then the chance that a new star will enter the field of view during this second wait period rises to ~92%. The 24 hr wait, therefore, provides a balance between a long enough duration to increase the chance that a new star will enter the field of view but not too long as to waste time[5].

A diagram showing the set definitions for M_i and U_i is given in Figure 7-3. In this example the set $N = \{n_1 \ldots n_{10}\}$ is reduced to the set $O_1 = \{n_2, n_4, n_5, n_6\}$ through application of the field of view defined by $\beta(t_1)$ in Equation (7.5). From O_1 the set $W_1 = \{\mathbf{v}_{1,2}, \mathbf{v}_{1,4}, \mathbf{v}_{1,5}, \mathbf{v}_{1,6}\}$ of edges is calculated. Applying the conditions in Equations (7.6) and (7.7) for node n_2 the edge $\mathbf{v}_{1,2}$ is removed because n_2 is outside the $\beta(t_2)$ field of view. Assuming no violations of Equations (7.5) and (7.6) occur for nodes n_4, n_5 and n_6 the final edge set is $U_1 = \{\mathbf{v}_{1,4}, \mathbf{v}_{1,5}, \mathbf{v}_{1,6}\}$, and hence, $M_1 = \{n_4, n_5, n_6\}$ is obtained. M_1 contains all the nodes (i.e. science tasks) that can be completed from node n_1.

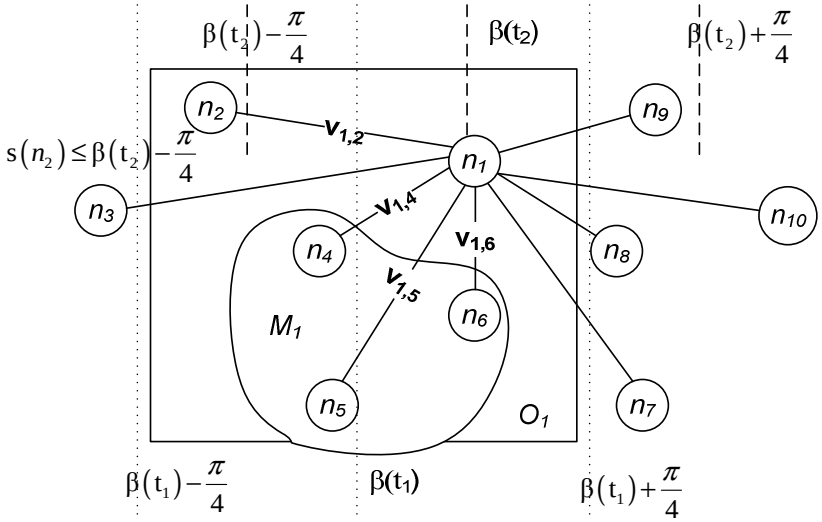

Figure 7-3 Set definitions for generating M_i and V_i in the Science Operations Module

After U_i is calculated a decision process (described in the following sub-sections) selects $v_{i,j} \in U_i$, and the edge is added to the tour. A file recording the completion status of each task

[5] The preceding percentages are obtained by calculating the differences in ecliptic longitude of adjacent stars in the catalogue as a function of $\dot{\alpha}$ to find the field of view slew times between adjacent stars.

(the taskflag file) is then updated, thus ensuring an up-to-date reference for the next iteration of the decision process. This process is repeated for n_j etc. until $L \geq T_{target}$, at which point the tour is terminated. A flow diagram showing this core algorithm can be found in Figure 7-4. Edge selection is the process that differs between the SOM and BSOM approaches to the STA problem.

The goal of the SOM is to maximize the number of edges (and hence the number of completed tasks) in the tour, *T*. Since there is no fixed tour time (T_{target} only represents a minimum tour time) the metric chosen to rate the performance of a tour is the task/time ratio, *R*:

$$R = \frac{(p - p_{waits})}{L} \tag{7.9}$$

where p_{waits} represents the total number of wait edges imposed on the tour (see preceding step 4). In the tour performance calculation p_{waits} is removed from the total number of edges, *p*, since the wait edges account for scheduled time not allocated to science observations. The time signature of the wait edges (24 hrs per wait) remains in the total tour time, *L*, because wait edges negatively affect the performance of the tour. In evaluating the performance of a tour a higher task/time ratio indicates a better use of time within the tour.

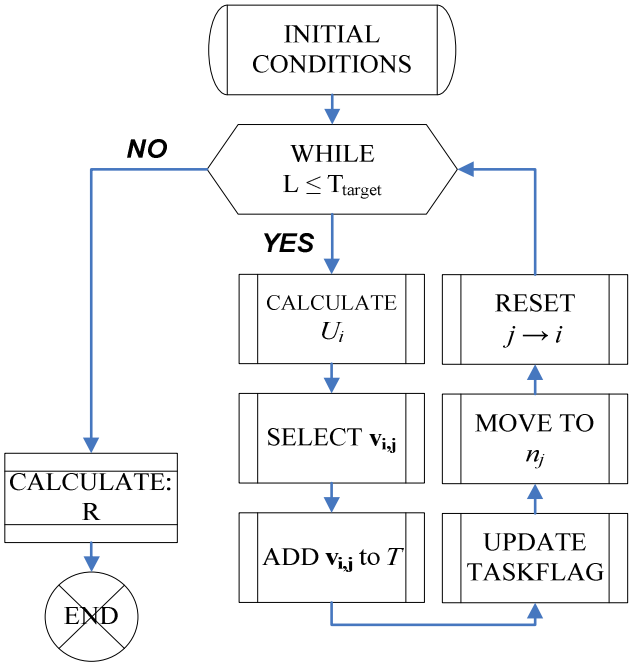

Figure 7-4 Flow diagram of the shell science operations planning algorithm

7.3.3 Benchmark Edge Selection Process

The benchmark edge selection algorithm implemented in the BSOM uses the principle of nearest neighbour selection (Hurkens and Woeginger, 2004). The nearest neighbour here refers to nearest neighbouring *tasks* (as a function of time) as opposed to nearest neighbouring *stars* (as a function of angular separation). The nearest neighbouring task is the

one that takes the shortest time to complete. At each node one edge needs to be selected from a group of all available edge for that node (i.e. the set U_i). For the benchmark, the selected edge is simply the one with the smallest completion time associated with it:

$$v_{i,j} = \min_{l_{i,j}} U_i \tag{7.10}$$

The BSOM is 'greedy' in nature in that it always selects the shortest edge regardless of how it affects the tour later on. The edge selection is locally optimal over one node but a tour constructed using the benchmark is unlikely to be globally optimal over T_{target} because the algorithm only examines a small section of the complete solution space. This method, however, is useful as a benchmark as it is intuitive, easy to implement and provides solutions quickly.

7.3.4 Optimised Edge Selection Process

The stochastic version of the edge selection algorithm implemented as the SOM is a simplified version of the basic ant colony optimisation meta-heuristic (Dorigo, 1999). As before, at each node one edge needs to be selected from a group of all available edges for that node (U_i). This edge is chosen using a weighted stochastic process. U_i is sorted in ascending order of $l_{i,j}$ such that:

$$V_{i,sort} \in \left\{ \min_{l_{i,j}} U_{i\,1}^{\,k}, \ldots\ldots\ldots, \max_{l_{i,j}} U_{i\,k}^{\,k} \right\} \tag{7.11}$$

Some of the nodes that the edges in $V_{i,sort}$ point to may have prioritisation flags (i.e. n_i^{f}). These edges are moved to the beginning of $V_{i,sort}$ in order of prioritisation to create the set V^*. Therefore V^* is the set U_i sorted into ascending order of edges length and rearranged to ensure prioritised edges are at the beginning. $\mathbf{v}_a^{*k}$ is the a^{th} element of V^* which has k elements (i.e. $a \in \{1,2\ldots\mathrm{k}\}$). If the first element in V^* represents a prioritised edge then the SOM selection is simply this element:

$$\mathbf{v_{i,j}} = \mathbf{v}_1^{*\mathrm{k}} \quad \text{if} \quad \mathrm{f} \neq \varnothing \quad \text{in} \quad n_j^{\mathrm{f}} \tag{7.12}$$

If there are no prioritised edges then a separate decision process is implemented.

The set E represents a number line that linearly maps to V^*:

$$E \rightarrow V^* \tag{7.13}$$

such that

$$E_0 \rightarrow v_1^{*k} \tag{7.14}$$

and

$$E_3 \rightarrow v_k^{*k} \tag{7.15}$$

A random number, *r*, is generated from a normal distribution with a mean of zero and a variable standard deviation (σ). When σ=1, there is a 99.7% chance that $|r|$ will lie between 0 and 3. The edge is chosen such that:

$$E_{|r|} \rightarrow v^{*\,k}_{\;a} = v_{i,j} \tag{7.16}$$

with

$$a = \left\lceil \frac{|r|}{3/k} \right\rceil \tag{7.17}$$

(i.e., *a* is rounded up to the nearest integer).

For example, suppose V^* contains 20 edges (i.e. $k = 20$) if $|r| = 0.5$, then $a = 4$. The selected edge is therefore the 4th edge in V^*:

$$E_{|0.5|} \rightarrow v^{*\,20}_{\;4} = v_{i,j} \tag{7.18}$$

The complete selection process can be visualised more clearly in Figure 7-5. Equation (7.16) and Equation (7.17) are invalid for cases where $|r| > 3$ thus a final constraint needs to be introduced,

$$v_{i,j} = \max_{l_{i,j}} U_i \equiv v^{*\,k}_{\;k} \quad : \quad |r| > 3 \tag{7.19}$$

Because $|r|$ is weighted towards zero, the selection of $v_{i,j}$ is weighted towards those $v_{i,j} \in V^*$ with shorter task completion times $l_{i,j}$. The strength of this weighting can be modified by altering σ. As $\sigma \rightarrow 0$, the probability of $|r|$ being a small number increases thus the likelihood of selecting shorter edges increases. Eventually however, σ will be so small that only one choice of edge becomes available and the algorithm will emulate the benchmark. σ is a tuneable parameter within this algorithm.

This stochastic approach to edge selection is not locally optimal but seeks to find better tours than the benchmark as it has the chance to examine a larger area of the solution-space. It is unlikely however that a better solution will be found on the first attempt and so this approach is enhanced by multiple iterations with the chosen tour being the best performing of the family of tours generated. The number of iterations, *N*, is another tuneable parameter for this algorithm and greatly affects its performance and efficiency.

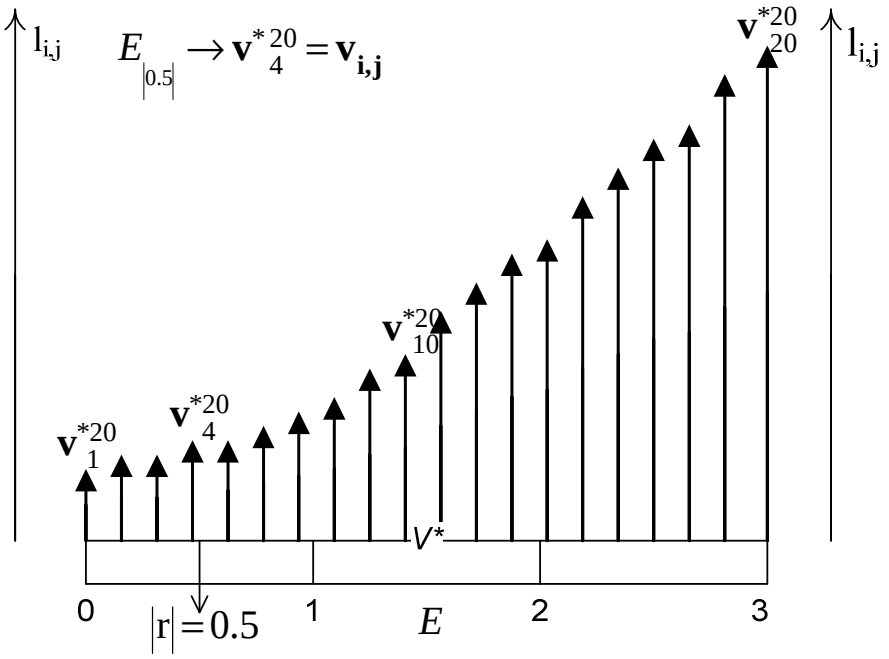

Figure 7-5 Example edge selection for the Science Operations Module. With 20 edges sorted into order of ascending duration a random number of 0.5 means the 4th edge is chosen.

The use of stochastic methods in a planning algorithm to be employed in an autonomous spacecraft environment is a potentially risky decision. The stochastic nature of the algorithm means that there is always the possibility that multiple 'bad' decisions can results in very poor performance of the algorithm. However, checks and balances can be added to the algorithm to ensure that a poor result is not acted upon. For example, the BSOM tour could always be made available, so if the SOM fails to find a better tour in the time given, the BSOM tour (that is the sum of locally optimal selections) would be considered the optimal tour.

7.4 Analysis Initial Set-up

In order to analyze the performance of the BSOM and SOM planning algorithms it is necessary to define some initial conditions that affect the calculations.

7.4.1 Star Catalogue and Observation Tasks

The star catalogue used for the algorithm contains coordinate data for the stars and task duration data for the observations to be performed on those stars. The catalogue of stars for the DARWIN mission was provided by den Hartog (2006). The catalogue contains 447 stars and the star distribution can be found in Figure 7-6. The star distribution is fairly even in both ecliptic longitude and latitude which will assist the SOM in avoiding large and time-consuming retarget manoeuvres. Task durations are also included in the catalogue for every star and every task using tables from Karlsson, et al. (2004). Durations for every task are set depending on the spectral class of the star as in Table 7-2. The task times for the F, G and K stars are taken direct from (Karlsson, et al., 2004) and the task times for the M stars are

interpolated (since Karlsson, et al., (2004) omits M stars from their analysis). This interpolation is based on the observation that as the star's spectral class changes from F to K the task observation duration decreases. Since M stars follow K stars in the spectral type sequence the task time for M stars is assumed less than that for K stars. This interpolation is very rough and in no way indicates the actual observational task durations required for M stars. The task times do however give a reasonable indication of typical task durations for initial analysis purposes. As there are 7 tasks for each of the 447 stars the star catalogue holds 3129 task duration values. However, since each detection task is required to be repeated three times the total number of potential tasks for the mission is 4023.

Table 7-2 Task times for stars of different spectral types (Karlsson, et al., 2004)

Spectral Type	**Task Time (days)**						
	Detection	**Spectroscopy**					
		<7.2 μm	8-9.2 μm	9.2-10 μm	10-13.2 μm	13.2-17.2 μm	> 17.2 μm
F	2.410	1676[6]	9.450	34.60	1.900	3.410	62.500
G	0.533	85.80	1.560	6.620	0.386	1.040	34.900
K	0.339	38.30	0.987	3.950	0.228	0.618	26.100
M	0.300	10.89	0.460	1.965	0.116	0.376	20.466

7.4.2 Simulating Planet Detection

In order for the algorithm to allow the scheduling of spectroscopy tasks some of the stars in the catalogue must have planets. As the existence of these planets is not known it becomes necessary to simulate this requirement. The simulated probability of a star having a planet is arbitrarily set to 10%. The distribution for the stars with planets can be seen in Figure 7-6 and shows there are 47 stars with planets, distributed fairly evenly in ecliptic longitude but biased in negative ecliptic latitude. It is reasonable to assume that no one portion of the sky will yield more planets than the next so this distribution can be viewed as fairly typical of what one might find in reality. The even distribution in ecliptic longitude needs to be noted as it will reduce the number of wait edges required during the latter stages of the mission.

[6] The F-type star <7.2 μm spectroscopy task duration is 1676 days as defined in Karlsson, et al. (2004). However Karlsson, et al. (2004) notes that this is too long and advises measures to decrease this time.

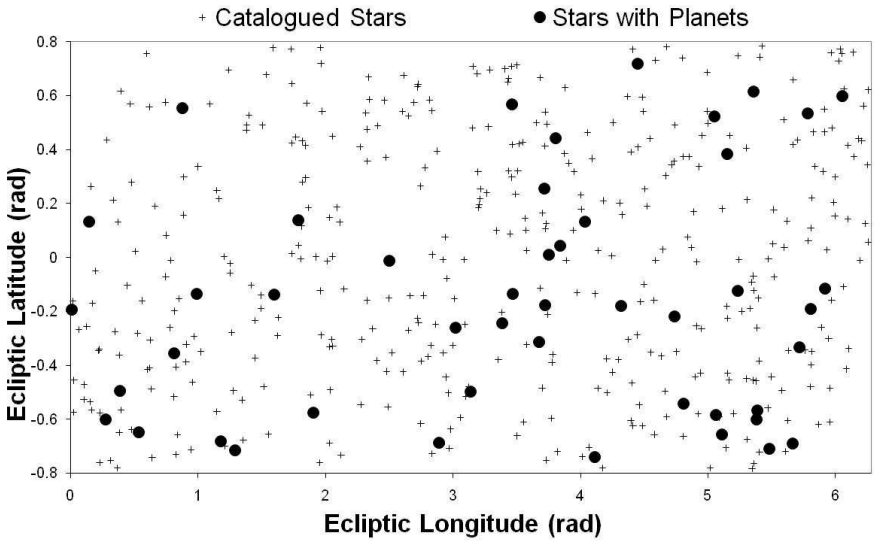

Figure 7-6 Distribution of catalogued stars (+) and stars with planets (•) for the Science Operations Module analysis.

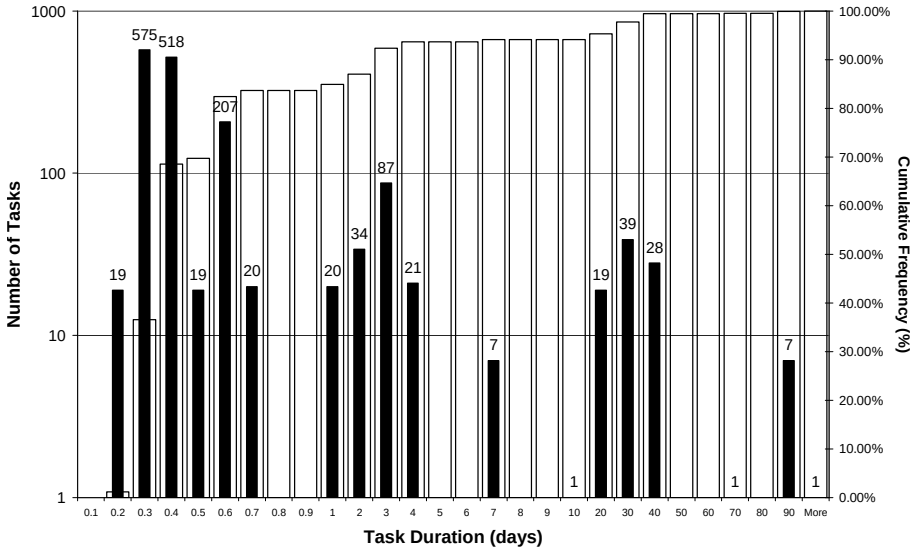

Figure 7-7 Histogram showing distribution of tasks by task duration and the cumulative frequency.

As previously calculated there are 4023 potential tasks for the mission. However this number was calculated assuming a 100 % planet to star probability. Using the 10 % planet to star probability shows that, for this analysis, there are 1623 tasks for the mission of summed duration ~14.8 years. Figure 7-7 shows a histogram of the task distribution by task duration in days (on the primary y-axis) and the cumulative frequency (on the secondary y-axis). The histogram is dominated by < 1 day task durations (~ 82% of all tasks). Of these, the vast majority are the detection tasks as can be seen by cross-referencing with Table 7-2. Regardless of the data in the planet file the task duration histogram will always look similar due to the detection tasks that will always need to be performed. The data also emphasises the small number of F-type stars present in the Darwin catalogue (~ 6.5% of the catalogue) and the inclusion of only one F-type star with a simulated planet (~ 2 % of the chosen stars).

7.4.3 Initialising Mission Stage Markers

The performance of both the planning algorithms will be affected by the time in the mission the algorithms are run. At the beginning of the mission, with no planets detected, the majority of tasks scheduled will be the shorter detection tasks. As the mission progresses, less detection tasks will be available and the longer spectroscopy tasks will start to be scheduled (since planets will have been found). By the end of the 5 year mission duration only spectroscopy tasks will be scheduled as all the detection tasks will have been performed. Performance analysis of the planning algorithms at different stages of the mission is carried out using taskflag files that record a snapshot of the current task distribution for every year of the mission as in Table 7-3. Here 'Year0' describes the beginning of the mission; 'Year1' describes the end of the 1st year of the mission and so on. The data in Table 7-3 were obtained using the BSOM with a target tour time of 365 days, 730 days etc. Table 7-3 shows by the beginning of the fourth year of the mission ('Year3' taskflag file) all the detection tasks have been completed and only spectroscopy tasks remain.

Table 7-3 Simulated task distribution for 5 mission stages

Taskflag	**Detection Tasks**		**Number of Planets**		**Spectroscopy Tasks**	
	Completed	Remaining	Found	Not Found	Completed	Remaining
Year0	0	1341	0	47	0	282
Year1	1044	297	38	9	60	222
Year2	1313	28	47	0	160	122
Year3	1341	0	47	0	201	81
Year4	1341	0	47	0	222	60

7.4.4 Calculation Count

To aid understanding of the amount of processor resources required to run the SOM the core algorithm and the two edge selection functions have been seeded with a calculations counter. This counter is incremented at various stages throughout the SOM. The counter does not record every calculation performed by the SOM as it is not present within embedded routines (like sin, cos, etc) and embedded MATLAB functions so its actual value has no meaning. However the calculations value for different SOM iterations does provide a comparison of processor resources used. Obviously including this counter does have a detrimental effect on the performance of the algorithm but this performance reduction is the same for each execution of the SOM so the comparison remains valid. A calculations counter was used as the comparison metric instead of the algorithm runtime as it allows analysis to be performed that can be independent of the type of processor used. Using the PC system described above a rate of ~ 6.8×10^6 calculations per second were recorded. For an autonomous SOM the tour must be generated before the end of the current observation. The shortest observation time is ~ 10000s (see Table 7-2) giving a calculation limit of ~ 6.8×10^{10} calculations to generate the tour for the PC system the analysis was performed on.

7.5 Science Operations Module – BSOM Tour Example

A sample BSOM tour is shown in Figure 7-8 generated using the Year2 taskflag file, T_{target} = 80 days and starting from star 200 in the catalogue. The star index numbers from the mission catalogue can be seen next to the star's ecliptic co-ordinates. The stars are joined by a dotted line and arrows represent the direction of travel. Multiple arrow-heads along one direction represent the number of times the direction is repeated. For clarity the tour data is given in Table 7-4. The stars are represented by their index numbers in the mission catalogue. Task '1' represents a detection task and task '4' represents the 9.2-10 μm spectroscopy task. The time row represents the cumulative completion time of the task at each star (incorporating the manoeuvre time to retarget to the star). The entire tour takes 81.73 days to complete and with 28 tasks gives R = 0.3436 tasks/day.

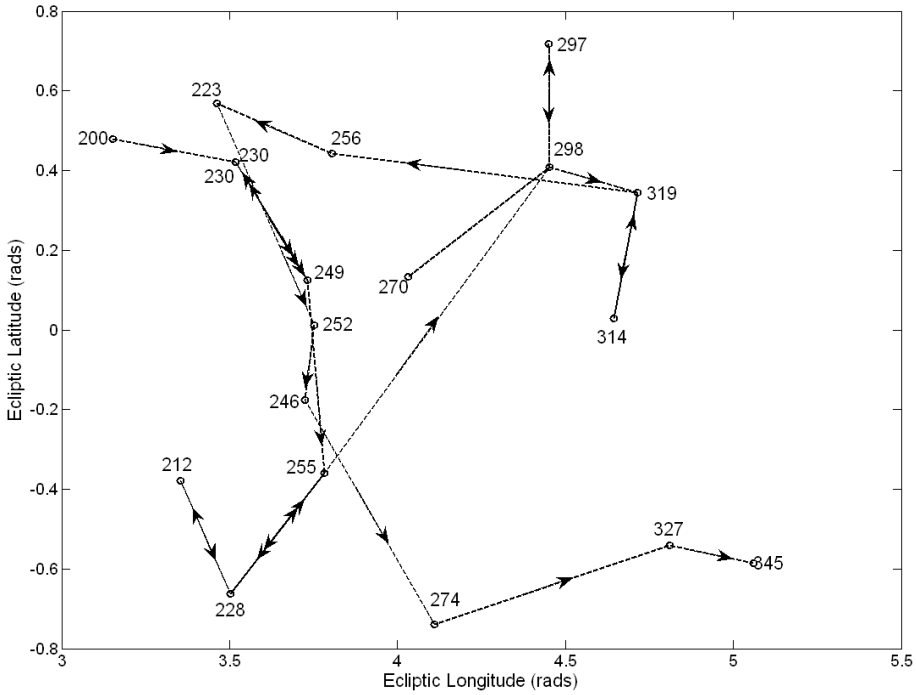

Figure 7-8 Sample Benchmark Tour (Year2, T_{target} = 80 days, Startstar = 200). The numbers represent the star index in the catalogue and the arrows show the direction of the star tour.

The tour is initially dominated by detection tasks on F-type stars and then the spectroscopy tasks on K-type stars. The tour is very dynamic with large retarget manoeuvres implemented both in increasing and decreasing ecliptic longitude. This is because the field of view (FOV) moves very slowly in relation to the frequency of the tasks scheduled and the manoeuvre time is significantly less than the observation time. There are a number of parts of the tour that involve multiple retargets between two stars (i.e. between stars 230 and 249 and between stars 255 and 288). This is due to the constraint that prevents detection tasks being scheduled consecutively. This implementation of the detection task scheduling constraint (Table 7-1) creates a scheduling situation that is unlikely to be replicated on a real mission as the detection task repetition rate would have to be far lower than that simulated here.

Table 7-4 Sample Benchmark Tour (Year2, T_{target} = 80 days, Startstar = 200)

Star	200	230	249	230	249	230	249	255	228	
Task	1	1	1	1	1	1	1	1	1	
Time (days)	0	2.42	4.84	7.26	9.69	12.11	14.53	16.96	19.38	
Star	212	228	255	228	255	298	297	298	270	298
Task	1	1	1	1	1	4	1	4	1	1
Time (days)	21.80	24.22	26.64	29.07	31.49	33.93	37.89	40.31	44.28	46.71
Star	319	314	319	256	223	252	246	274	327	345
Task	1	1	4	4	4	4	4	4	4	4
Time (days)	49.12	51.54	53.96	57.94	61.90	65.87	69.83	73.80	77.77	81.73

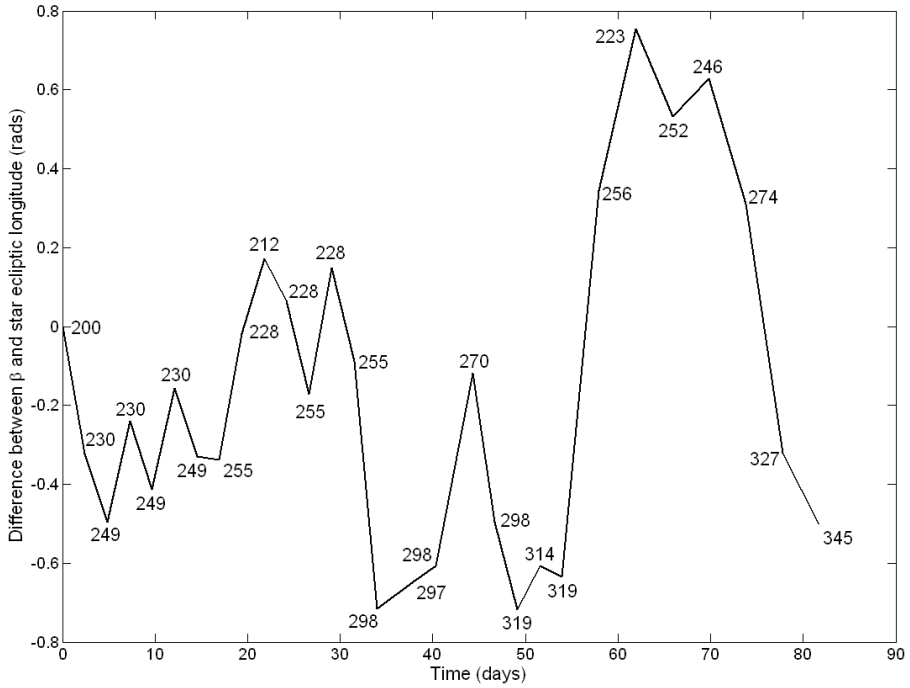

Figure 7-9 Difference between the longitude of the anti-sun vector (β) and the ecliptic longitude of the formation pointing direction at task completion throughout the benchmark sample tour. Negative differences indicate a pointing direction ahead of β whilst positive differences indicate a pointing direction behind β.

The dynamic nature of the tour as illustrated in Figure 7-8 is emphasised in Figure 7-9. Here the difference between the ecliptic longitude of the anti-sun vector (β) and the ecliptic longitude of the formation pointing direction at task completion is plotted against the cumulative task completion time. Longitude values close to the leading edge of the FOV are represented on the negative y-axis whilst those longitudes close to the trailing edge of the FOV are represented on the positive y-axis. The data show that for this tour the full longitudinal range of the FOV is used but never breached (the spectroscopy task on star 223 ends at less than 2° from the trailing edge of the FOV).

7.6 BSOM Analysis

The BSOM proves to be extremely quick in finding solutions even for large values of T_{target}. Table 7-3 shows the task distribution of the taskflag files generated using the BSOM. A 365 day tour typically takes less than 5 min to generate on the hardware detailed in sub-section 6.6. As expected the majority of the detection tasks get completed within the first year of the mission. This is partly because no spectroscopy tasks can be scheduled until a planet has been detected (after three detection tasks on the same star) and because the detection tasks typically have a shorter duration than the spectroscopy tasks. The detection tasks are fully complete by the end of year 3. The Year2 taskflag is of note because although all the planets have been found there are still 28 detection tasks remaining to be scheduled. Obviously in a real mission environment the user will not know if all the planets have been detected and so all the detection tasks must be completed regardless. By the end of year 4 ~ 96% of all the tasks have been completed with only 60 spectroscopy tasks remaining. However these last 4% of tasks have very long observation times associated with them and it would be impossible to complete all the remaining tasks by the end of year 5.

Table 7-5 Performance data for the BSOM

Taskflag	Target tour time (days)	Task/time ratio (tasks/days)	Mean task duration (days)
Year0	5	3.2977	0.3032
Year1	40	1.5533	2.3858
Year2	80	0.1997	5.3014
Year3	160	0.0510	10.3836
Year4	320	0.0389	25.3966

Running the SOM for T_{target} = 365 days for comparison purposes with the data in Table 7-3 is impractical in terms of calculation time so a different data set is required for the comparison. This is shown in Table 7-5 with the data from the BSOM. The difference in T_{target} for the different taskflag files is partly to allow long enough tours to be generated and partly to ease the computation burden on the SOM (i.e. Year0 T_{target} = 5 days rather than 20 days) since the same T_{target} values are used in the SOM analysis, see Section 7.7. The data show an apparent performance decrease as the mission progresses. This is to be expected as the availability of shorter tasks decreases over time. This results in an increase in the mean task duration for the generated tours as shown Table 7-5. Comparison of the Year2 performance in Figure 7-9 and that of the sample benchmark tour shows that the sample tour has a performance increase of over 60%. This is due to the differing starting positions for the tour (the sample tour starts at star 200 whilst in Table 7-5 the tour starts at star 1). This emphasises the importance of fixing the tour starting position when comparing tour performance.

A comparison of processor resources required for each taskflag can be seen in Table 7-6. The data show that the computational burden of the BSOM decreases throughout the lifetime of the mission even though T_{target} increases. The Year0 result is ‘low’ due to its T_{target} value.

Increasing the Year0 T_{target} to 20 days yields a calculation count of ~ 18.8×10^6 and a tour length of 68 tasks. The decrease of calculations observed is expected since with a reduction in the number of possible tasks later on in the mission there are fewer edges to calculate at each node. There is an anomaly in the calculations data for the Year4 taskflag since the calculation count is larger than Year3. This is due to the Year4 tour having more tasks scheduled than the Year3 tour. Four of those extra tasks are wait edges with durations of 24hrs. This a relatively short period for a typical Year4 task and essentially imposes an extra calculation burden on the algorithm that would otherwise not be present if there were more tasks available for scheduling. The task availability from the end of year 4 however is sparse (see Table 7-3) and so the inclusion of wait edges is much more likely at this stage of the mission. Finally the calculations data show that for all the taskflag files the calculations required to generate the benchmark tours are well below the limit of ~ 6.8×10^{10} calculations imposed by the autonomy requirements[7].

Table 7-6 Calculations count comparison for the BSOM and task data

Taskflag	Target Tour Time (days)	'Calculations' count	No. of tasks[8]	No. of wait edges
Year0	5 (20)	4,828,631 (18.8×10^6)	17 (68)	0 (0)
Year1	40	6,822,655	63	0
Year2	80	1,192,842	18	0
Year3	160	261,573	9	0
Year4	320	524,314	17	4

7.7 Science Operations Module Analysis

The stochastic version of the planning algorithm has two variable parameters that require analysis: the number of iterations required to produce a high performing tour and the standard deviation (σ) for the random number generator. These are investigated separately in the following sub-sections.

7.7.1 Number of Iterations

For this analysis the same T_{target} for each taskflag file is used as in the BSOM analysis and $\sigma = 1$. A typical data set can be found in Figure 7-10. The tour generated is a T_{target} = 80 days tour using the Year2 taskflag file. The *x*-axis (displayed on a logarithmic scale) represents the number of iterations with the iteration count doubling (i.e. 20, 40, 80, etc.). The *y*-axis represents the tour performance, *R*, of the highest performing tour found in the family generated. The data points (×) represent generated tours and there are 20 for each iteration count. The data points (●) represent the mean performance for each iteration count and the line represents a logarithmic fit of the mean data.

[7] Except the 20 days Year0 tour.

[8] Including wait edges

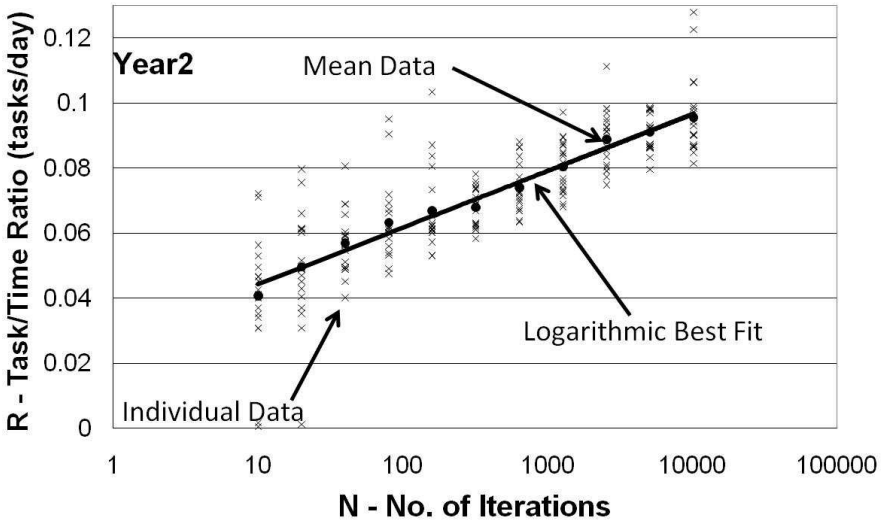

Figure 7-10 Tour performance vs. no. of iterations for Year2 taskflag, T_{target} = 80 days. Individual data points are represented by a (x), mean data points by a (•) and the solid line is a logarithmic best fit for the mean data. The chart shows that for tours calculated in Year2 of the mission the tour performance increases logarithmically with increasing iterations.

The data show that the SOM can produce very different highest performing tours when identical initial conditions are used (indicated by the spread of data points for each iteration count). A general trend formed, however, is that the performance increases with increasing iteration count. The logarithmic nature of this trend introduces a problem since increasing the iteration count decreases the performance gains achievable. This is expected because for large enough numbers of iterations the algorithm will eventually find the true optimal tour and a maximum tour performance. The number of iterations required to achieve this, however, is likely to exceed the time limitations imposed on the calculation.

Figure 7-11 shows the calculation count for increasing values of iteration count for the same tours generated for Figure 7-10. The *x* and *y*-axes are linear. The data points (●) represent the mean calculation count from the 20 generated tours at each iteration count and the line represents a linear best fit of the mean data. The individual tour data (x) are not shown as their spread is indiscernible from the size of the data point markers (●). The slope of the best fit line is also included indicating that on average ~ 131,145 calculations are performed for every iteration. The data from Figure 7-11 show that the relationship between the number of iterations and the number of calculations performed is clearly linear. So at high values of iteration count whilst doubling the number of iterations doubles the computational burden the resulting tour performance is only marginally increased. This is a clear example of the law of diminishing returns and care must be taken to select an iteration count value that can find a good performing tour without using up too many processor resources. Comparing the mean calculation per iteration with the benchmark Year2 calculation count shows that the SOM calculation count is much lower (~11% of the benchmark). This is because for a large number of the tours generated by the SOM fewer tasks are scheduled than for the tour generated by the BSOM. Scheduling those fewer tasks results in fewer calculations.

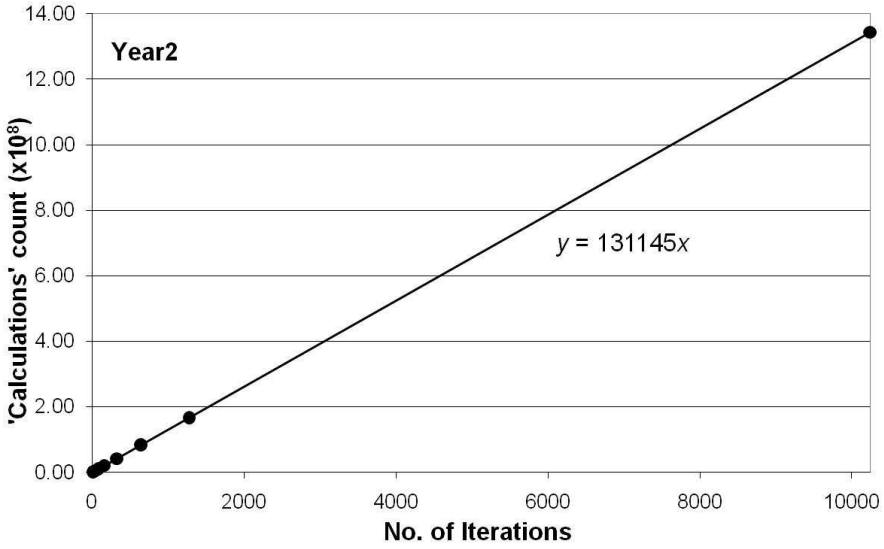

Figure 7-11 Calculations count vs. no. of iterations for Year2 taskflag, T_{target} = 80 days. Mean data points (•) and a linear best fit of the mean data point (solid line) are shown. The chart shows that the number of calculations executed during the Year2 mission stage increases linearly with the number of iterations.

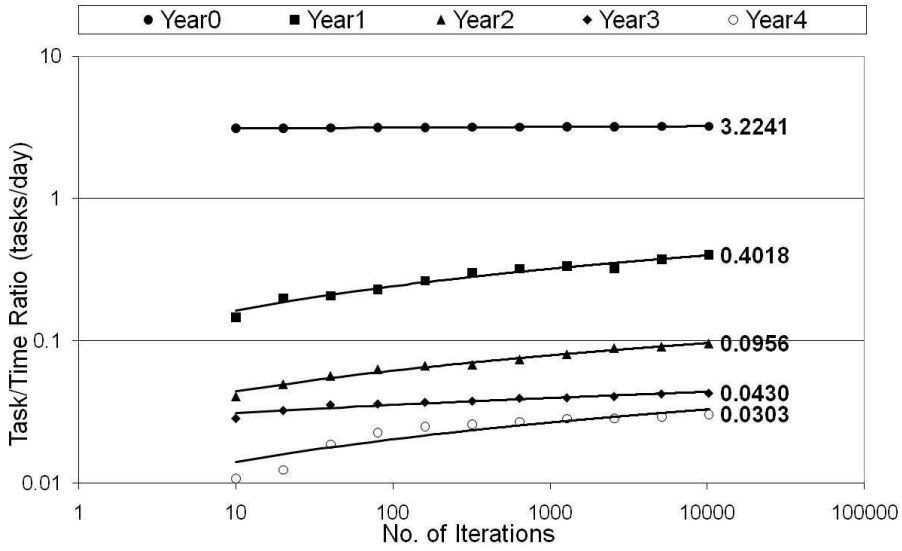

Figure 7-12 Mean tour performance vs. no. of iterations for all taskflag files. Mean data points are represented by markers whilst the solid lines represent logarithmic best fits of the mean data. The chart shows that for all missions stages the maximum tour performance logarithmically increase with increasing iterations.

The trends present in Figure 7-10 and Figure 7-11 are the same for all the stages of the mission. Figure 7-12 shows how the mean tour performance (from 20 generated tours) varies with increasing iteration count for each taskflag file. Both the *x* and *y*-axes are on a logarithmic scale to aid visualisation of the data. The lines represent the logarithmic best fit for each mean data set. For each taskflag the performance shows a logarithmic increase with increasing iteration count (this includes Year0 though it is barely discernible from Figure 7-12 due to the scale). The numbers next to each 10240 iterations data point show the mean

performance ratio for each taskflag at 10240 iterations. As for the BSOM, the performance of the SOM appears to decrease with mission time however, for reasons explained previously, this is expected. Comparing the mean performance data with the BSOM reveals that the SOM at 10240 iterations cannot find better performing tours than the BSOM for any taskflag for the initial conditions provided. This can be put down to two factors:

- There are not enough iterations of the algorithm being performed
- The standard deviation ($\sigma = 1$) is too large

7.7.2 Fine-Tuning the Number of Iterations

One way of achieving better performing tours using the SOM is to increase the number of iterations the algorithm performs. As seen previously however this follows the law of diminishing returns and could be limited by the calculations limit imposed on the simulation defined earlier. Figure 7-13 shows how the mean calculation count (from 20 generated tours) varies with increasing iteration count for each taskflag file. Both the *x* and *y*-axes are on a logarithmic scale to aide visualisation of the data. The lines represent the linear best fit for each mean data set. Though on a logarithmic scale the taskflag lines all have the same gradient, on a linear scale the gradients are different. Also shown on Figure 7-13 is the calculations limit defined earlier in the chapter. As seen in the BSOM analysis the number of calculations decreases as the mission progresses but each taskflag calculation count increases linearly with an increasing number of iterations. From the data it is possible to calculate the requirements for the SOM to find the benchmark tour whilst maintaining the original initial conditions.

The logarithmic best fit of the mean performance data for each taskflag can be used to calculate the number of extra iterations that would be required by the SOM to find the benchmark tour (or a tour performing equally as well). This calculation assumes that at high iteration counts the logarithmic fit of the data is still applicable. The results can be seen in Table 7-7. The best fit equations were calculated using the Microsoft Excel '*trendline*' function and have *x* representing the number of iterations and *y* representing the tour performance. For all taskflag cases the number of iterations required is significantly more than the 10240 iterations used in the previous analysis. How these figures relate to the calculation limit is shown in Table 7-8. The gradients of the linear best fit lines for each taskflag (from Figure 7-13) can be used to calculate how many calculations would be required to perform the necessary iterations to find the benchmark tour. With the calculation limit set at ~ 6.8×10^{10} the data show that only the Year3 and Year4 tours could find the benchmark tour (or a similarly performing one).

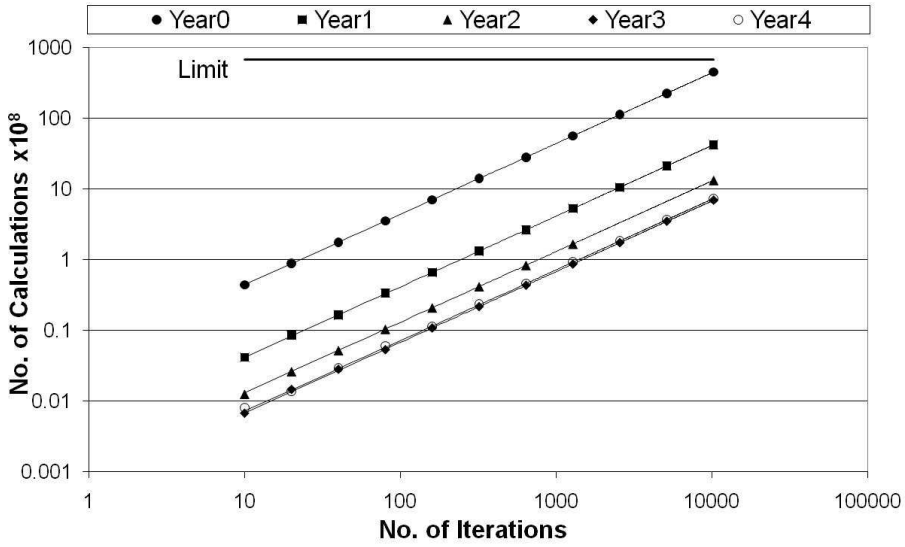

Figure 7-13 Mean calculation count vs. no. of iterations for all taskflag files. Mean data points are represented by markers whilst the solid lines represent linear best fit lines to the mean data. The chart shows that none of the taskflag files reach the calculations limit within 10^3 iterations.

Table 7-7 SOM data extrapolation to equal BSOM performance

Taskflag	Logarithmic Fit Equation	Benchmark Task/Time Ratio (tasks/day)	Iterations required to reach benchmark
Year0	$y = 0.0154\ln(x) + 3.0853$	3.2977	~ $9.76\text{x}10^5$
Year1	$y = 0.034\ln(x) + 0.0858$	1.5533	~ $5.55\text{x}10^{18}$
Year2	$y = 0.0076\ln(x) + 0.0269$	0.1997	~ $7.49\text{x}10^9$
Year3	$y = 0.0018\ln(x) + 0.0271$	0.0510	~ $5.84\text{x}10^5$
Year4	$y = 0.0027\ln(x) + 0.0078$	0.0389	~ $1.00\text{x}10^5$

Table 7-8 SOM data extrapolation to compare calculations count with calculations limit. The data show that only at Year3 and Year4 mission stages would the SOM be able to emulate the BSOM tour performance

Taskflag	Calculations/ Iteration	Calculations required to reach benchmark	Within limit?
Year0	4,000,000	$3.90\text{x}10^{12}$	×
Year1	416,103	$2.31\text{x}10^{21}$	×
Year2	131,145	$9.82\text{x}10^{14}$	×
Year3	68,155	$3.98\text{x}10^{10}$	✓
Year4	72,366	$7.28\text{x}10^9$	✓

Even though the data show that it is possible for the benchmark to be found by increasing the number of iterations (assuming the best fit equations still hold) the method is not useful for all taskflag cases and involves a heavy computational burden. Analysis of the standard deviation (σ) will reveal whether the SOM can find better performing tours than the BSOM.

To reduce the computational burden caused by the iteration count however, further analysis is required.

For the σ analysis it is desirable to set the number of iterations to a fixed value that is great enough to allow the SOM to find good performing tours but limited to avoid computational burden. As introduced previously, the calculation time is to be limited to ~10000 seconds and this could be used as a limiting factor in the iteration count. For Year3 and Year4 taskflags however, this time limit would result in a number of iterations much greater than necessary to achieve good results. An additional limiting metric is therefore introduced to expedite the σ analysis (by lowering the number of iterations) whilst still maintaining the integrity of the results. A differing performance from the 10240 iteration mean performance of 1 task/year (~0.0027 tasks/day) is chosen as a tolerable loss to expedite the analysis. The iteration count that allows this loss is found simply from:

$$(R^{tf}_{10240} - 0.0027) = m^{tf} \ln(x) + c^{tf} \tag{7.20}$$

where R^{tf}_{10240} is the mean tour performance at 10240 iterations for taskflag $tf = \{0,1,\ldots,4\}$, x is the required number of iterations and m^{tf} and c^{tf} are taken from the equations in Table 7-7 for each taskflag. The resulting iteration counts (rounded to the nearest 10 iterations) for each taskflag are given in Table 7-9. Using these iterations counts for each taskflag significantly reduces the calculation burden of the algorithm for only a 1 task/year loss compared with using 10240 iterations. The count for Year1 remains high due to the size of T_{target} and the number of available tasks in the taskflag. These iteration counts are used in the σ analysis that follows.

Table 7-9 Calculated iteration count for each taskflag used in the σ analysis.

Taskflag	Calculated Number of Iterations
Year0	6,870
Year1	10,030
Year2	5,880
Year3	1,540
Year4	1,510

7.7.3 Standard Deviation (σ)

From the analysis of the number of iterations it is clear that when σ=1 the SOM is unable to find tours performing better than the benchmark unless a very large number of iterations are used. This is because the size of $\boldsymbol{\sigma}$ determines the size of the solution-space that algorithm has to search though. If the solution-space is too large then the algorithm cannot find the 'good' tours due to all the 'bad' ones also present. Reducing the size of σ decreases the size of the solution-space making it less likely for 'bad' tours to be generated. σ however cannot be reduced indefinitely as there is a minimum boundary σ-value that causes the SOM to emulate the benchmark. This can be seen in Figure 7-14. The chart shows mean tour performance for varying σ for the Year1 taskflag with T_{target} = 40 days. σ is plotted on a

logarithmic scale whilst *R* is plotted on a linear scale. The data points (●) are the mean performance from twenty repetitions of the algorithm. The individual repetition data is not shown as the spread is indiscernible from the mean data point markers (●). Also shown is the benchmark performance for Year1 at 1.5533 tasks/day represented by a dashed line.

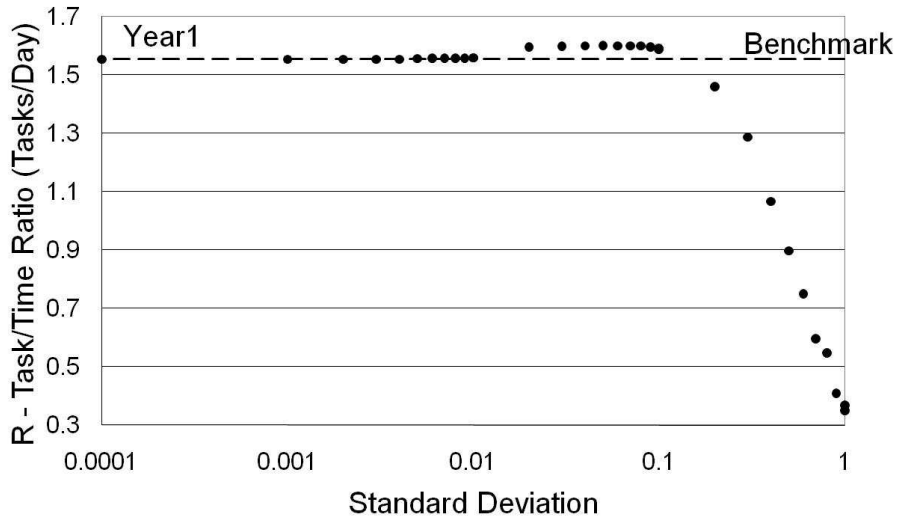

Figure 7-14 Mean tour performance vs. standard deviation for Year1 taskflag. The mean data points (•) and benchmark tour performance (dashed line) are shown. The chart shows that the SOM for the Year1 mission stage is able to perform better than the benchmark for some values of σ.

The data show that the performance does increase with decreasing σ surpassing the benchmark (between σ=0.1 and σ=0.2). Decreasing σ past this point increases the performance to a maximum (σ=0.07) before it falls again to the benchmark value. This shows that (with careful setting of σ) the SOM is able to find better performing tours than the BSOM. These results are seen for the other taskflag files in Figure 7-15. The data here is plotted on a log-log scale so that all the taskflags can be represented and although some detail is lost due to the scale the same trends as in Figure 7-14 are apparent with a few minor differences. The data points represent the mean performance from twenty repetitions of the algorithm whilst the data points (*) represent the maximum mean performance for that taskflag. Each taskflag benchmark is represented by a dashed line.

As seen in Figure 7-14 the taskflag data for all mission stages exhibits an increasing tour performance with decreasing σ, rising to a maximum and then falling to emulate their respective benchmark. The σ that gives the maximum observed mean tour performance decreases as the mission progresses. This is due to task availability at each stage of the mission. Due to the number of available tasks early on in the mission, the long duration tasks are very unlikely to be selected and so a relatively large σ can be maintained. Later on in the mission task availability is greatly reduced and so the likelihood of selecting a long duration task is greatly increased. Reducing σ lowers the chance of a long duration task being chosen. If the same σ is maintained throughout the mission then tour performance will suffer in the later stages. If σ is allowed to decrease however, maximum tour performance can be

maintained. Another point to note is that the benchmark SOM is emulated for values σ<0.004 for all mission stages.

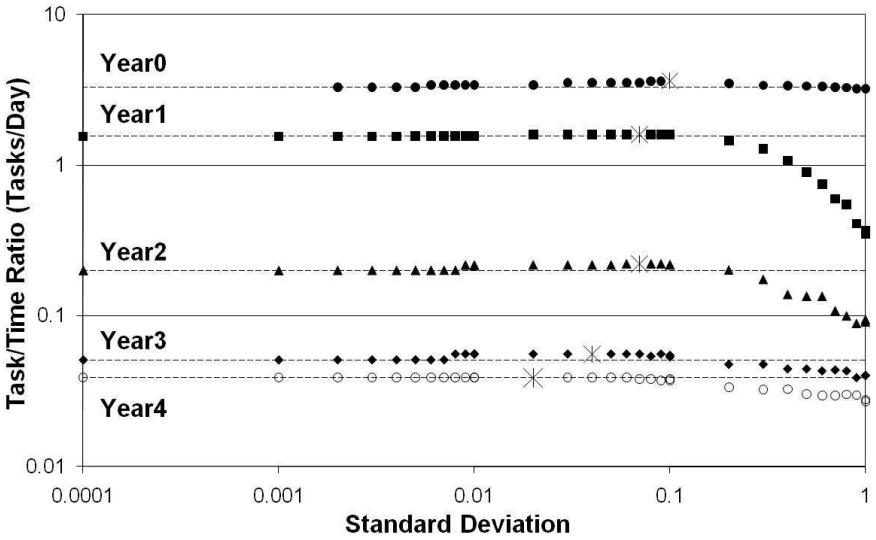

Figure 7-15 Mean tour performance vs. standard deviation for all taskflags. Mean data points are represented by markers whilst the respective benchmark performance is shown as a dashed line. The chart shows that for all mission stages there exists a σ value where the SOM can perform better then the BSOM given enough iterations.

Table 7-10 Comparison data between the SOM and BSOM tours

	Year0	Year1	Year2	Year3	Year4
Target tour time (days)	5	40	80	160	320
Chosen number of iterations	6,870	10,030	5,880	1,540	1,510
Optimal standard deviation (σ)	0.1	0.07	0.07	0.04	0.03
Mean SOM ratio (tasks/day)	3.6241	1.5996	0.2204	0.0557	0.038905
BSOM ratio (tasks/day)	3.2976	1.5533	0.1972	0.0509	0.038904
Performance increase (%) of SOM	9.9	3.0	11.8	9.4	0.003
Ratio Difference (tasks/day)	0.3265	0.0463	0.0232	0.0048	$1x10^{-6}$
Mean Task Duration (days)	0.3032	2.3858	5.3014	10.383	25.3966
Extra observation time (hr/day)	~2.4	~2.7	~3.0	~1.2	~0

A comparison of the data from the BSOM analysis and the SOM analysis can be found in Table 7-10. The table details the T_{target}, iteration count and σ-value used to generate the data, the performance ratios for the BP and SOM and the performance increase observed using the SOM. Also shown is the ratio difference, the mean task duration for each taskflag (from Table 7-5) and the calculated mean extra observation time, *O*, when using the SOM:

$$O = 24 \times \left(\bar{R}_{SOM} - \bar{R}_{BSOM} \right) \bar{D} \tag{7.21}$$

where $\bar{R}_{SOM}$ and $\bar{R}_{BSOM}$ are the maximum mean performance ratios for the SOM and BSOM respectively and $\bar{D}$ is the mean task duration (from the benchmark generated tours, Table 7-5). The comparison shows that significant time gains can be obtained when using the SOM algorithm (up to 3hrs/day for the Year2 taskflag). The Year4 results however show very poor performance gains compared to the other taskflag data.

As mentioned previously during the benchmark tour analysis the performance of a tour can vary greatly depending on the initial starting star of the tour. Using this observation, the Year4 taskflag tour is repeated but with varying starting stars. The results can be seen in Figure 7-16. The chart shows the variation of tour performance with standard deviation for the Year4 taskflag with four different starting stars. Each star is separated in ecliptic longitude by ~90°. The performance scale is linear whilst the standard deviation scale is logarithmic. The data points represent one run-through of the algorithm for each star with T_{target} = 360 days, 1510 iterations and σ=0.03. The respective benchmark performance is shown by a dashed line. The data show a variation in the benchmark performance between starting stars however there is only 5% difference between the highest and lowest performing benchmarks. For all the starting stars maximum performance is observed that betters the respective benchmark. The data emphasises that the SOM is able to find better performing tours than the BSOM but the starting star for the tour has a great affect on the tour's performance. This is expected since the difference between the ecliptic co-ordinates of the stars is not uniform (and even less so for the stars with planets). This result is important but not problematic for the SOM as the mean data always indicates that at the very least the SOM is able to emulate (but more than likely perform better than) the BSOM.

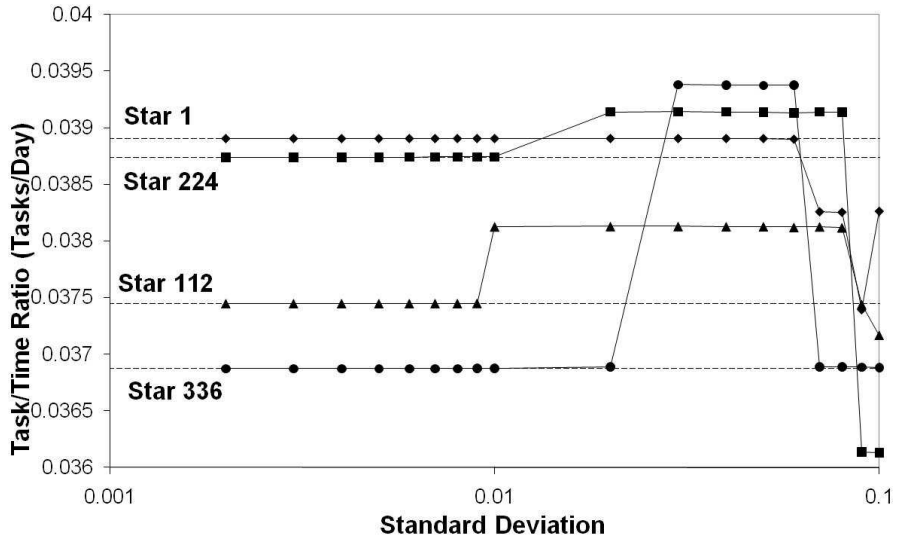

Figure 7-16 Year4 σ-analysis for different starting stars. Mean data points are represented by markers and the respective benchmark tour performance is given by a dashed line. The chart shows that the SOMs ability to find better performing tours than the BSOM is partially dependant of the starting star of the tour.

7.7.4 SOM Analysis Summary

The results from the analysis of the SOM indicate that it is able to consistently better the benchmark SOM provided that the correct tuning parameters are used. The first parameter, the number of iterations, needs to be large enough to allow a large area of the solution space to be searched whilst small enough to avoid exceeding the maximum amount of time available for the calculation. In reality this would probably be achieved by simply limiting the available calculation time but for this analysis a further metric was used. The second parameter, the standard deviation, needs to be large enough to allow a fair proportion of the solution space to be searched but not too large to avoid the selection of extremely long tasks. Additionally, the data show that the optimal values for these tuning parameters need to change as the mission progresses.

7.8 Science Operations Module and Benchmark Planner Comparison

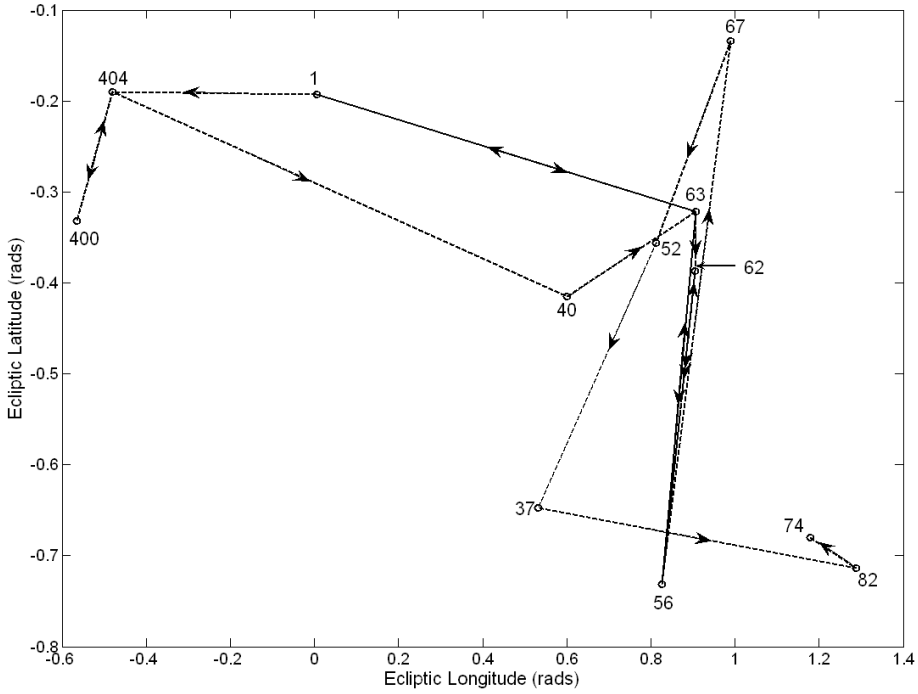

Figure 7-17 Sample optimised tour (Year2, T_{target} = 80 days, Startstar = 1, 5880 iterations, σ = 0.07). The numbers represent the star index from the catalogue and the arrows represent the direction of the star tour.

Using the observed optimal tuning data from the previous analysis (Table 7-10) a tour can be generated using the SOM for comparison purposes. The sample optimised tour is shown in Figure 7-17 where the star index numbers are shown next to each star. The stars are joined by a dotted line that blends to a solid line when retarget manoeuvres are repeated. The tour data can be seen in the left hand side of Table 7-11. Star index numbers are given along with the tasks selected to be performed on the star. Task '1' represents a detection task, task '2' the <7.2 μm spectroscopy task, task '3' the 8-9.2 μm spectroscopy task, task '4' the 9.2-10 μm spectroscopy tasks and task '6' the 13.2-17.2 μm spectroscopy task. The optimised tour has 18 tasks completed in 81.68 days giving a performance R = 0.22 tasks/day. As in the

previous benchmark tour example multiple retarget manoeuvres between detection tasks is observed (stars 56, 62 and 63).

Table 7-11 Comparison of tour data using identical initial conditions

Stochastic SOM Tour				**Benchmark SOM Tour**			
Star	Task	Task Duration (days)	Cumulative Duration (days)	Star	Task	Task Duration (days)	Cumulative Duration (days)
1	1	0	0	1	1	0	0
404	6	3.42	3.42	40	1	2.43	2.43
400	4	3.95	7.38	56	1	2.42	4.85
404	3	9.45	16.83	404	6	3.45	8.30
40	1	2.44	19.28	56	1	2.45	10.75
63	1	2.41	21.70	62	1	2.42	13.17
62	1	2.41	24.11	63	1	2.41	15.58
56	1	2.42	26.53	62	1	2.41	18.00
62	1	2.42	28.95	63	1	2.41	20.41
56	1	2.42	31.37	56	1	2.42	22.83
63	1	2.42	33.80	63	1	2.42	25.26
1	4	3.97	37.78	67	4	3.95	29.21
63	1	2.43	40.22	1	4	3.98	33.19
56	1	2.42	42.64	52	4	6.64	39.84
67	4	3.97	46.61	74	4	6.63	46.48
52	4	6.62	53.24	82	2	10.89	57.37
37	2	10.90	64.14	37	2	10.90	68.28
82	2	10.90	75.05	117	2	10.92	79.20
74	4	6.62	81.67	106	2	10.91	90.12
Tour Performance		0.220376	tasks/day	Tour Performance		0.199727	tasks/day

The right hand side of Table 7-11 shows the comparison benchmark tour generated using the same initial conditions. Comparing the SOM and the BSOM tours shows that many of the same tasks get performed in both tours but the order in which they are performed and the time where they are performed differs greatly. For this example both algorithms generate tours with 18 tasks however the SOM is able to complete those tasks in ~82 days whereas the BSOM requires ~90 days. This makes the SOM tour a better performing tour by the chosen comparison metric, *R* and frees up an extra 8 days (~10% of T_{target}) for additional tasking.

Figure 7-18 shows that the SOM tour is just as dynamic as the benchmark tour when utilising the full FOV. The chart shows how the difference between the ecliptic longitude of the anti-sun vector and the ecliptic longitude of the pointing direction of the formation at task completion varies throughout the tour. The optimised tour is represented by a solid line whilst the benchmark comparison tour is represented by a dashed line. Star index numbers are given at each task point. Longitude values close to the leading edge of the FOV are represented on

the negative *y*-axis whilst those longitudes close to the trailing edge of the FOV are represented on the positive *y*-axis. The chart shows both tours fully utilising the longitudinal range of the FOV without violating the FOV boundary conditions.

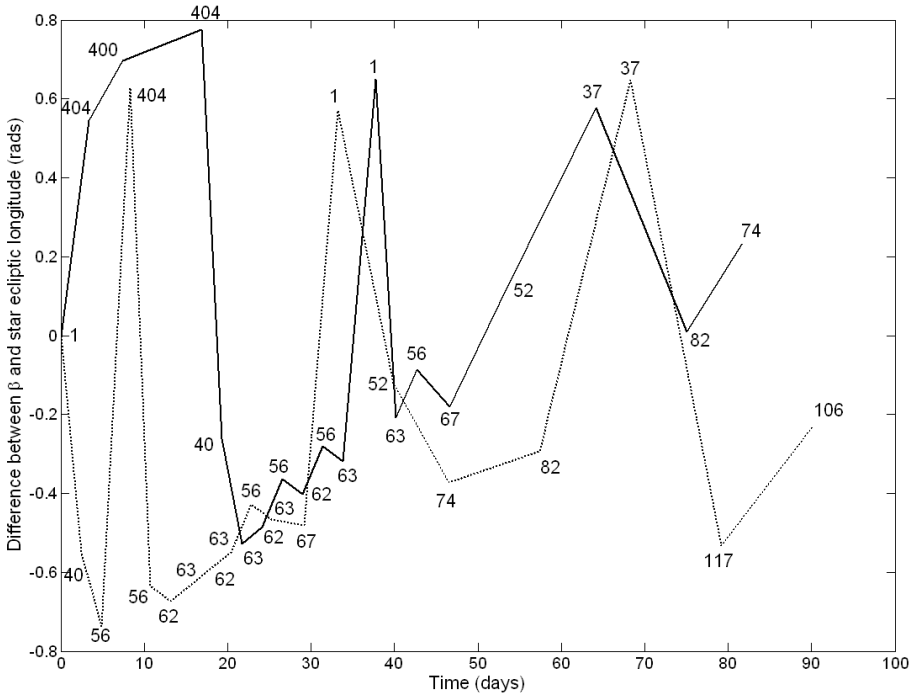

Figure 7-18 Difference between β and the ecliptic longitude of the formation pointing direction at task completion throughout the optimised (solid line) and benchmark (dashed line) comparison sample tours. Negative differences represent pointing directions ahead of β whilst positive differences represent pointing directions behind β.

7.9 Further Work

The SOM introduced in this chapter has been designed using a deterministic model of the DARWIN mission and the implementation of the manoeuvre planning constraints. There are a number of additional ways that the SOM and its implementation could be improved upon.

7.9.1 DARWIN Mission Constraints

Obtaining tours that reflect the DARWIN mission more realistically can be achieved by fully implementing all seven manoeuvre planning constraints introduced at the beginning of this chapter since, for the SOM developed, only one is fully implemented and three partially implemented. How these constraints are implemented affects the nature of the tour generated by the SOM.

Fully implementing these constraints would provide a much more accurate tour reflective of the DARWIN mission and could be achieved as follows:

- Detection task scheduling and the planetary orbital characteristics could be modelled by giving each of the tasks a time range in which they can be executed. This range would be calculated for each task to ensure the correct order of tasks performed (i.e.

detection before spectroscopy), appropriate task separation (i.e. between detection tasks) and correct planetary alignment (i.e. to avoid occultation or transits during observations). The time range would be referenced during the core algorithm (Section 7.3.2, the reduction of V_i to U_i) and only tasks whose time range included t_i would be selected for U_i. Should a time range pass and the task not completed then the time range would need to be updated for the next appropriate range.

- The calibration time constraint could be implemented by adding this duration to Equation (7.1).
- Planet detection analysis time could be implemented in two ways. In one implementation spectroscopy tasks are simply un-schedulable (through the use of a flag) until detection has been confirmed. This is simple to implement but will give less accurate schedules due to the unknown analysis duration. In another implementation the time range concept introduced above can be used with the time range offset to account for a standard or expected analysis time.
- A more accurate manoeuvre time can be modelled by including a manoeuvre planning algorithm (such as the Position Assignment Module (Chapter 8) or the Trajectory Design Module (Chapter 9)) but would require the adoption of different manoeuvre planning strategy to the one chosen.
- Finally the science task weighting can be modelled as suggested in the problem formulation (though not implemented in the analysis).

7.9.2 Comprehensive Data Analysis

The results obtained for SOM tours through the selection of the mission stage markers and the minimum tour duration, the tuning of the number of iterations and the standard deviation are specific to the problem. These selections are likely unrepresentative of the optimal parameters to adopt throughout the mission (as highlighted by the requirement for the standard deviation to decrease as the mission progresses). The results do show however that optimal values can be found for specific cases. A much more comprehensive data analysis study is required to refine the coarse data analysed here to see if further patterns emerge that can be used. Some further questions to answer include: how does the value of T_{target} affect tour performance and calculation performance throughout the mission; how does the optimal standard deviation change with the tour starting star; how does the algorithm perform when presented with a higher (or lower) than 10% planet probability?

7.9.3 Coding Optimisation

The performance of the SOM is affected to some extend by the capabilities of the computer system it is executed on. For autonomous operation on-board a spacecraft in the year 2030 computing capabilities are unlikely to be greater than the desktop PC used in this analysis (see section 6.6). Performance increases could be achieved by optimising the calculation processes within the coding of the algorithm. One such method involves the use of the memory resources.

As the SOM is running each tour generated is stored in the memory so that when the highest task/time ratio is found, the tour corresponding to that ratio can be retrieved and used. As a tour is being generated no reference is made to any of the previously generated tours and so with thousands of tours being generated there is a likelihood that tours may be duplicated. This can be avoided by using 'weighted edge memory'. Imagine a tour being generated using the SOM. At node n_i, there may be a completed tour, C, within the memory that has an identical edge sequence up to node n_i. The selection of the edge, $v_{i,j}$, will be weighted by the random number $|r|$. However an additional weighting could be included by cross-referencing with the $v_{i,j}$ found in C. The selection could be weighted against repeating the same selection as found in C for node n_i. As in simulated annealing (Chapter 4), at the beginning of the tour the weighting would have to be slight to allow for as large a solution space as possible. But as the tour progresses the weighting could slowly be increased so that near the end of the tour the algorithm is highly unlikely to select an edge that has been selected before in an identical tour.

Using this method would virtually eliminate repetition of entire tours whilst allowing similar tours to be generated. The number of unique tours would increase for the same iteration count as would the number of good performing tours. The only disadvantage to this method is that continual cross-referencing to the memory will increase the average time required to generate a tour. Any performance advantages gained by not repeating tours maybe lost by the performance loss though the use of additional memory resources. This method is one possible way of improving the performance of the SOM algorithm through coding optimisation but there are undoubtedly many more that the author is unaware of.

7.10 Chapter Summary

This chapter dealt with the design and implementation of the Science Operation Module (SOM), its output and its comparison with a benchmark operations scheduler (BSOM). A bespoke optimisation algorithm, based on ant colony optimisation, was designed to tackle the STA problem. This algorithm, the SOM, was described in detail in the text along with the elements of the STA problem that were implemented. A comparison algorithm, the BSOM, based on nearest neighbour selection, was also defined as an intuitive algorithm to compare performance.

The performance comparison between the SOM and the BSOM underlines the importance of the optimisation of science operations scheduling for the DARWIN mission as up to 3hrs per day science observation time is gained using the SOM rather than the BSOM. This performance however can only be obtained through careful selection of the optimisation parameters prior to starting the algorithm. Increasing the number of iterations of the algorithm logarithmically increases the performance of the optimised tour but linearly increases the time required to find a solution. Decreasing the standard deviation increases the performance of the optimised tour up to a maximum before decreasing again to emulate the BSOM tour's performance. This is complicated further with the affects of the optimisation parameters depending on when in the 5-year mission the SOM is being run. Fixing the value of the number of iterations for each mission year yielded optimal values for the standard deviation that gave the highest performing tours. These fixed number of iterations were

chosen firstly to limit the calculation time to less than 10^3 sec and secondly to reduce the number of iterations such that a mean performance decrease of only 1 task/year was observed. The data show a decrease in the value of the optimal standard deviation as the mission progresses. This indicated that to find the best tours required the standard deviation to decrease over the mission lifetime.

The limitations of the SOM were discussed in the final section of this chapter. Fully implementing all elements of the STA problem would make the SOM tour calculation more accurate but could incur large penalties on calculation time and thus affect the SOMs ability to find good, optimised tours. Methods to incorporate these additional elements were presented and limitations in the analysis of the SOM were discussed with comprehensive data analysis suggested to provide a greater insight into the effects the variables have on the performance of the algorithm. Finally, the limitations of the coding of the algorithm were examined and a method presented which reduces the chance of tour repetition and therefore would increase the number of unique tours found during each calculation.

The SOM represents one way of providing solutions to the science task assignment (STA) problem within the calculation time restricted environment imposed on the simulation. The SOM integrates well into the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) from sub-section 6.4. The following chapter describes the next optimisation module within the SepM-MPA data flow, the Position Assignment Module (PAM).

8. POSITION ASSIGNMENT MODULE

Once the target of the next observation has been chosen by the Science Operations Module (SOM) the next stage of the manoeuvre planning process is to plan the manoeuvre. This is achieved though a two stage optimisation procedure. In the first stage the post-manoeuvre spacecraft positions are found. This gives the position boundary conditions for the manoeuvre. In the second stage the trajectories for each spacecraft are found. The first stage is performed by the Position Assignment Module (PAM) whilst the second is performed by the Trajectory Design Module (TDM). Figure 8-1 shows the position of these modules within the Separate Modular Manoeuvre Planning Architecture (SepM-MPA). This chapter concerns the PAM only.

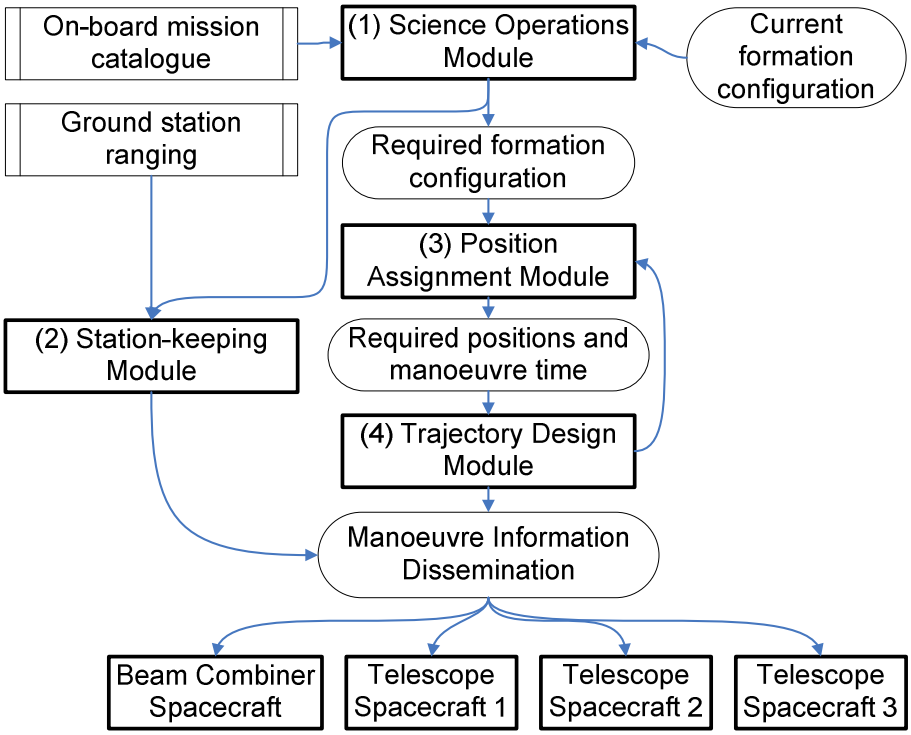

Figure 8-1 Separate Modular Manoeuvre Planning Architecture (reproduced from Figure 6-10)

As introduced in section 6.4.1 the goal of this module is to find the relative final spacecraft positions that, at the end of the manoeuvre, satisfy the science task formation configuration requirements whilst optimising for fuel minimisation, manoeuvre duration minimisation and fuel balancing.

8.1 Previous Contributions

A free-space position optimisation scheme was developed for separated spacecraft interferometry manoeuvres in the 2-dimensional planar case by Beard, McLain and Hadaegh (1998). Here the formation is viewed as a rigid body and the manoeuvre is a rotation about a point in space. A third party optimisation algorithm is used to find the position of this rotation

point so that the final spacecraft fuel states are equalised with as little fuel expended as possible. A bang-coast-bang acceleration profile is used to characterise the manoeuvre for fuel calculation purposes. Two algorithms are considered, a closed-loop fixed rotation point approach and an open-loop dynamic rotation point approach. The results show that the dynamic approach does not perform better than the fixed approach at fuel equalising due to the amount of fuel expended by spacecraft chasing the dynamic point of rotation.

This work is expanded to include 3-dimensional unconstrained rotations (Beard and Hadaegh, 1999) and 3-dimensional constrained rotations (Beard and McLain, 2000). In Beard and Hadaegh (1999) a fixed rotation point is used to rotate a 'rigid-body' formation to point at a new direction in space. For fuel calculation purposes however the spacecraft perform the manoeuvre unconstrained using straight-line trajectories with a bang-coast-bang thrust profile. The position of the fixed rotation point is optimised to consider fuel minimising / fuel balancing goals. Fuel balancing is demonstrated using a manoeuvre optimised by the given algorithm. In addition, time minimising / fuel minimising goals can also be included at the user's discretion. In Beard and McLain (2000) the exact same technique is demonstrated however for fuel calculation purposes the formation is constrained during the manoeuvre. This means the relative positions of the spacecraft within the formation remain fixed throughout the manoeuvre and the spacecraft follow an arcing trajectory in free space. Again the position of the fixed rotation point is optimised and fuel balancing is demonstrated using the technique.

All three of these papers by Beard et al. (1998, 1999 and 2000) demonstrate very similar and effective manoeuvre planning techniques for formations in free space. However they are constrained in a number of factors. In modelling the formation as a rigid body only retarget and rotation manoeuvres can be considered. There is no scope for the inclusion of the reconfiguration or resize manoeuvres that are essential for separated spacecraft interferometry. In addition, no analysis appears to have been published regarding the affects of varying the time/fuel minimising weighting factors or the fuel minimising/balancing weighting factors. Neither is any data available on the performance of the algorithm on sequences of manoeuvres.

More complex optimisation schemes need to be adopted in cases where the solution space is much greater than the 3-space of a rotation position vector. Maihle and Guzman (2004) use genetic algorithms and primer vector theory in a two stage optimisation process to minimise the ΔV for formation initialisation in an unperturbed two-body environment. With each spacecraft using a two-burn Lambert transfer the optimisation goal is to find the initial and final spacecraft positions and velocities that minimise the manoeuvre ΔV whilst satisfying the mission constraints. This is achieved by employing a genetic algorithm to search a 6 variable solution-space for optimal conditions. Primer vector theory is then employed to find the optimal location and number of manoeuvres required to reduce the delta-V even further.

Position assignment is also considered using mixed integer linear programming (Richards, et al., 2001 and Richards, et al., 2002). Here the spacecraft are assumed identical and the relative dynamics are governed by Hill's equations. A discretised target formation configuration is given and the optimisation routine is required to find the spacecraft

assignment within the configuration that minimises the fuel consumption of the reconfiguration manoeuvre. The problem and objectives are defined using a number of logical equality and inequality constraints so that the cost function can be defined as a linear function. The optimisation is then performed using third-party software.

A similar assignment optimisation method employs integer programming techniques (Tillerson, Inalhan and How, 2002). Again, with dynamics governed by Hill's equations, the optimisation algorithm is required to find the fuel optimal spacecraft assignment within a discretised formation configuration using identical spacecraft. The initial calculations are performed in a distributed manner where every spacecraft calculates its 'delta-V map' of fuel costs from its present position to each discretised target position. A 'co-ordinator' algorithm using integer programming techniques then uses every delta-V map to find a position assignment for every spacecraft that globally minimises fuel consumption. The authors also describe a way to incorporate fuel balancing into the algorithm.

These latter papers; Maihle and Guzman (2004), Richards, et al. (2002) and Tillerson, Inalhan and How (2002), all perform their optimisations based on a dynamic model unsuitable for the Darwin mission, namely in Earth orbit and manoeuvres are calculated towards known discretised positions. The dynamic environment at L2 allows for much less restrictive position constraints for spacecraft manoeuvres and therefore requires an optimisation technique that can handle such flexibility.

8.2 Position Assignment Module

The approach presented in this is inspired by Beard and Hadaegh (1999) but adapted to include the complexities of the DARWIN mission concept. Major differences of this approach to that published by Beard and Hadaegh (1999) are highlighted at the end of this section.

8.2.1 Model Definition

8.2.1.1 Spatial Geometry

Let F_0 be a co-ordinate frame with orthonormal basis vectors $\{\mathbf{i_0}, \mathbf{j_0}, \mathbf{k_0}\}$, where $\mathbf{i_0}$ lies in the ecliptic plane and points towards the J2000.0 vernal equinox, $\mathbf{j_0}$ lies in the ecliptic plane normal to $\mathbf{i_0}$ such that $\mathbf{j_0} = \mathbf{k_0} \times \mathbf{i}_0$ and $\mathbf{k_0}$ is normal to the ecliptic plane and in the same direction as the angular momentum vector of the orbit of the Earth around the Sun. F_0 is designated the inertial reference frame. Let F_S be a co-ordinate frame with orthonormal basis vectors $\{\mathbf{i_S}, \mathbf{j_S}, \mathbf{k_S}\}$, where $\mathbf{i_S}$ points towards the star S. F_S is a quaternion rotation of F_0 by an angle θ_S about an axis $\mathbf{z_S}$ where

$$\sin\theta_S = \frac{|\mathbf{z_S}|}{|\mathbf{i_S}||\mathbf{i_0}|} \tag{8.1}$$

$$\mathbf{z_S} = \mathbf{i_S} \times \mathbf{i_0} \tag{8.2}$$

F_S is designated the initial formation reference frame. Both F_0 and F_S are centred on the position of the BCS prior to the manoeuvre. Let F_F be a co-ordinate frame with orthonormal basis vectors $\{\mathbf{i_F}, \mathbf{j_F}, \mathbf{k_F}\}$, where $\mathbf{i_F}$ points towards the star F. F_F is a quaternion rotation of F_0 by an angle θ_F about an axis $\mathbf{z_F}$ (using Eqns. (8.1) and (8.2)) and substituting the subscript **S** for **F**) and a translation of F_0 by a vector $\mathbf{r_F}$. F_F is designated the final formation reference frame. This can be seen more clearly in Figure 8-2. Finally let $\mathbf{r_{0,i}}$ and $\mathbf{r_{f,i}}$ denote the initial and final position vectors of spacecraft *i* in F_0 respectively, where *i={BCS,TS1,TS2,TS3}.*

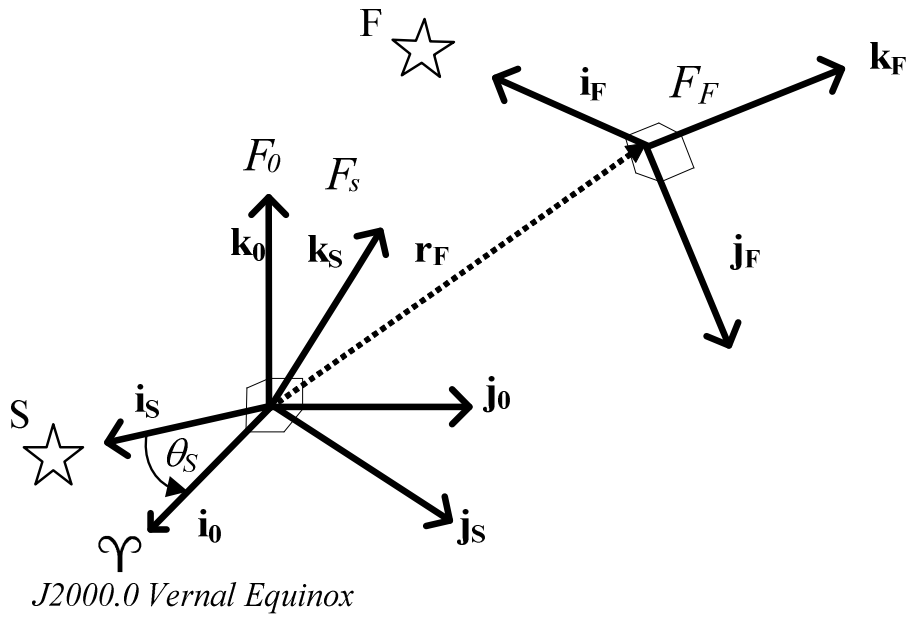

Figure 8-2 Spatial geometry of the PAM model

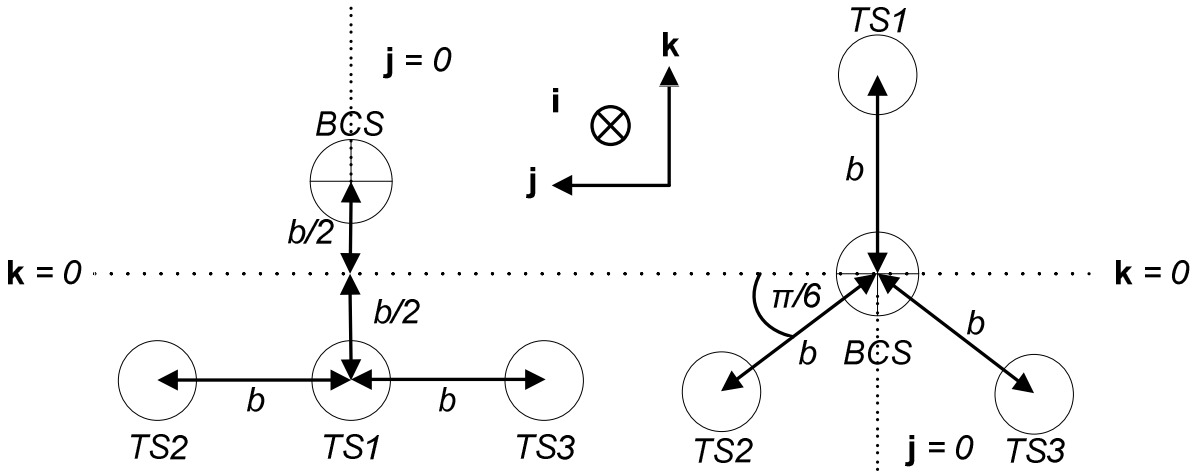

Figure 8-3 Geometry for the linTTN (left) and triTTN (right) with baseline *b*

8.2.1.2 Formation Geometry of the PAM Model

The formation planes are determined by the vectors ($\mathbf{j_S}$, $\mathbf{k_S}$) and ($\mathbf{j_F}$, $\mathbf{k_F}$) for the initial (Π_S) and final (Π_F) formation planes respectively. The spacecraft positions within these planes are shown in Figure 8-3 for the linTTN and triTTN. The plane is defined by the **k** and **j** vectors whilst the **i** vector points into the page completing the Cartesian set. For the linTTN the

formation takes a ‘T’ shape with each arm separated by baseline *b*. The centre of this configuration is set along the **k**-axis halfway between the BCS and TS1. The centre has been defined this way maximise the efficiency of the comparison benchmark algorithm that is used to compare against the PAM (see subsection 8.2.5). The benchmark utilises a rotation about the centre of the formation. If that centre was defined as the initial position of the BCS then the TS spacecraft would be required to perform very large manoeuvres whilst the BCS remained stationary. This would increase the fuel consumption and fuel differences between the spacecraft unduly with respect to the optimised PAM method. With the centre displaced halfway between the BCS and TS1 these increases are minimised allowing a less biased comparison to be made. For the triTTN the formation takes a ‘Y’ shape with each arm separated by baseline *b* and each TS separated by a 120º angle. The centre of this formation is the initial position of the BCS.

8.2.1.3 Model Assumptions

The following assumptions have been made to aide model simplification:

- The formation is in free space.
- Each spacecraft is modelled as a point mass
- Each spacecraft has mass that is time-invariant.
- The position of each spacecraft can be determined.

These assumptions provide a representation of the DARWIN spacecraft and mission and are good generalisations that aid the simplification of the PAM algorithm.

Though the formation is to follow a trajectory within the vicinity of the Sun-Earth L_2 libration point the gravity gradient between the spacecraft (i.e. at 100 m separation) is negligible (see section 6.3.2). Furthermore the reconfiguration manoeuvre time (~6000 sec) is much shorter than the destabilisation frequencies of libration point trajectories so the spacecraft will not noticeably drift apart during the manoeuvre due to the varying gravity gradient. Modelling the formation in a full gravity gradient model within the SepM-MPA is deemed unnecessary especially since the L_2 trajectory station-keeping planning will be performed by the Station-keeping Module (SKM).

Modelling the spacecraft as point masses removes the requirements for spacecraft attitude manoeuvres to be considered for reconfiguration manoeuvre planning. This allows the translational manoeuvre planning to be de-coupled from the attitude manoeuvre planning (which is spacecraft specific). It should be noted that the PAM can still plan formation retargeting manoeuvres (i.e. changes in the direction of the formation plane) as they rely on spacecraft translational manoeuvres, not attitude manoeuvres.

The spacecraft masses are 900-1100 kg and using FEEP thrusters the total fuel amount is likely to only be 5 kg for the whole mission (Karlsson, et al., 2004). Each manoeuvre will only use a fraction of that fuel and so to model the spacecraft with time-varying mass is deemed unnecessary.

8.2.1.4 Translational and Fuel Dynamics

From the assumptions given in subsection 8.2.1.3 the translational and fuel dynamics for each spacecraft are given by:

$$\begin{aligned} m_i\ddot{\mathbf{r}}_i &= \begin{cases} \alpha_i \mathbf{u_i}; & mf_i(t) > 0 \\ 0; & otherwise \end{cases} \\ \dot{mf_i} &= \begin{cases} -\gamma\alpha_i; & mf_i(t) > 0, |\mathbf{u_i}| \neq 0 \\ 0; & otherwise \end{cases} \end{aligned} \tag{8.3}$$

where m_i is the mass (in kg) of the ith spacecraft, $\ddot{\mathbf{r}}_\mathbf{i}$ is the acceleration vector in F_0 (ms^{-1}), α_i is the thrust saturation limit (N), $\mathbf{u_i}$ is a unit force vector (N), $\dot{mf_i}$ is the fuel consumption rate (kgs^{-1}), mf_i is the amount of fuel (kg) and γ is a proportionality constant

$$\gamma = 1/I_{sp}g \tag{8.4}$$

where I_{sp} is the specific impulse of the thruster and g is the acceleration due to gravity at sea level.

8.2.2 Position Assignment Module Optimisation

The goal of the PAM is to find the post-manoeuvre spacecraft positions that satisfy the interferometry requirements (formation pointing direction, configuration and baseline size) whilst optimising for time, fuel usage and fuel balancing. In the model the time optimal manoeuvre is a straight-line trajectory from initial to final positions using a bang-bang thrust profile. Each spacecraft however will have a different distance to travel and so will complete their individual manoeuvres at different times. The time it takes the formation to complete the manoeuvre is constrained by the spacecraft that takes the longest to complete its individual manoeuvre. Fuel savings can be achieved for the other spacecraft by adopting a bang-coast-bang thrust profile so that all spacecraft in the formation complete their individual manoeuvres at the same time. This is the concept adopted for the PAM.

By adopting a bang-coast-bang thrust profile time optimisation of the formation manoeuvre is achieved for any combination of initial and final spacecraft positions. In addition, fuel savings are also accomplished. Addressing the fuel balancing goal can be found by optimising the final spacecraft positions. This can be achieved through the minimisation of a cost function

$$J = \min_X \left\{ \sum_i^n \left(mf_i(t_0) - mf_i(t_f) \right) + \mu_{PAM} \sum_i^n \sum_{j:j\neq i}^n \left| mf_i(t_f) - mf_j(t_f) \right| \right\} \tag{8.5}$$

where $i,j=\{BCS,TS1,TS2,TS3\}:i\neq j$ are the spacecraft identifiers, $mf_i(t_0)$ is the initial fuel amount of spacecraft i, $mf_i(t_f)$ is the final fuel amount and X is an independent vector. The

first part of the objective function represents the total amount of fuel used by the formation during the manoeuvre. The second part is the sum of the post-manoeuvre fuel differences across the fleet. The weighting of the second term is governed by the parameter μ_{PAM}. As $\mu_{PAM} \rightarrow 0$ the objective function minimises to generate a fuel minimising manoeuvre. As $\mu_{PAM} \rightarrow \infty$ fuel balancing is achieved. It should be noted that the fuel balancing part of the cost function aims to reduce the sum of the fuel differences for the formation at the end of the manoeuvre not encourage balanced fuel use amongst the spacecraft during the manoeuvre. To find the cost J of a manoeuvre it is necessary to find $mf_i(t_f)$ in terms of input parameter X.

Whereas in Beard and Hadaegh (1999) the independent variable, X, was the position of a rotation vector, this prohibits the optimisation of more complex reconfiguration manoeuvres. The geometric constraints of the linTTN and the triTTN allow greater optimisation flexibility by adding only one extra term to X. The four terms X are defined as:

- 3 unconstrained position co-ordinates, $\mathbf{r_F}$, representing the final position of the BCS in the inertial reference frame F_0

- 1 unconstrained angle, θ_1, representing the relative angular position slot of TS1 around the BCS in F_F.

X is therefore defined as:

$$X = \left\{ \mathbf{r}_{\mathbf{F},\mathbf{x}}, \mathbf{r}_{\mathbf{F},\mathbf{y}}, \mathbf{r}_{\mathbf{F},\mathbf{z}}, \theta_1 \right\} \tag{8.6}$$

X can be seen more clearly in Figure 8-4. The vector $\mathbf{r_F}$ allows the formation to be translated anywhere in space from the initial formation position. The reference axis is an arbitrary vector in F_F and lies in the formation plane, Π_F. TS1 can be placed anywhere in the formation plane using the angle θ_1 and the required baseline b. TS2 and TS3 then link directly to the required relative formation configuration geometry from Figure 8-3. Since the telescope spacecraft are identical, their positions within the formation plane are interchangeable. For each X there are six possible TS position combinations that can be assessed. The combination that returns the lowest cost in Equation (8.5) is assigned to that X. This TS assignment provides additional flexibility for the optimisation algorithm.

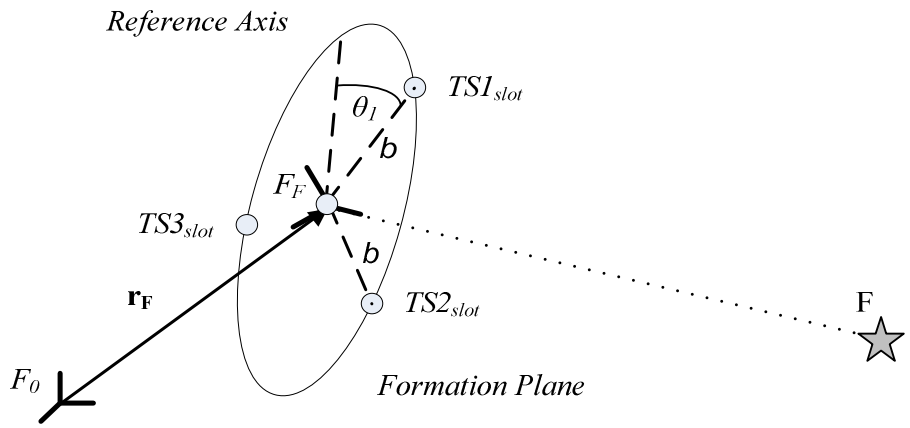

Figure 8-4 Geometry of the final spacecraft position calculation (triTTN shown). The independent variable, *X*, decides the position of the BCS through $\mathbf{r_F}$ and the position slots of the TSs through angle θ_1.

8.2.3 Position Assignment Module Algorithm

The optimisation procedure is summarised in the flow chart in Figure 8-5. The optimisation routine starts with an initial estimate for the independent vector, X_0. Using Figure 8-3 and Figure 8-4 as a guide the final spacecraft positions are calculated. These positions, $\mathbf{r_{f,i}}$, are input to the assignment routine. Within the assignment routine, each set of TS combinations are tested and a cost calculated. The assignment routine returns the lowest cost of the six possible TS iterations. This cost (and related independent variable *X*) is stored in the optimisation algorithm's database. If the stopping conditions for the optimisation algorithm have not been met then a new *X* is generated and the process iterates. When the stopping conditions are met the optimisation algorithm terminates and outputs the variable, *X*, that minimises the cost, *J*. This *X* can then be used to calculate the final positions of the spacecraft that satisfy the observation requirements and the optimisation goals.

8.2.4 Calculating the Cost

For any set of spacecraft positions the minimal time trajectory is always implemented through the bang-bang thrust profile. Then, fuel savings are made by only implementing the bang-bang thrust profile on the spacecraft with the longest trajectory, with bang-coast-bang implemented on the other spacecraft. The cost function, Equation (8.5), gives the cost of any manoeuvre with respect to the initial and final fuel masses of each spacecraft in the formation. This sub-section describes how the final fuel masses, $mf_i(t_f)$, are calculated from the initial ($\mathbf{r_{0,i}}$) and final ($\mathbf{r_{f,i}}$) spacecraft positions. The calculation is a modification of the method in Beard and Hadaegh (1999).

The trajectory length for each spacecraft is simply

$$\mathbf{r_{traj,i}} = \mathbf{r_{f,i}} - \mathbf{r_{0,i}} \tag{8.7}$$

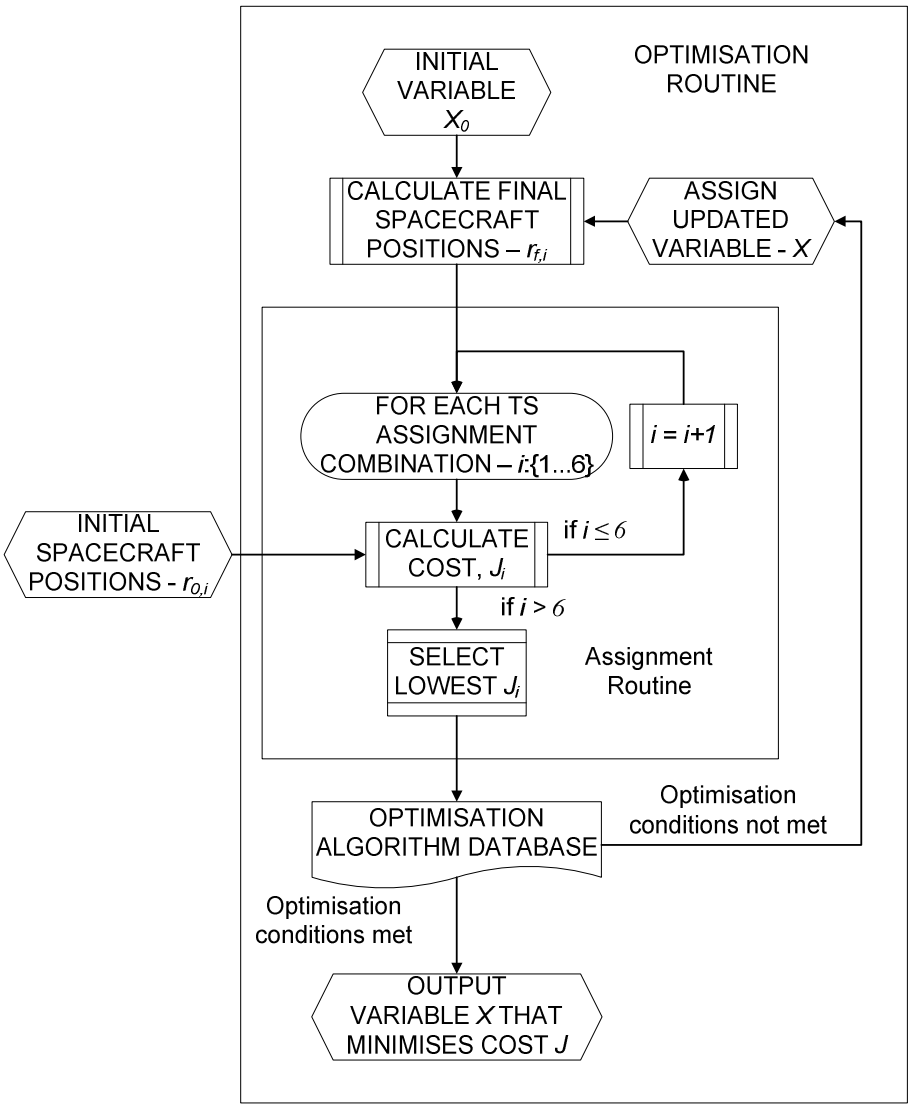

Figure 8-5 Optimisation routine for the Position Assignment Module

The maximum acceleration of each spacecraft is α/m so a manoeuvre time using a continuous thrust profile can be found by integrating twice along the direction of travel

$$\ddot{r} = \frac{\alpha}{m}; \quad \dot{r} = \frac{\alpha}{m}t; \quad r = \frac{\alpha}{m}\frac{t^2}{2} \tag{8.8}$$

with

$$\dot{r}_0 = r_0 = 0 \tag{8.9}$$

where r is the distance the spacecraft travels in time t. Equation (8.8) can be rearranged to find the duration, t, of any manoeuvre using a continuous thrust over a distance, r.

$$t = \sqrt{2r\frac{m}{\alpha}} \tag{8.10}$$

For a bang-bang thrust profile the continuous thrust acts in one direction for $\mathbf{r}_{\mathbf{traj,i}}/2$ and the opposite direction for $\mathbf{r}_{\mathbf{traj,i}}/2$. Therefore the time, t_i, to travel the complete trajectory, $\left|\mathbf{r}_{\mathbf{traj,i}}\right|$, with a bang-bang thrust profile is:

$$t_i = 2\sqrt{\left|\mathbf{r}_{\mathbf{traj,i}}\right|\frac{m_i}{\alpha_i}} \tag{8.11}$$

Let spacecraft β be the spacecraft that has the longest manoeuvre time, t_β, using a bang-bang thrust profile. β is found from the condition

$$\beta = \arg\max_{1\le i\le 4} t_i \tag{8.12}$$

The formation manoeuvre duration, t_f, cannot be less than the manoeuvre duration of spacecraft β using a bang-bang thrust profile, therefore

$$t_f = t_\beta \tag{8.13}$$

To include the provision for a bang-coast-bang thrust profile for the other spacecraft we introduce the time intervals T_1, T_2 and T_3 as

$$\begin{aligned} T_1 &= [0, t_{\omega i}] \\ T_2 &= [t_{\omega i}, (t_f - t_{\omega i})] \\ T_3 &= [(t_f - t_{\omega i}), t_f] \end{aligned} \tag{8.14}$$

where $t_{\omega i}$ is the thrust pulse width for spacecraft i. Let $y_i(t)$ be the distance spacecraft i has travelled along $\mathbf{r}_{\mathbf{traj,i}}$ in time t. The manoeuvre profile for spacecraft i is therefore

$$\begin{array}{lccc} & t \in T_1 & t \in T_2 & t \in T_3 \\ \ddot{y}_i(t) = & \dfrac{\alpha_i}{m_i}; & 0; & -\dfrac{\alpha_i}{m_i}; \\ \dot{y}_i(t) = & \dfrac{\alpha_i}{m_i}t; & \dfrac{\alpha_i}{m_i}t_{\omega i}; & \dfrac{\alpha_i}{m_i}(t_f - t); \\ y_i(t) = & \dfrac{1}{2}\dfrac{\alpha_i}{m_i}t^2; & \dfrac{\alpha_i}{m_i}\left[t_{\omega i}t - \dfrac{t_{\omega i}{}^2}{2}\right]; & \dfrac{\alpha_i}{m_i}\left[t_{\omega i}(t_f - t_{\omega i}) - \dfrac{(t_f - t)^2}{2}\right]; \end{array} \tag{8.15}$$

where T_1, T_2 and T_3 represent the durations of the bang-coast-bang thrust profile respectively. The top row of Equation (8.15) shows the spacecraft acceleration during these time periods. The middle row shows the spacecraft velocity and the bottom row shows the spacecraft position. A schematic showing the thrust profiles and timings is given in Figure 8-6.

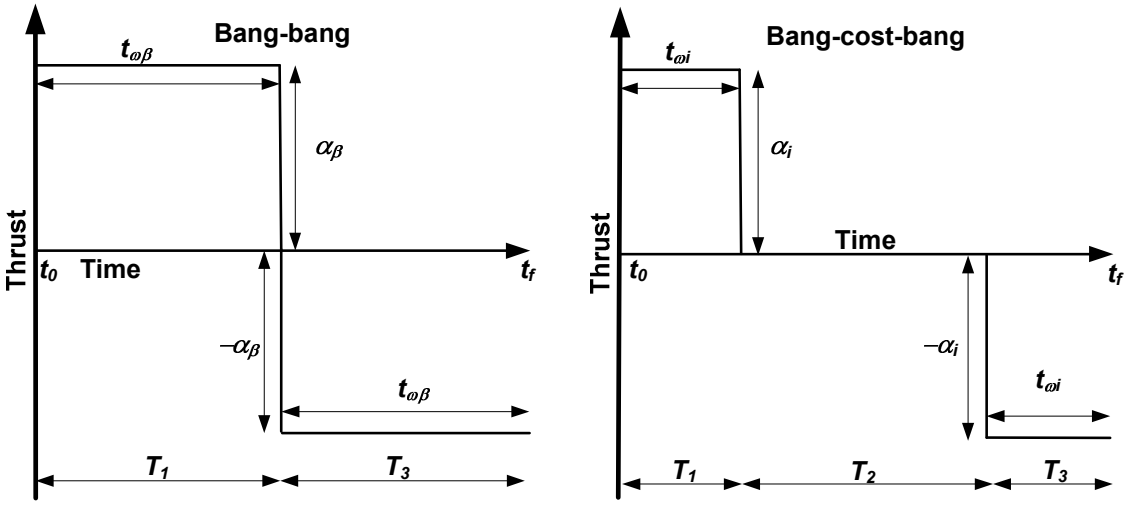

Figure 8-6 Schematic of the bang-bang (left) and bang-coast-bang (right) thrust profiles

When $t = t_f$ the trajectory length of spacecraft *i* is

$$y_i\left(t_f\right) = \frac{\alpha_i}{m_i} t_{\omega i}\left(t_f - t_{\omega i}\right) \tag{8.16}$$

Rearranging this becomes

$$t_{\omega i}{}^2 - t_f t_{\omega i} - \frac{m_i}{\alpha_i} y_i\left(t_f\right) = 0 \tag{8.17}$$

and solving for $t_{\omega i}$ gives[9]

$$t_{\omega i} = \frac{t_f}{2} - \sqrt{\frac{t_f{}^2}{4} - \frac{m_i}{\alpha_i} y_i\left(t_f\right)} \tag{8.18}$$

This equation satisfies the condition

$$t_\beta = 2t_{\omega\beta} = t_f \tag{8.19}$$

With two periods of thrust for each manoeuvre the fuel consumed by spacecraft *i* is

[9] here the negative part of the root discriminate is used as $t_{\omega i} \le t_f / 2$

$$F_i = -2t_{\omega i} m\dot{f} = 2t_{\omega i}\gamma\alpha_i \tag{8.20}$$

So the final fuel levels for each spacecraft are

$$mf_i\left(t_f\right) = mf_i\left(t_0\right) - F_i \tag{8.21}$$

8.2.5 Comparison Benchmark Algorithm

For the effectiveness of the PAM to be demonstrated it needs to be compared to another manoeuvre planning method, dubbed the benchmark position assignment module (BPAM). The BPAM is based on the most intuitive way to plan the manoeuvres. The retarget is a rigid body quaternion rotation in 3-D space by an angle θ_B about an axis $\mathbf{z_B}$ where

$$\sin\theta_B = \frac{\left|\mathbf{z_B}\right|}{\left|\mathbf{i_S}\right|\left|\mathbf{i_F}\right|} \tag{8.22}$$

$$\mathbf{z_B} = \mathbf{i_F} \times \mathbf{i_S} \tag{8.23}$$

This moves the formation plane from F_S to F_F. Resize and reconfiguration can then be performed within F_F using the geometry in Figure 8-3. Comparison between the PAM and BPAM is given in Table 8-1. The additional flexibility of the BCS translation, rotation within the formation plane and the TS assignment gives the PAM the ability to plan manoeuvres with greater optimisation of fuel consumption and fuel balancing than the BPAM.

Table 8-1 Comparison between the Benchmark Planner and the PAM

	BPAM	**PAM**
Retarget	✓	✓
Resize	✓	✓
Reconfigure	✓	✓
BCS translation	×	✓
Rotation in formation plane	×	✓
TS assignment	×	✓

8.2.6 Differences of the Position Assignment Module to Beard and Hadaegh's Method

Although the paper by Beard and Hadaegh (1999), hereafter abbreviated to BH, has been the inspiration for the development of the PAM, there are a number of significant differences that have been made to specifically accommodate the constraints of the DARWIN mission and improve the optimisation performance.

8.2.6.1 Cost Function and Optimisation Parameter

The cost function chosen for the PAM is given in Equation (8.5) but repeated here for clarity

$$J_{PAM} = \min_{X}\left\{\sum_{i}^{n}\left(mf_i(t_0) - mf_i(t_f)\right) + \mu_{PAM}\sum_{i}^{n}\sum_{j:j\neq i}^{n}\left|mf_i(t_f) - mf_j(t_f)\right|\right\} \tag{8.24}$$

The first part of the cost function is the total amount of fuel used to perform the manoeuvre whilst the second part is the sum of the spacecraft fuel differences. The cost function for BH however is slightly different

$$J_{BH} = \min_{R}\left\{\sum_{i=1}^{N}\left(mf_i(t_0) - mf_i(t_f)\right)^2 + \mu_{PAM}\sum_{i=1}^{N}\frac{mf_i(t_f)}{\sum_{j=1}^{N} mf_j(t_f)}\log\frac{mf_i(t_f)}{\sum_{j=1}^{N} mf_j(t_f)}\right\} \tag{8.25}$$

Here the first part of the cost function is the sum of the squares of each spacecraft's fuel consumption for the manoeuvre. The second part is described as "*the negative entropy of a probability distribution*" (Cover and Thomas, 1991). This is minimised for a uniform distribution, i.e. when $f_i(t_f) = f_j(t_f)$ for all $i, j \in \{1,\ldots,N\}$.

In the first part of Equation (8.25), each spacecraft's fuel consumption is squared and then summed to give a cost. This is the fuel minimising part of the cost function (i.e. when $\mu_{PAM}=0$). When the individual fuel consumption is squared this means outlier fuel consumption values become overly dominant in determining the cost for the manoeuvre. To minimise this function the outliers will become suppressed and the spacecraft fuel consumption will be more balanced. The first part of Equation (8.25) is therefore unable to find a truly fuel minimal manoeuvre. For this reason the squared term is dropped for the PAM cost function as in Equation (8.24). The second part of Equation (8.25) is a complex function based on a probability distribution and returns a dimensionless value. This is the fuel balancing part of the cost function (i.e. when $\mu_{PAM}\neq 0$). With this function it is difficult to assess how well the function actually balances the fuel distribution without looking at another (dimensionalised) metric (i.e. the second part of Equation (8.24)). In simulation the fuel balancing part of Equation (8.24) performs better than the fuel balancing part of Equation (8.25) when using the dimensionalised metric, contrary to that experienced in Beard and McLain (2001)[10]. The dimensionalised metric is chosen for the PAM as it provides a 'real' fuel balancing term for comparison purposes.

The optimisation parameters are also different for the PAM and BH. The PAM optimisation parameter, *X*, is a 4-space vector describing the position vector of the BCS and

[10] Though in Beard and McLain (2001) the manoeuvres being optimised are more akin to the BPAM method than the PAM method

the rotation of the formation plane. The BH optimisation parameter, *R*, is a 3-space vector describing the position of the rotation origin for the rotation vector. The additional dimension in *X* allows the optimisation routine to perform a much more flexible calculation but at the cost of increased calculation time.

8.2.6.2 Manoeuvre Time Optimisation

In the PAM the formation manoeuvre time is governed by the spacecraft with the longest manoeuvre time using a bang-bang thrust profile. This allows the use of a bang-coast-bang thrust profile for the other spacecraft in the formation, thus saving fuel. In the BH however, there is an additional minimal-time / minimal-fuel trade-off implemented. This factor allows the formation manoeuvre time to be increased by decreasing the thrust output of the spacecraft making further fuel savings possible. Although this makes the BH optimisation more flexible the minimum time/fuel trade-off has been dropped for the PAM in favour of always calculating the time minimal aspect of the manoeuvre. This makes the PAM slightly less complex and aids the optimisation goal of optimising observation time by minimising manoeuvre time.

8.3 Analysis

8.3.1 Analysis Setup

For all the manoeuvres generated in this section the following initial parameters were used:

$$\begin{aligned} i &= \{BCS, TS1, TS2, TS3\} \\ m_i &= \{1100, 900, 900, 900\}\, kg \\ \alpha_i &= \left\{6\times10^{-3}, 6\times10^{-3}, 6\times10^{-3}, 6\times10^{-3}\right\} N \\ I_{sp} &= 3300 s \end{aligned}$$

The mass data is representative of the estimated masses of each spacecraft type in the DARWIN mission (Karlsson, et al., 2004) and the thruster saturation data and thruster specific impulse describe a thruster configuration using the RIT-10 FEEP thruster (D'Arcio, 2005). The PAM and BPAM routines are written and executed within MATLAB® and the PAM is optimised using the Mesh Adaptive Direct Search (MADS) algorithm found within MATLAB's Genetic Algorithm and Direct Search Toolbox (GADS). The metrics used in the comparison of the PAM and BPAM planners are the total fuel used, Equation (8.26), the sum of the fuel differences, Equation (8.27), and the manoeuvre duration.

$$\sum_i \left(mf_i(t_0) - mf_i(t_f) \right) \tag{8.26}$$

$$\sum_i^n \sum_{j:j\neq i}^n \left| mf_i(t_f) - mf_j(t_f) \right| \tag{8.27}$$

8.3.2 Example Manoeuvre

An example manoeuvre showing the benchmark planner and the PAM can be seen in Figure 8-7 and Figure 8-8. For simplicity the manoeuvre involves no retarget (i.e. the formation remains pointing in the same direction) but resizes from a baseline of 50m to a baseline of 75m maintaining the triTTN configuration. The initial fuel amount for each spacecraft is $mf_i(t_0) = \{5.0, 4.7, 5.0, 5.0\}\, kg$. The left-hand chart shows the BPAM manoeuvre and the right-hand chart shows the PAM manoeuvre. Initial positions are given by open circles (○) and final positions by closed circles (●). The trajectories linking the positions are given as solid lines. The BPAM manoeuvre in Figure 8-7 and Figure 8-8 is the same. In the BPAM manoeuvre the BCS remains stationary whilst the TSs move radially away from the BCS until the baseline requirements are met. For μ_{PAM} =0 (Figure 8-7) the PAM has found the fuel-minimal positions by shifting the whole formation slightly in the negative *y* and *z* directions. For μ_{PAM} =10^3 (Figure 8-8) the fuel balancing positions are found by shifting the whole formation in the negative **z** direction. This fuel balancing manoeuvre eliminates the distance TS1 has to travel since TS1 is the spacecraft with the least amount of fuel at the beginning of the manoeuvre. Since TS1 uses no fuel and the other spacecraft do, fuel differences across the fleet are reduced.

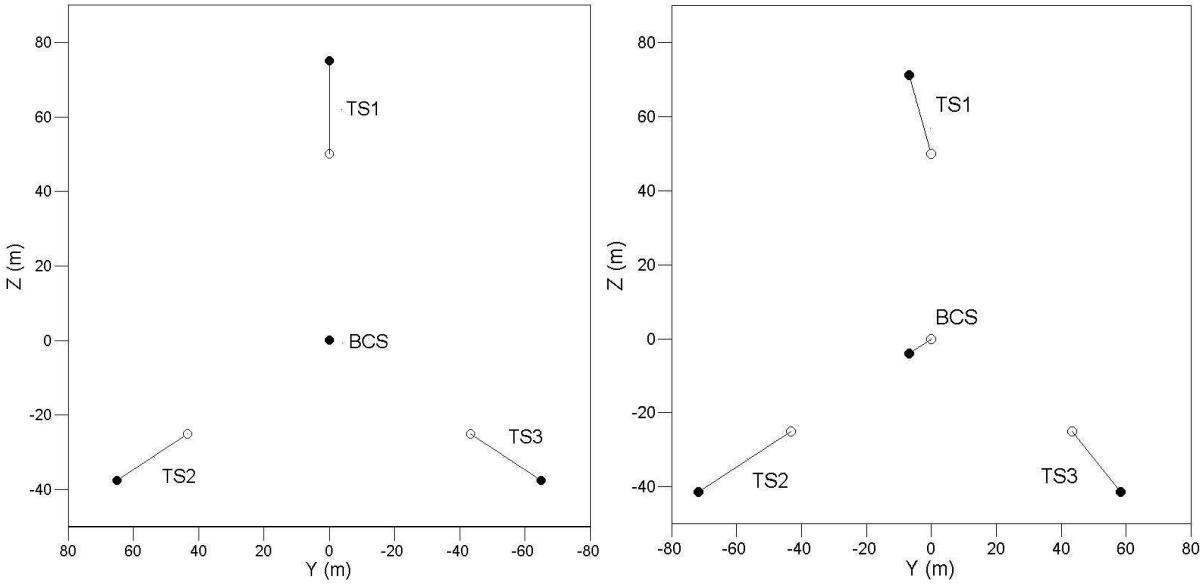

Figure 8-7 Example Manoeuvre: BPAM (left) and PAM with μ_{PAM} =0 (right). The diagram shows how the PAM displaces the BCS from its original position to minimise the fuel consumption of the formation.

Table 8-2 shows the manoeuvre data for the example manoeuvre. The numbers show agreement with the aims of the PAM. When μ_{PAM} =0 the PAM is able to plan a manoeuvre using less fuel than the BPAM with an increased manoeuvre duration of only a few minutes. When μ_{PAM} =10^3 the PAM is able to plan a manoeuvre that reduces the fuel difference amongst the fleet members. This can only be achieved however at the cost of increased fuel consumption and manoeuvre duration. The level to which the PAM achieves these goals equates to an ~23% decrease in fuel consumption for a 15% increase in manoeuvre time for

μ_{PAM} =0 and an ~1 % decrease in the sum of the fuel differences relative to the initial sum of the fuel differences for an 84% increase in manoeuvre time for μ_{PAM}=10^3. The manoeuvre shown here is simplified for visualisation purposes, the more complex manoeuvres required by Darwin are shown in the next sub-section.

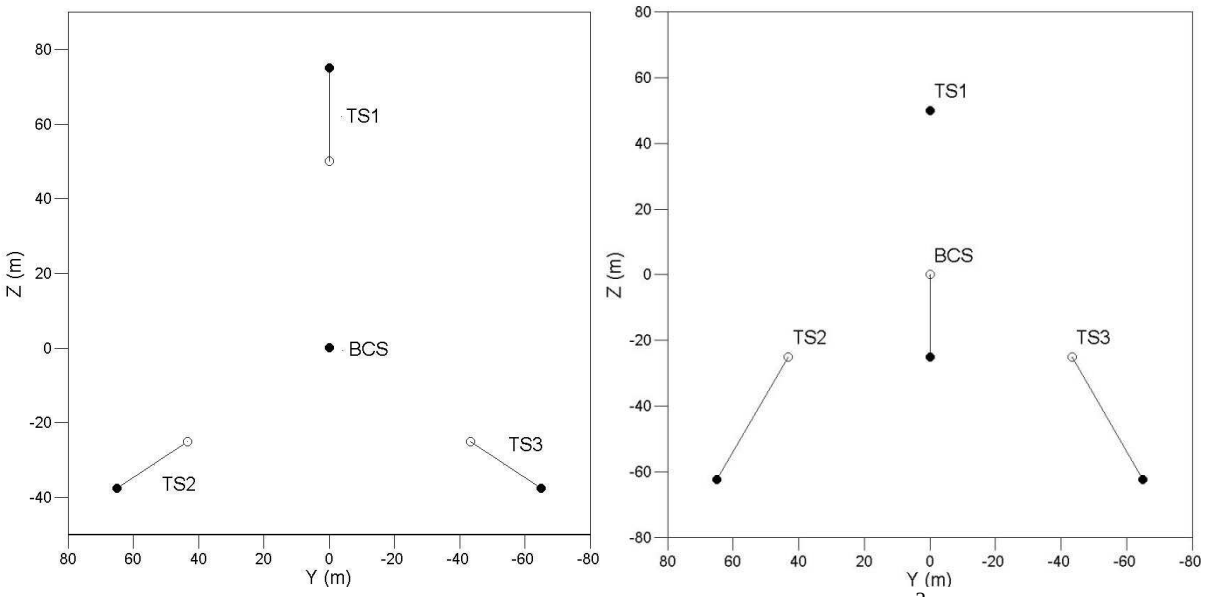

Figure 8-8 Example Manoeuvre: BPAM (left) and PAM with μ_{PAM} =10^3 (right). The diagram shows how the PAM displaces the BCS so that TS1 (the spacecraft with least amount of fuel) does not manoeuvre at all. This has the effect of improving fuel balancing across the formation.

Table 8-2 Manoeuvre data for the example manoeuvre

	Total Fuel Used (kg)	**Sum of Fuel Differences (kg)**	**Manoeuvre Duration (s)**
BPAM	0.005745	1.811489	3873
PAM μ_{PAM} =0	0.004427	1.806094	4435
PAM μ_{PAM} =10^3	0.011633	1.785409	7112

8.3.3 DARWIN Manoeuvre Analysis

The manoeuvre analysis of the PAM involves examining a set of 16 possible DARWIN manoeuvres. The initial and final target stars remain the same so that each manoeuvre can be assessed independently from the magnitude of the retarget. The initial star's unit vector $\mathbf{S} = [1, 0, 0]$ and final star's unit vector $\mathbf{F} = \left[\sqrt{1/3}, \sqrt{1/3}, \sqrt{1/3}\right]$ means the manoeuvres involve an ~55° retarget in 3-D space. The manoeuvres use combinations of configurations, {triTTN, linTTN}, and baseline, {50m, 75m} as the initial and final formation configurations. For example, a manoeuvre designated tri50-lin75 describes a formation in the triTTN configuration with a baseline of 50m and pointing at **S** performing a manoeuvre so that the final configuration is the linTTN with a baseline of 75m pointing at **F**.

Comparison of the PAM manoeuvres with BPAM optimised manoeuvres and differing values of μ_{PAM} reveal a number of trends. Figure 8-9 shows the total fuel used for each

manoeuvre using the benchmark planner and the two PAM manoeuvres, μ_{PAM} =0 and $\mu_{PAM} = 10^3$. Comparing the μ_{PAM} =0 case to the benchmark shows that in all manoeuvres the PAM finds more fuel efficient spacecraft positions with varying fuel reductions from 4%-14% with an average of 6.5%. The PAM: μ_{PAM} =0 case appears more effective for manoeuvres involving a formation reconfiguration (i.e. triTTN to linTTN or vice versa) because the nature of the BPAM means it performs better for more symmetrical manoeuvres (i.e. triTTN-triTTN, linTTN-linTTN). Looking at the total fuel used for the fuel balancing PAM: μ_{PAM} =10^3 case shows a great increase in fuel consumption to achieve the fuel balancing requirements (on average 70% increase over the BPAM).

Comparison of the sum of the fuel differences is shown in Figure 8-10. The initial value is 1.8 kg. As can be seen both the BPAM and PAM: μ_{PAM} =0 cases show an increase in the fuel differences for all manoeuvres. This is due to the nature of the formation configurations which will tend to limit the movement of the BCS compared to the other spacecraft in the formation. When μ_{PAM} =10^3 the flexibility of the PAM can force a BCS manoeuvre in order to minimise fuel expenditure by one of the other spacecraft. For the majority of manoeuvres the effect of this is a reduction in the fuel differences and thus a more fuel balanced manoeuvre (on average reduction in the sum of the fuel differences of 0.35% relative to the initial sum of the fuel differences). This is not the case however for triTTN-triTTN manoeuvres where an increase is still observed. This increase in the sum of the fuel differences however is at a value less than that of the benchmark and PAM: μ_{PAM} =0 cases.

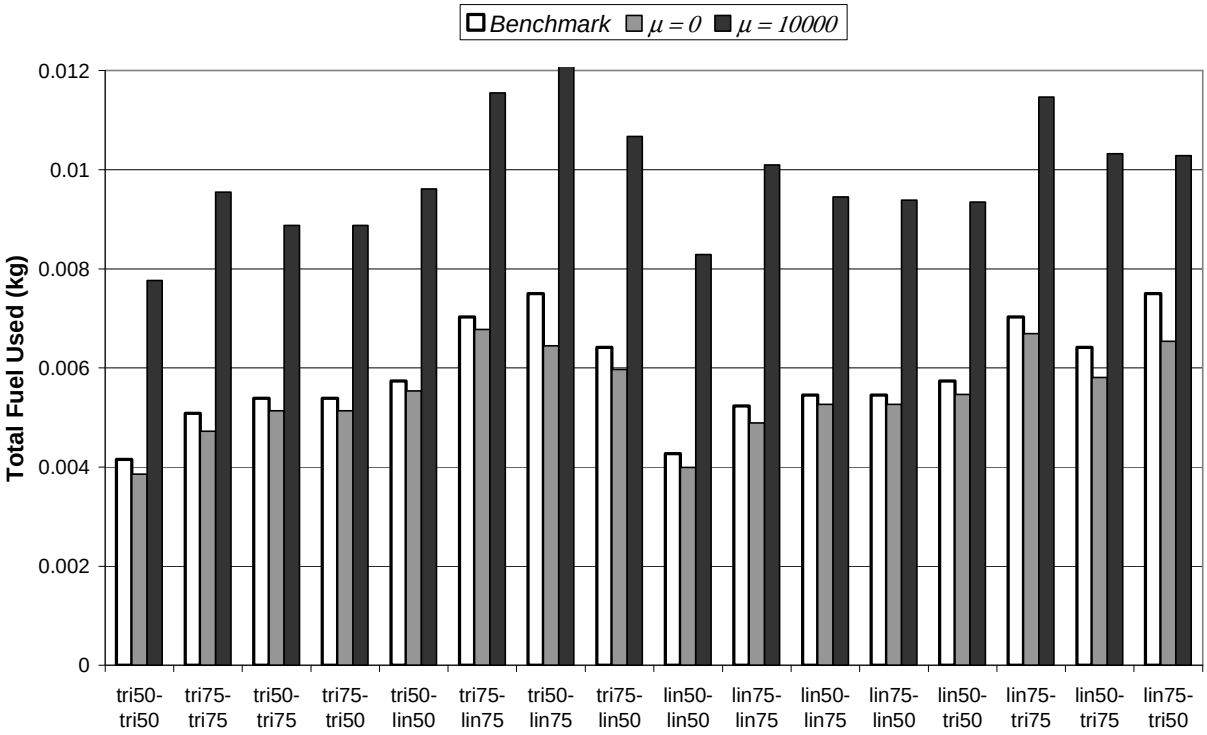

Figure 8-9 Total fuel manoeuvre comparison for 16 DARWIN-like manoeuvres. The chart shows that when $\mu_{PAM} = 0$ the PAM always finds spacecraft positions that require less fuel to achieve than the BPAM. However when $\mu_{PAM} = 10^3$ the fuel consumption is much greater in order to achieve better fuel balancing (see Figure 8-10)

Comparison of the manoeuvre duration is shown in Figure 8-11. For all manoeuvres using PAM: μ_{PAM} =0 optimisation the manoeuvre duration is increased over the BPAM. This ranges between an increase of 7-20% with an average of 10%. For the PAM: μ_{PAM} =10^3 case the manoeuvre duration difference varies with one third of the fuel balancing manoeuvres taking less time than the equivalent benchmark manoeuvre. The manoeuvre duration ranges from an increase of 40% to a decrease of 19% over the BPAM.

The results from the manoeuvre analysis of the PAM and BPAM show the use of the PAM does have its advantages. Using the PAM for fuel minimising gives a decrease in fuel consumption for a small increase in manoeuvre duration. Using the PAM for fuel balancing decreases the fuel differences across the fleet at the expense of an increase in fuel cost with variable performance gains and losses in manoeuvre duration. The extent to which the PAM performs fuel balancing can be varied through the parameter μ_{PAM} thus it is possible to trade-off fuel balancing gains to reduce fuel consumption. A further trade-off may be necessary when planning manoeuvres to ensure that time optimisation goals are not compromised too much to attain fuel optimisation goals.

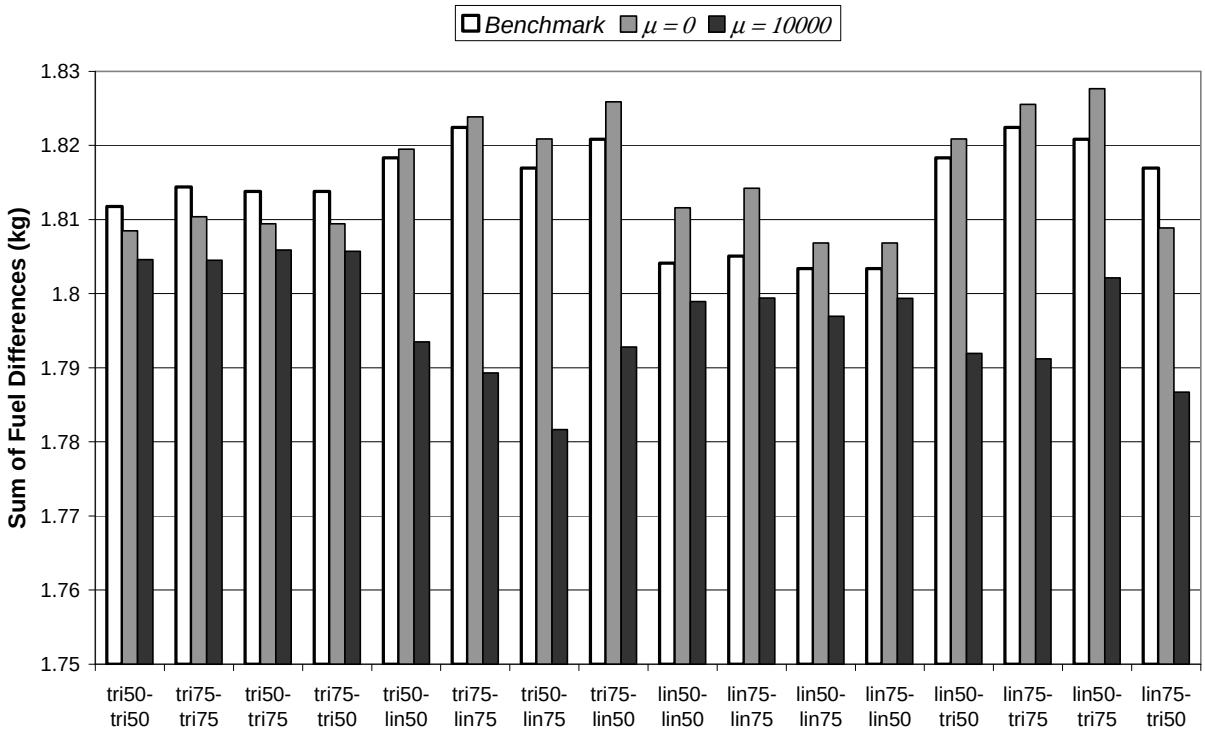

Figure 8-10 Fuel difference manoeuvre comparison for 16 DARWIN-like manoeuvres. The chart shows that when $\mu_{PAM} = 10^3$ the post-manoeuvre sum of the fuel differences is lower than for the BPAM or when $\mu_{PAM} = 0$. This shows that the PAM with $\mu_{PAM} = 10^3$ is better able to find spacecraft positions that improve fuel balancing within the formation.

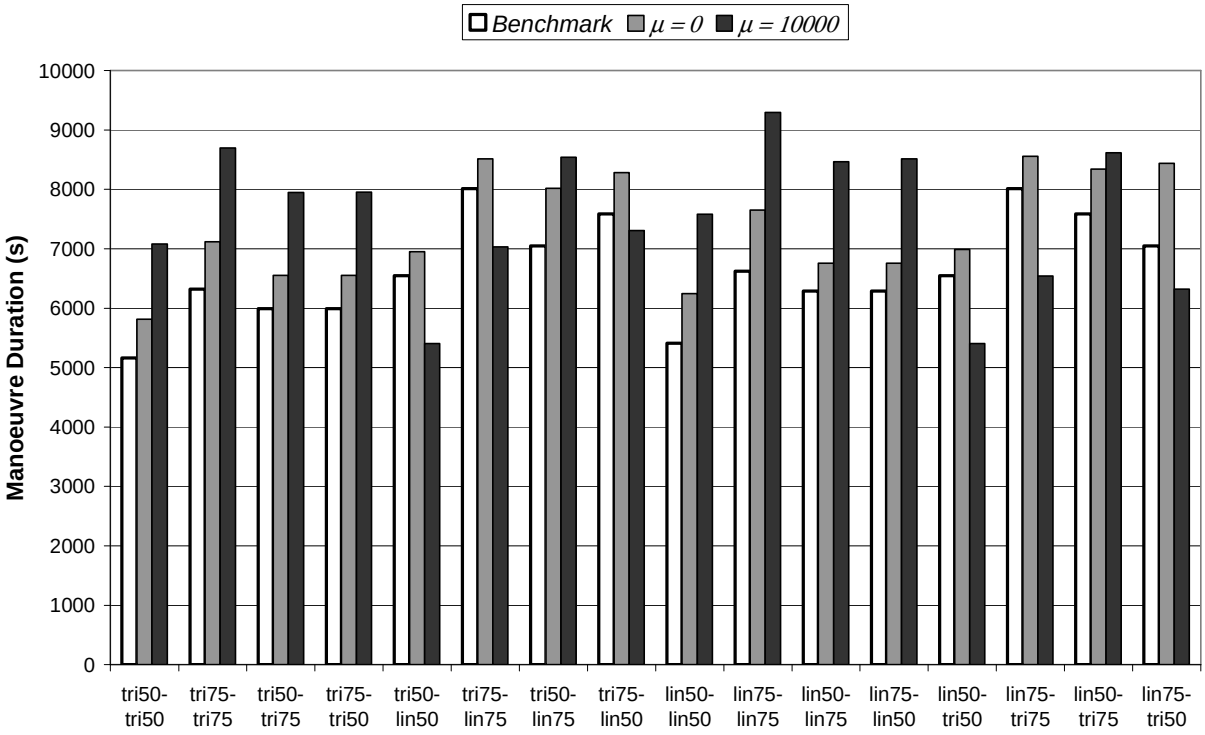

Figure 8-11 Manoeuvre duration comparison for 16 DARWIN-like manoeuvres. The chart shows that in general the PAM optimised manoeuvres have a higher manoeuvre time than the BPAM manoeuvres. The exception is some PAM $\mu_{PAM} = 10^3$ manoeuvres.

8.3.4 Analysis of μ_{PAM}

The value of μ_{PAM} in Equation (8.5) determines the extent to which fuel balancing is optimised during the calculation. When $\mu_{PAM} = 0$ the fuel balancing term in Equation (8.5) becomes zero and thus only fuel minimisation is sought by the objective function. For $\mu_{PAM} > 0$ the fuel balancing term affects the total cost of the manoeuvre. Analysing how the value of μ_{PAM} affects the fuel balancing, fuel consumption and manoeuvre time of a manoeuvre is important in determining appropriate values of μ_{PAM} to apply.

8.3.4.1 Initially Balanced Fuel

The affect of μ_{PAM} on random manoeuvres for an initially fuel balanced formation is shown in Figure 8-12. The data points represent a random manoeuvre for a particular value of μ_{PAM} with each μ_{PAM} repeated twenty times to obtain the data set. The solid filled line shows the mean for each value of μ_{PAM} . The *x*-axes are logarithmic while the *y*-axes are linear. The initial fuel levels of the formation are $mf_i(t_0) = \{5.0, 5.0, 5.0, 5.0\}\, kg$. The upper chart in Figure 8-12 shows the formation fuel consumption for the manoeuvre, the middle chart shows the sum of the fuel differences for each manoeuvre whilst the lower chart shows the manoeuvre time.

Figure 8-12 (upper) clearly shows that for $\mu_{PAM} < 0.1$ and $\mu_{PAM} > 1$ changing μ_{PAM} has very little effect on the mean fuel consumed during the manoeuvre. Within the effective

range however ($0.1 \leq \mu_{PAM} \leq 1$) the amount of fuel consumed changes significantly. The data is slightly different for Figure 8-12 (middle) and Figure 8-12 (lower). Again the charts show an effective range for μ_{PAM} however, for both charts, this range is extended to $0.01 \leq \mu_{PAM} \leq 1$. This data indicates that for a range $0.01 \leq \mu_{PAM} \leq 1$ some fuel balancing is obtainable without sacrificing fuel consumption.

Figure 8-12 also shows other interesting trends relating to the initial fuel balanced state of the formation. Firstly, Figure 8-12 (middle) shows that, although the initial fuel is balanced, the PAM is unable to maintain that balance even when optimising for fuel balancing. This is not surprising due to the complexity of the manoeuvres being performed. The mean sum of the fuel differences for fuel balancing (i.e. when $\mu_{PAM} > 1$) is only ~1g and so the PAM demonstrates a remarkable ability to maintain fuel balancing within <1g difference between all the spacecraft in the formation over a 6000 sec manoeuvre. Secondly, for an initially fuel balanced formation, Figure 8-12 (lower) shows that the fuel balancing manoeuvres are more time efficient than the fuel minimising manoeuvres. Again this is expected due to the nature of the manoeuvre planning. Since one spacecraft will always adopt a bang-bang thrust profile, to obtain better fuel balancing the other spacecraft, adopting the bang-coast-bang thrust profile, will have their coast period reduced (and hence their thrust periods increased to maintain a uniform manoeuvre time). To lower the fuel consumption of the more fuel depleted spacecraft the spacecraft positions with lower manoeuvre times associated with them are selected.

8.3.4.2 Initially Unbalanced Fuel

The affect of μ_{PAM} on manoeuvres where the initial spacecraft fuel amounts are unbalanced is given in Figure 8-13. Table 8-3 shows the initial fuel amounts for the five cases examined. The cases are chosen to reflect a number of different fuel difference scenarios. Case 1 is the same as the previous sub-section with the initial fuel balanced. This case emulates fuel amounts at the start of the mission. Cases 2 and 3 are defined with one spacecraft fuel deficient with 0.3kg and 0.5kg fuel deficiency respectively. These cases emulate potential fuel amount conditions after recovery from safe-mode. Cases 4 and 5 are defined such that they give the same 'sum of the fuel differences' results, Equation (8.27), as cases 2 and 3 but with more than one spacecraft fuel deficient. These cases emulate potential fuel conditions during the mission. In Figure 8-13 the solid markers (■, ▲, ♦, × and *) represent the mean value over twenty random manoeuvres for each initial formation fuel distribution case and the solid line represents the mean for each value of μ_{PAM} over all the initial formation fuel distribution cases. The *x*-axes are logarithmic while the *y*-axes are linear. The upper chart in Figure 8-13 shows the formation fuel consumption for the manoeuvre, the middle chart shows the sum of the fuel differences for each manoeuvre whilst the lower chart shows the manoeuvre time.

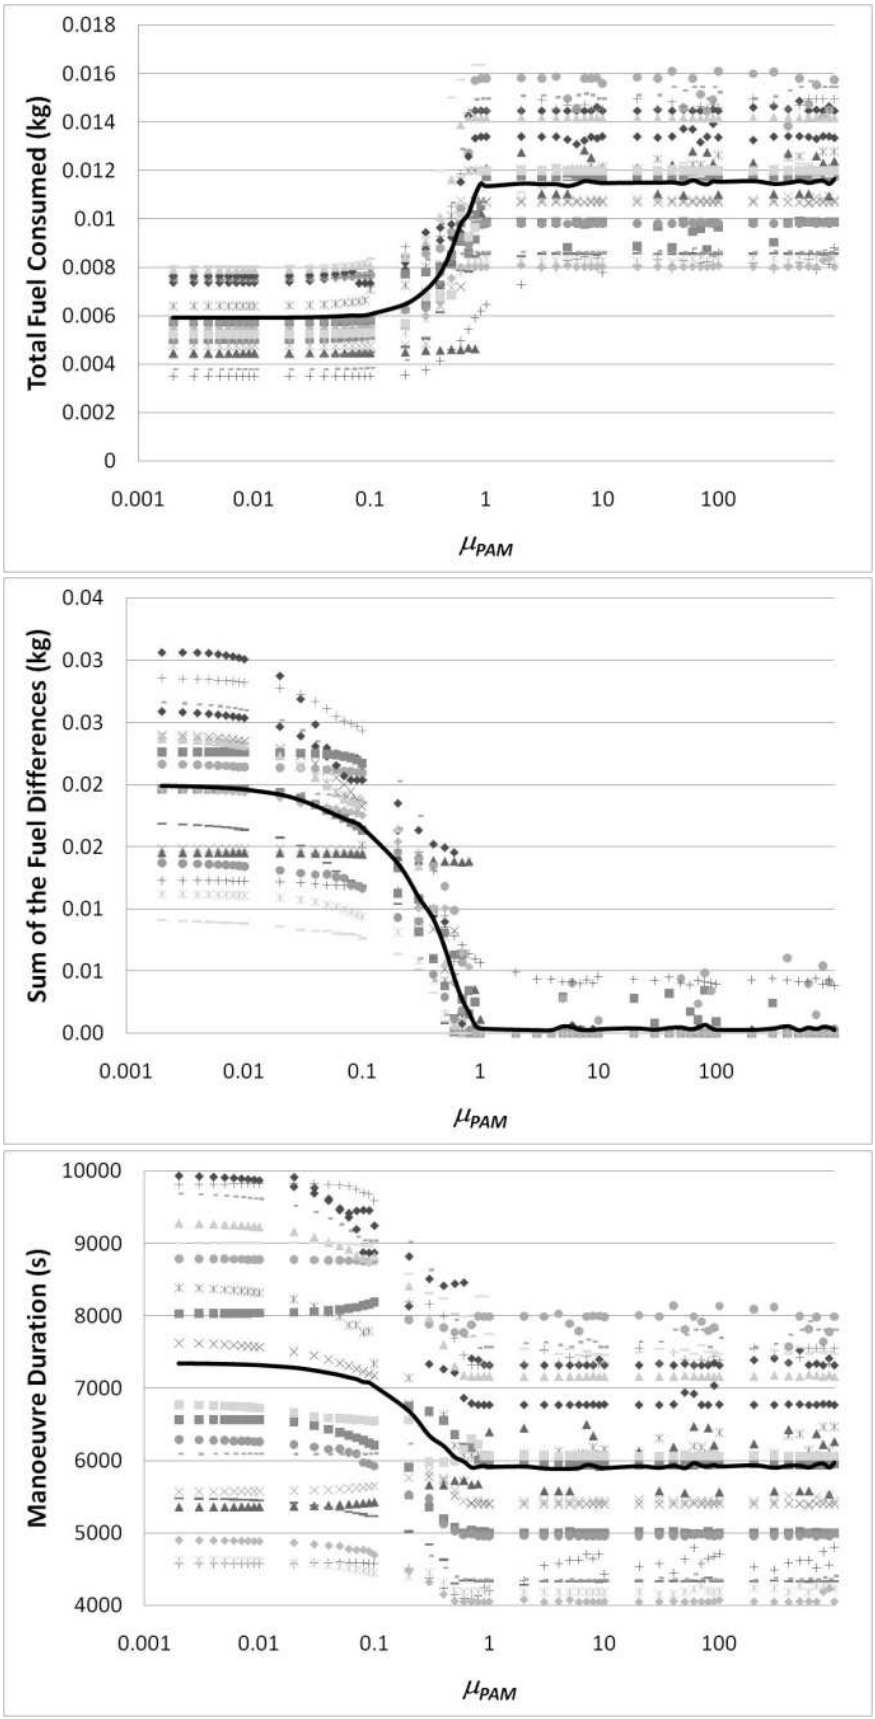

Figure 8-12 The affect of μ_{PAM} on the formation fuel consumption (upper) , the sum of the fuel differences (middle) and the manoeuvre time (lower) for an initially fuel balanced formation. The markers represent the data points while the solid lines represent the mean.

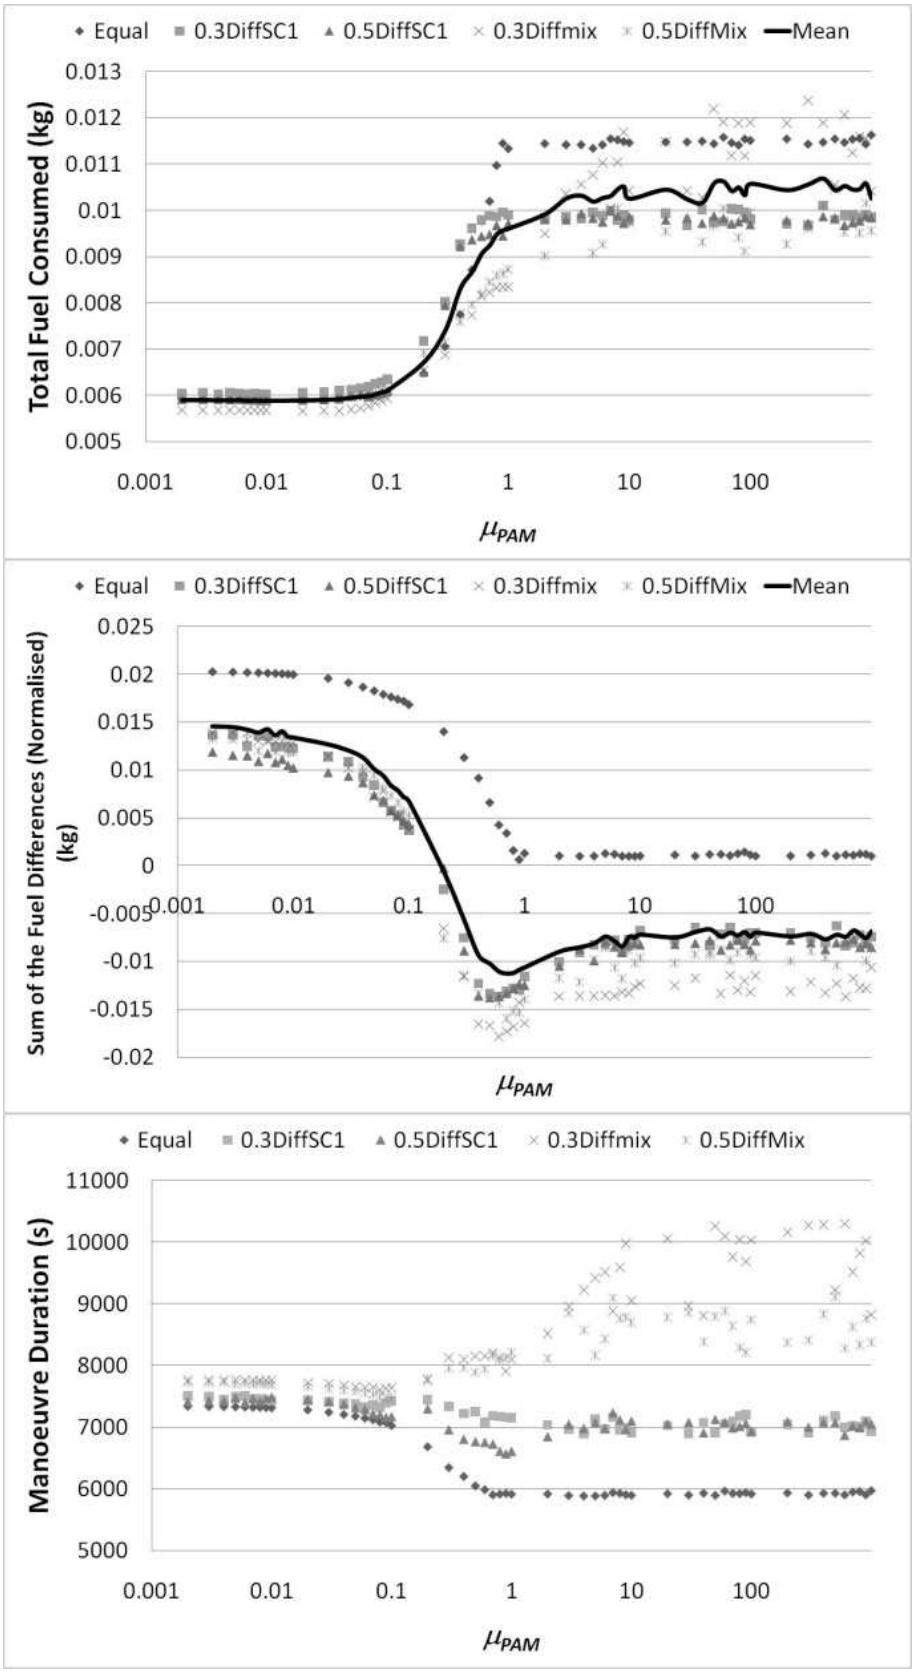

Figure 8-13 The affect of μ_{PAM} on the formation fuel consumption (upper) , the sum of the fuel differences (middle) and the manoeuvre time (lower) for an initially fuel un-balanced formation

Table 8-3 Initial fuel amounts for the μ_{PAM} -analysis

Case	Code	Initial Fuel Amounts	Initial sum of fuel differences (kg)
1	Equal	$f_i(t_0) = \{5.0, 5.0, 5.0, 5.0\} kg$	0.0
2	0.3DiffSC1	$f_i(t_0) = \{5.0, 4.7, 5.0, 5.0\} kg$	1.8
3	0.5DiffSC1	$f_i(t_0) = \{5.0, 4.5, 5.0, 5.0\} kg$	3.0
4	0.3MixSC1	$f_i(t_0) = \{5.0, 4.7, 4.8, 4.8\} kg$	1.8
5	0.5MixSC1	$f_i(t_0) = \{5.0, 4.5, 4.7, 4.7\} kg$	3.0

From the fuel consumption chart (Figure 8-13 upper) the effective range of μ_{PAM} is extended to approximately $0.05 \le \mu_{PAM} \le 10$ when the different initial fuel difference values are included in the analysis. Outside of this range there is very little change observed in fuel consumption for each initial formation fuel distribution case. When $\mu_{PAM} < 0.05$ there is very little difference in the total amount of fuel consumed between the five cases. For this range of μ_{PAM} the difference between the cases using the maximum and minimum amount of fuel is about 0.1g. However during the effective range of μ_{PAM} the fuel usage between the cases diverges so that when $\mu_{PAM} > 10$ the difference between the cases using the maximum and minimum amount of fuel has risen to approximately 1.5g. There appears to be no correlation relating the type of initial fuel case that results in higher manoeuvre fuel consumption. When $\mu_{PAM} < 0.05$, case 4 shows the lowest fuel consumption but for $\mu_{PAM} > 10$ case 4 shows the highest fuel consumption. This also appears to have little to do with the mixed initial fuel differences cases since case 5 shows the lowest fuel consumption for $\mu_{PAM} > 10$. A final point to note from Figure 8-13 (upper) is that the initially balanced fuel difference case requires the use of significantly more fuel for fuel balancing than the majority of the initially unbalanced fuel difference cases.

In Figure 8-13 (middle) the sum of the fuel differences for each case has been normalised to aid representation on the chart. This has been achieved by subtracting the initial sum of fuel differences for each case (from Table 8-3) for every data point. In this representation, positive values show a decrease in fuel balancing whilst negative values show an increase in fuel balancing. The data show the extended effective range of μ_{PAM} as approximately $0.005 \le \mu_{PAM} \le 10$ independent of the initial spacecraft fuel difference. As in Figure 8-12 this indicates that there is a range of μ_{PAM} values ($0.005 \le \mu_{PAM} \le 0.1$) where small fuel balancing gains can be achieved without sacrificing an increase in the amount of fuel consumed during the manoeuvre. The data in Figure 8-13 (middle) also appears to indicate that for all the initial fuel difference cases maximum fuel balancing can be achieved using $0.5 \le \mu_{PAM} \le 0.7$. This would appear counter-intuitive as from the cost function, Equation (8.5), the fuel balancing term becomes more dominant as $\mu_{PAM} \to \infty$. However the value of the cost function is affected by both the fuel minimising term, Equation (8.26), and the fuel balancing term, Equation (8.27), and reveals some interesting properties of the cost function. This is covered in more detail in sub-section 8.3.4.3.

Comparing how fuel balancing in Figure 8-13 (middle) is affected by different initial fuel difference cases shows that when the fuel differences are mixed amongst the spacecraft (cases 4 and 5) the PAM on average is able to find better fuel balancing positions than when the initial fuel difference is on a single spacecraft (cases 2 and 3). This is because when the initial spacecraft fuel differences are mixed the fuel differences between individual spacecraft are smaller. This means that small changes in the individual spacecraft fuel differences result in large changes to the sum of the fuel differences. But in the cases where the initial fuel difference is on a single spacecraft, small differences in the individual spacecraft fuel differences only result in small changes in the sum of the fuel differences. Thus more solutions that give better manoeuvre fuel balancing will occur when the initial spacecraft fuel differences are mixed amongst the spacecraft.

Figure 8-13 (lower) shows how the manoeuvre duration is affected by differing values of μ_{PAM} for each of the five initial fuel difference cases. For all the cases there is little difference in manoeuvre duration for fuel minimising values of μ_{PAM} (i.e. $\mu_{PAM} < 0.05$). However during the effective range of μ_{PAM} the initial fuel difference case manoeuvre durations diverge and either increase or decrease. Here there is a clear difference between the different initial fuel cases. For case 1 (the initially fuel balanced case) a decrease in manoeuvre duration over fuel minimising values of μ_{PAM} is shown. Cases 2 and 3 follow the same pattern to a lesser degree. However for cases 4 and 5 (when initial fuel differences are mixed amongst the spacecraft) the pattern is reversed with an observed increase in manoeuvre duration for fuel balancing values of μ_{PAM}. This correlates with the data in Figure 8-13 (middle) indicating that for cases 4 and 5 the increased performance in fuel balancing comes at a cost of increased manoeuvre duration. The manoeuvre duration data also shows that for the observed best fuel balancing value of μ_{PAM} (≈0.6) the manoeuvre duration increases are only small (a 3-5% increase in manoeuvre duration) compared with the increase observed for $\mu_{PAM} > 10$ (an 11-25% increase) for initial fuel difference cases 4 and 5. This means that $0.5 \le \mu_{PAM} \le 0.7$ not only appears to give the best fuel balancing it does so with little increase in manoeuvre duration for initial fuel difference cases 4 and 5 and a reduction in manoeuvre duration for cases 1, 2 and 3. This however is tempered by an average increase in fuel consumption by about one third.

8.3.4.3 $0.5 \le \mu_{PAM} \le 0.7$ Analysis

The data in Figure 8-13 appears to indicate that when $0.5 \le \mu_{PAM} \le 0.7$ maximum fuel balancing can be achieved regardless of the initial fuel differences within the formation. This is reinforced when the μ_{PAM} analysis is repeated for different values of thrust saturation, α, and thruster *Isp*. Figure 8-14 shows the sum of the fuel differences for different values of μ for three separate values of α and *Isp*. The data was obtained using the same method as in sub-sections 8.3.4.1 and 8.3.4.2. The data represented is the mean data from twenty repetitions of the analysis. Figure 8-14 (upper left) shows the case where $\alpha = 6\times10^{-2}$ N and $I_{sp} = 3300\text{s}$, the upper right chart shows $\alpha = 6\times10^{4}$ N and $I_{sp} = 3300\text{s}$, and the lower chart shows $\alpha = 6\times10^{-3}$ N and $I_{sp} = 3300\text{s}$. From Figure 8-14 it is clear to see that regardless of the

input thrust capabilities of the formation when $0.5 \leq \mu_{PAM} \leq 0.7$ optimal fuel balancing is achieved. It is currently unknown why this fuel balancing minima should occur for this value of μ_{PAM} however it is a useful value to use whilst analysing the PAM further.

8.3.5 DARWIN Tour Analysis

The analysis performed in the previous sub-sections has been carried out on a single manoeuvre basis. In this sub-section the performance of the PAM on a tour of different manoeuvres is analysed to assess the PAM's performance over time. The tour consists of 133 manoeuvres and was generated using the benchmark Science Operations Module (BSOM). The tour has a target tour time of 100 days, uses the Year1 taskflag and starts at the first star in the catalogue. The tour covers all the separated spacecraft interferometry manoeuvres (retarget, reconfiguration, resize and rotate) and combinations thereof. The fuel data generated during the tour is only representative of the fuel used during the formation reconfiguration manoeuvres within a gravity-free force model. Any fuel consumption due to attitude control, station-keeping requirements, formation-keeping requirements and on-observation manoeuvres is not included in this data set.

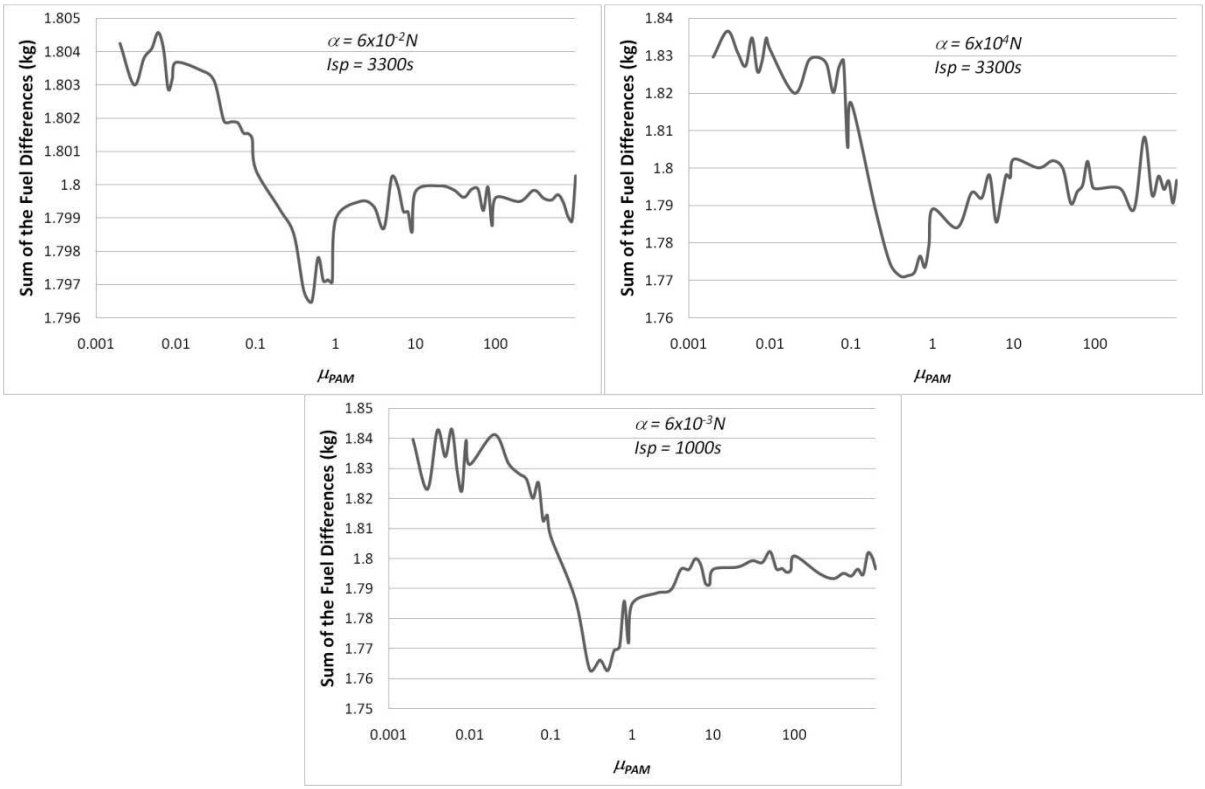

Figure 8-14 Sum of the fuel differences vs. μ_{PAM} for three different thrust capabilities. The charts show that regardless of the thrust capabilities the fuel balancing performance of the PAM with respect to μ_{PAM} follows the same pattern as in Figure 8-13

8.3.5.1 Tour Analysis

Figure 8-15, Figure 8-16 and Figure 8-17 show the results from the tour analysis for initial formation fuel distribution cases 1, 2 and 4 respectively. In all the charts the x and y-axes are linear and data is given for four values of μ_{PAM}:

- $\mu_{PAM} = 0$ fuel minimising
- $\mu_{PAM} = 0.05$ slight fuel balancing without sacrificing fuel consumption
- $\mu_{PAM} = 0.6$ observed best fuel balancing
- $\mu_{PAM} = 10$ highest boundary of effective μ_{PAM} range from μ_{PAM} analysis

These values of μ_{PAM} were chosen as they represent interesting values found during the single manoeuvre analysis performed in sub-section 8.3.4.

The tour data for the initially balanced fuel case is given in Figure 8-15. When $\mu_{PAM} = 0$ the formation clearly uses less fuel during the tour than for other values of μ_{PAM} however for $\mu_{PAM} = 0.05$ after 133 manoeuvres the total fuel difference is only ~3g, an increased fuel consumption of less than 2%. For the $\mu_{PAM} = 0.6$ and $\mu_{PAM} = 10$ cases however the increased fuel consumption is 65% and 123% respectively. Examining the sum of the fuel difference data shows that for $\mu_{PAM} = 0$ and $\mu_{PAM} = 0.05$ the sum of the fuel differences gets progressively higher as the tour advances. However the rate at which this occurs for the $\mu_{PAM} = 0.05$ case is significantly slower such that by the end of the tour the sum of the fuel differences is 20% less than the $\mu_{PAM} = 0$ case. This emphasises the effectiveness of using $\mu_{PAM} = 0.05$ for an initially fuel balanced formation since a 20% increase in fuel balancing can be achieved for only a <2% increase in fuel consumption. For $\mu_{PAM} = 0.6$ and $\mu_{PAM} = 10$ the sum of the fuel difference data mirrors that found in Figure 8-12 (middle). The PAM is unable to return to completely balanced fuel but maintains the sum of the fuel differences to a very low level (<2g for this data set).

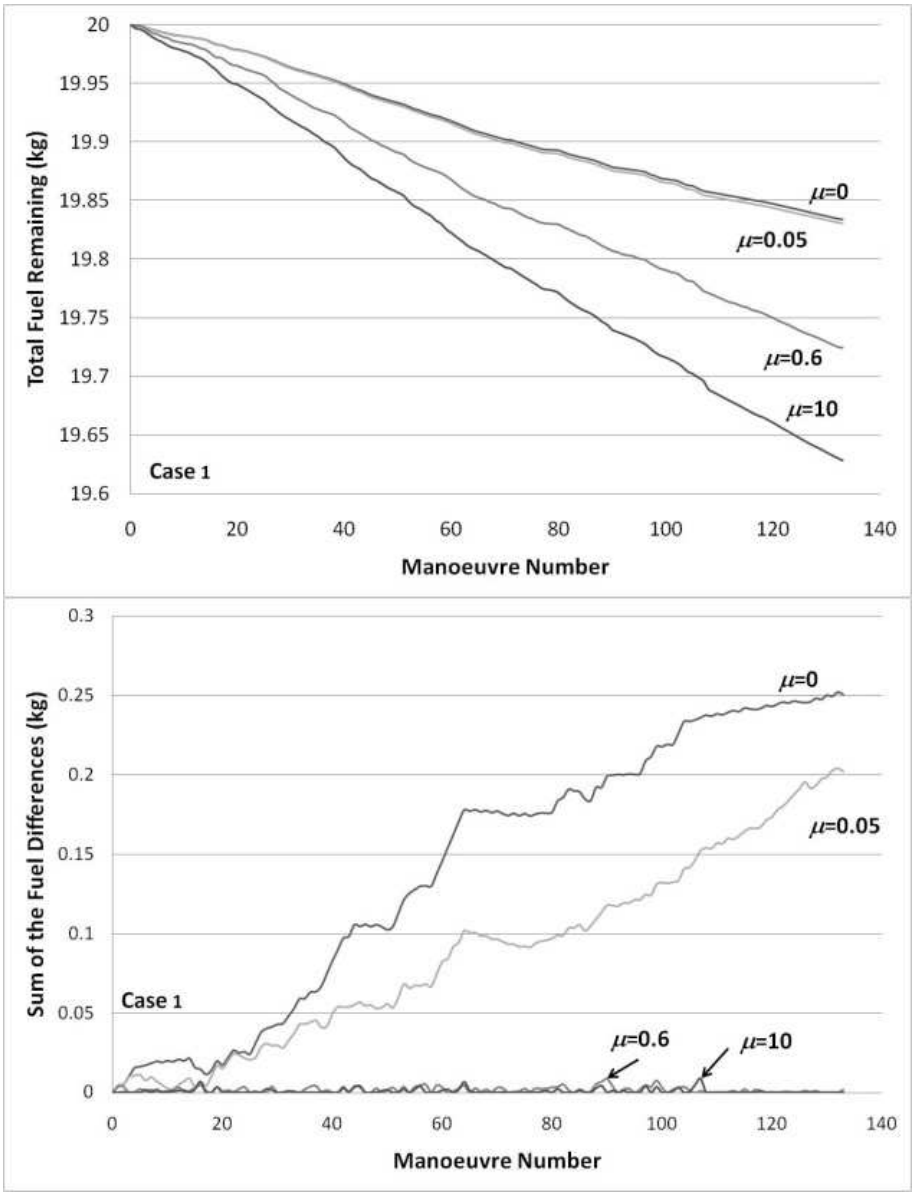

Figure 8-15 PAM tour analysis for case 1 – total fuel remaining (upper) and sum of the fuel differences (lower). The charts show that when $\mu_{PAM} = 0$ the least fuel is consumed to the detriment of fuel balancing. When $\mu_{PAM} = 0.6$ and $\mu_{PAM} = 10$ the PAM is able to maintain excellent fuel balancing to the detriment of fuel consumption.

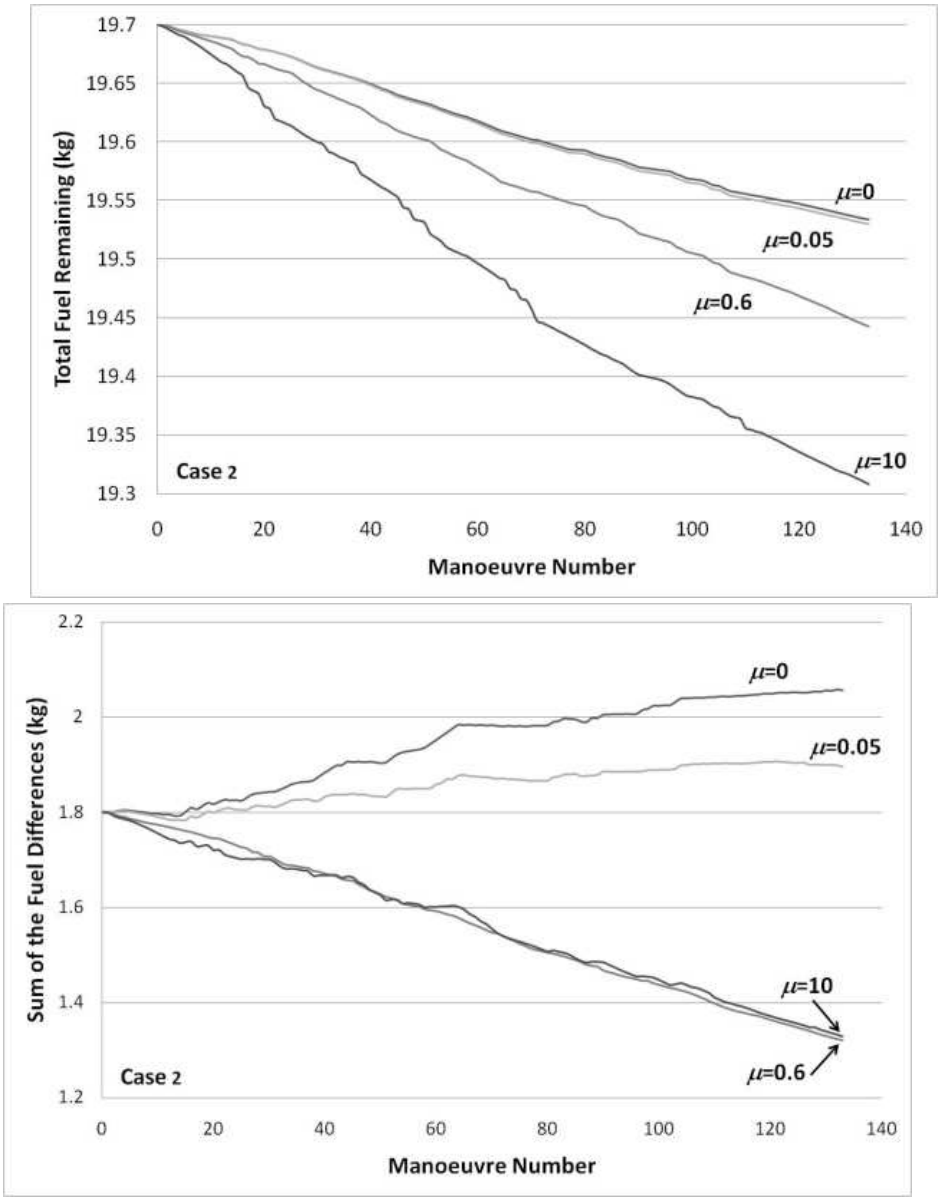

Figure 8-16 PAM tour analysis for case 2 – total fuel remaining (upper) and sum of the fuel differences (lower). The charts show that when $\mu_{PAM}=0$ the least fuel is consumed to the detriment of fuel balancing and that when $\mu_{PAM}=0.6$ the best fuel balancing is achieved to the detriment of fuel consumed (but not as much as the $\mu_{PAM}=10$ tour).

The tour data for the initial fuel difference case 2 (where the initial fuel levels are $f_i(t_0)=\{5.0,4.7,5.0,5.0\}kg$) can be seen in Figure 8-16. Here the initial total fuel is 19.7kg and the initial sum of the fuel differences is 1.8kg. The fuel consumption data shows that once again the $\mu_{PAM}=0$ and $\mu_{PAM}=0.05$ cases consume significantly less fuel with the $\mu_{PAM}=0.05$ case only consuming ~2% more fuel than when $\mu_{PAM}=0$. The $\mu_{PAM}=0.6$ and $\mu_{PAM}=10$ cases however consume 54% and 135% more fuel respectively. The fuel balancing data again shows that for $\mu_{PAM}=0$ the sum of the fuel differences steadily increases as the tour progresses so that by the tour end the fuel balancing is 14% worse than

at the beginning of the tour. For $\mu_{PAM} = 0.05$ the sum of the fuel differences slightly increases over the tour duration so that the fuel balancing is worse by ~5%. Though not shown in Figure 8-16 initial fuel difference case 3 shows similar results to those found for case 2.

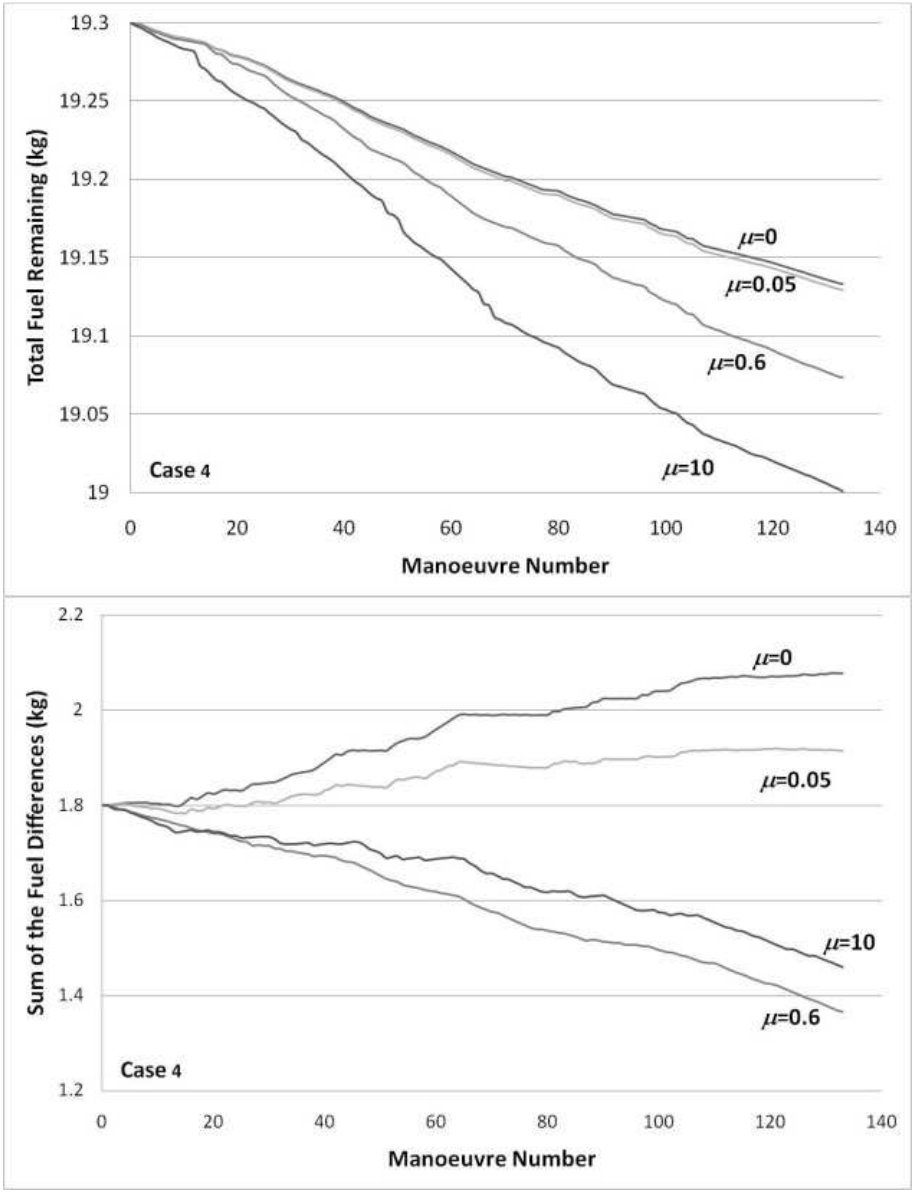

Figure 8-17 PAM tour analysis for case 4 – total fuel remaining (upper) and sum of the fuel differences (lower). The charts show similar results to Figure 8-16.

Figure 8-17 shows the tour data for initial fuel difference case 4 (where the initial fuel levels are $f_i(t_0) = \{5.0, 4.7, 4.8, 4.8\}\, kg$). Here the initial total fuel is 19.3kg whilst the initial sum of the fuel differences is 1.8kg. As in the previous two cases the fuel consumption data show the similarity between the $\mu_{PAM} = 0$ and $\mu_{PAM} = 0.05$ cases with the $\mu_{PAM} = 0.05$ case consuming only ~2% more fuel than when $\mu_{PAM} = 0$. The fuel consumption data for the $\mu_{PAM} = 0.6$ case shows an increase of ~36% and the $\mu_{PAM} = 10$ tour shows an increase of ~80%. The data is also similar to the initial fuel difference case 2 for fuel balancing since

when $\mu_{PAM} = 0$ fuel balancing is ~15% worse and for $\mu_{PAM} = 0.05$ ~6.5% worse. An increase in fuel balancing is observed when $\mu_{PAM} = 0.6$ and $\mu_{PAM} = 10$ with increases of ~24% and ~19% respectively. This again emphasises the benefit of using $\mu_{PAM} = 0.6$ for fuel balancing over $\mu_{PAM} = 10$. Though not shown in Figure 8-17 initial fuel difference case 5 shows similar results to those found for case 4.

8.3.5.2 Tour Fuel Balancing Performance Analysis

The fuel balancing performance is the ratio between the increased percentage in fuel balancing per percentage increase in fuel consumption:

$$fuel\ balancing\ performance = \frac{\%\ increase\ in\ fuel\ balancing}{\%\ increase\ in\ fuel\ consumption} \tag{8.28}$$

This gives a metric to compare the fuel efficiency of each fuel balancing tour. This performance is important to examine for a number of reasons:

- Fuel balancing requires more fuel consumption per manoeuvre
- The rate of fuel consumption increases with μ_{PAM}
- Fuel balancing is desirable but not to the detriment of minimising fuel consumption

In assessing tour fuel balancing performance two different approaches are used depending on the initial formation fuel distribution case. For both approaches the percentage increase in fuel consumption is relative to the fuel minimising tour (i.e. $\mu_{PAM} = 0$). For initially un-balanced fuel distributions (cases 2-5) the percentage increase in fuel balancing is relative the initial sum of the fuel differences for each case. However for case 1 the initial sum of the fuel differences is zero so this metric cannot be used. For case 1 therefore the percentage increase in fuel balancing is calculated relative to the sum of the fuel differences of the $\mu_{PAM} = 0$ tour.

Figure 8-18 shows how the fuel balancing performance of 10 separate tours varies with μ_{PAM}. The tours were generated using the BSOM, a target tour time of 100 days, the Year1 taskflag and 10 different starting stars. The markers show the data for the ten tours and the solid line represents the mean. For clarity not all the data points are shown but their values are reflected in the mean data. Positive performance indicates an increase in fuel balancing whilst negative performance indicates a decrease (i.e. an increase in the sum of the fuel differences). The performance slowly increases as μ_{PAM} decreases until a sharp maximum is reached at $\mu_{PAM} \cong 0.03$. As μ_{PAM} decreases further the individual tour data diverges greatly but the mean shows a sharp decrease of performance (sometimes negative). The data in Figure 8-18 shows that relative to a $\mu_{PAM} = 0$ tour the best fuel balancing performance in

relation to the amount of fuel consumed can be achieved when $\mu_{PAM} = 0.03$ for initial formation fuel distribution case 1.

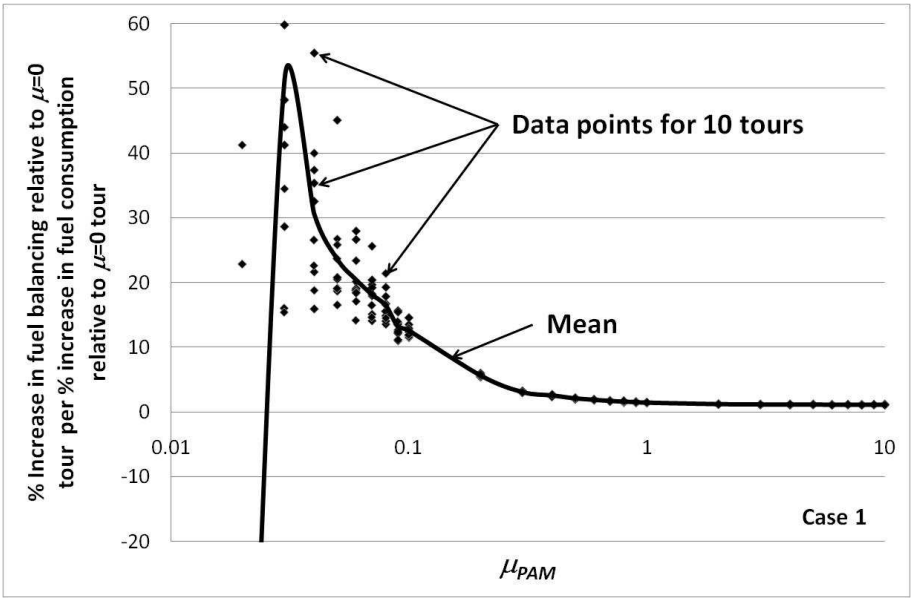

Figure 8-18 Tour performance data for 10 separate tours for initial fuel difference case 1 relative to the $\mu_{PAM} = 0$ tour. Tour performance reaches a maximum at $\mu_{PAM} = 0.03$ indicating that the most efficient fuel balancing is achieved at this value.

Figure 8-19 shows how the tour performance relative to the initial sum of the fuel differences changes with respect to μ_{PAM} for initial fuel difference cases 2 and 4. The markers represent data from the 10 difference tours used previously and the solid line represents the mean. Both charts show similar patterns to those in Figure 8-18 however this time the peak performance is found for $\mu_{PAM} = 0.1-0.2$. Additionally, for both charts, no performance increase is observed for $\mu_{PAM} \leq 0.09$. The observed peaks in Figure 8-18 represent the μ_{PAM} values that give the largest reduction in the sum of the fuel differences relative to the amount of extra fuel used. The position of these peaks is very favourable since comparing with the fuel remaining data for the same tours (Figure 8-13 – upper) these μ_{PAM} values represent the lower end of the extra fuel consumption (~7-8% extra fuel consumption over the $\mu_{PAM} = 0$ case). Had the peaks been found at higher μ_{PAM} values the advantage of the fuel balancing performance might have been outweighed by the actual fuel consumption requirements to achieve that performance.

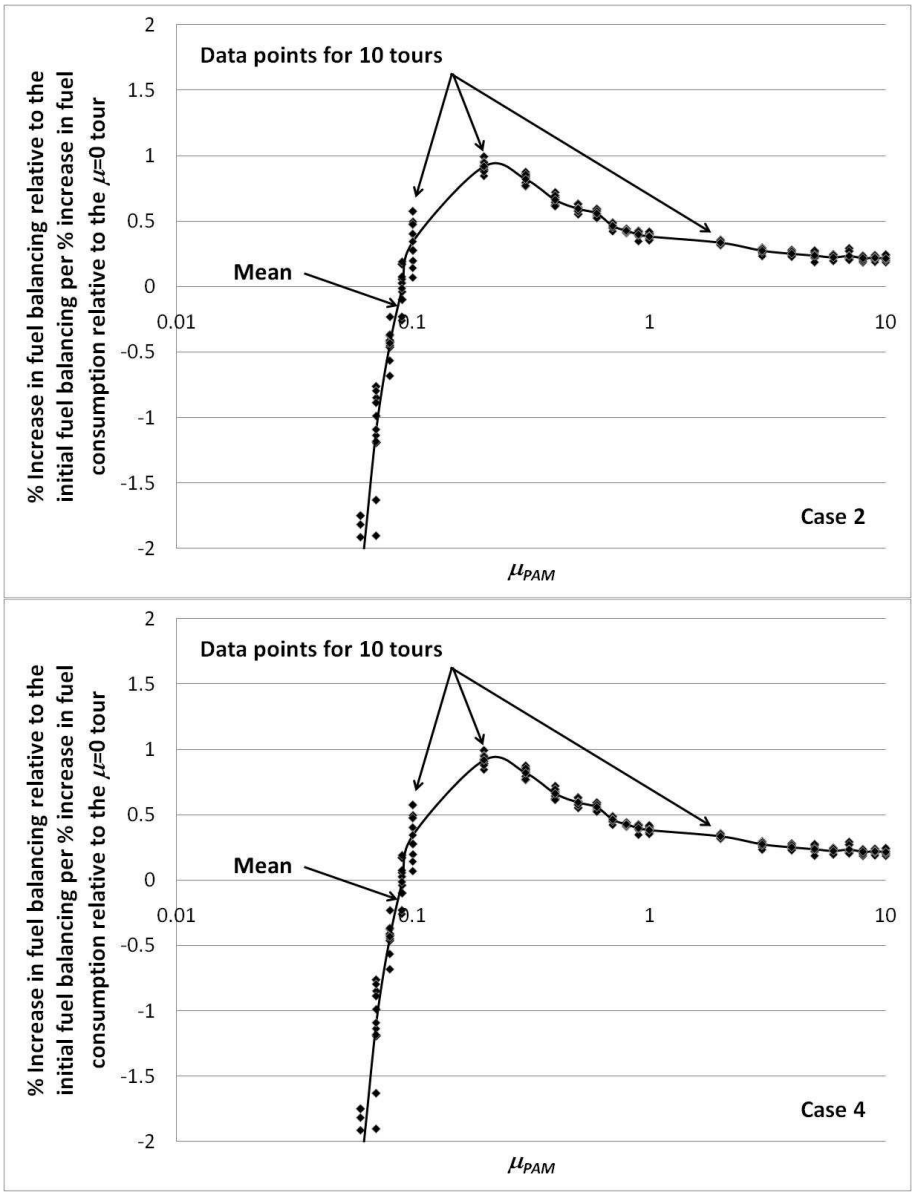

Figure 8-19 Tour performance data for 10 separate tours for initial fuel case 2 (upper) and case 4 (lower) relative to the initial sum of the fuel differences. The charts both show that maximum performance is found for $\mu_{PAM} = 0.1-0.2$ indicating that the most efficient fuel balancing is achieved at this value.

8.3.6 Observed Manoeuvre Anomalies within the Position Assignment Module

A few manoeuvre planning anomalies have been observed whilst examining the hundreds of manoeuvres for this analysis of the PAM.

8.3.6.1 Large Single Manoeuvre Fuel Consumption

The first anomaly is shown in Figure 8-20 and shows, on the upper chart, the total fuel remaining and, on the lower chart, the sum of the fuel differences for μ_{PAM} comparisons of a 133 manoeuvre tour. The tour is exactly the same to that in section 8.3.5 however due to the stochastic nature of the PAM optimisation routine the results are slightly different. The initial

fuel difference is case 4, $f_i(t_0) = \{5.0, 4.7, 4.8, 4.8\} kg$. The major anomaly from this iteration of the tour however is for the $\mu_{PAM} = 10$ tour. At manoeuvre three in the tour the PAM plans a manoeuvre that uses significantly more fuel (1.287kg) than the average for the tour (8.2g). The lower chart in Figure 8-20 shows the affect of this manoeuvre on the tour's fuel balancing performance. The manoeuvre reduces the initial fuel imbalance by 67% so that by the middle of the tour the fuel is as near completely balanced as possible in stark contrast to the other μ_{PAM} values examined in the same chart and in Figure 8-17. Whilst fuel balancing is important the balance between fuel consumption vs. fuel balancing is clearly not optimal in this example.

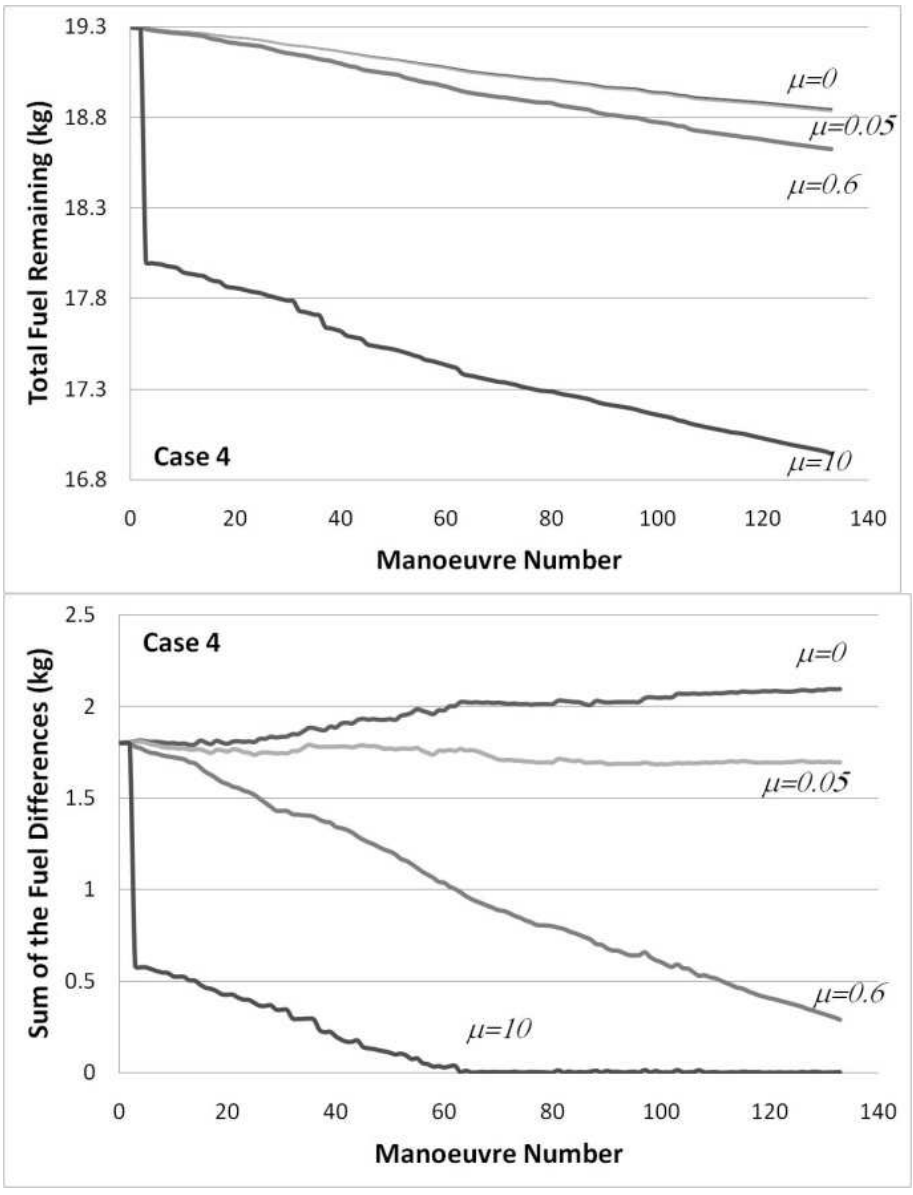

Figure 8-20 PAM manoeuvre planning anomaly – large manoeuvre fuel consumption. The upper chart shows a sharp drop in fuel remaining after one particular manoeuvre in the $\mu_{PAM} = 10$ tour. The sum of the fuel differences in the lower chart shows a corresponding sharp drop for the same manoeuvre.

The large single manoeuvre fuel consumption manoeuvre is shown in Figure 8-21. The manoeuvre involves a retarget from star 48 to star 51 in the catalogue but the linTTN configuration and the baseline of 27.8m remains. The initial positions are given by the open circles (○), the final positions by the closed circles (●) and the trajectories are the straight lines. All other initial conditions are the same as for this entire section, i.e. the spacecraft masses, $m_i = \{1100, 900, 900, 900\} kg$, but the initial fuel for the manoeuvre is $f_i(t_0) = \{4.9977, 4.6988, 4.7960, 4.7960\} kg$ since it is manoeuvre 3 in the tour. Figure 8-21 clearly shows the planned manoeuvre involving a large translation of the entire formation. The translation is a result of the optimisation function trying to balance the formation fuel since $\mu_{PAM} = 10$. Although the fuel consumption is large the sum of the fuel differences gains achieved by the manoeuvre result in the fuel balancing term in the cost function, Equation (8.5), dominating the fuel consumption term so that overall the cost for the manoeuvre is the lowest in the solution space.

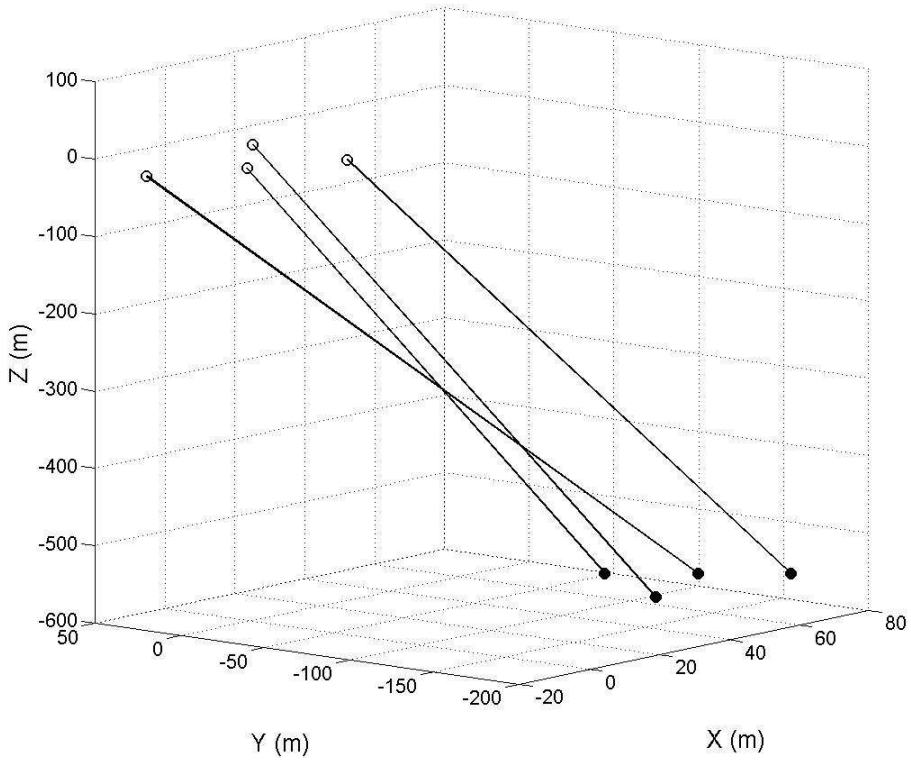

Figure 8-21 Large single manoeuvre fuel consumption anomaly example manoeuvre. The chart shows how the PAM has displaced the entire formation, much more than necessary to complete a manoeuvre, in order to achieve better fuel balancing across the formation.

The large single manoeuvre fuel consumption anomaly is a consequence partly of the nature of the optimisation cost function, the values being optimised (i.e. fuel) and the initial conditions of the manoeuvre. In this manoeuvre the optimisation routine translates the formation in order to improve the fuel balancing of the formation. With the BCS having the most fuel at the beginning of the manoeuvre and being the more massive spacecraft a large translation manoeuvre will cause the BCS to use more fuel than the TSs so eventually the fuel balancing across the formation will be improved. This, coupled with the nature of the cost

function, allows the manoeuvre to have an ultimately lower cost than other manoeuvres in the solution space even though the fuel consumption term, Equation (8.26), is comparatively much larger.

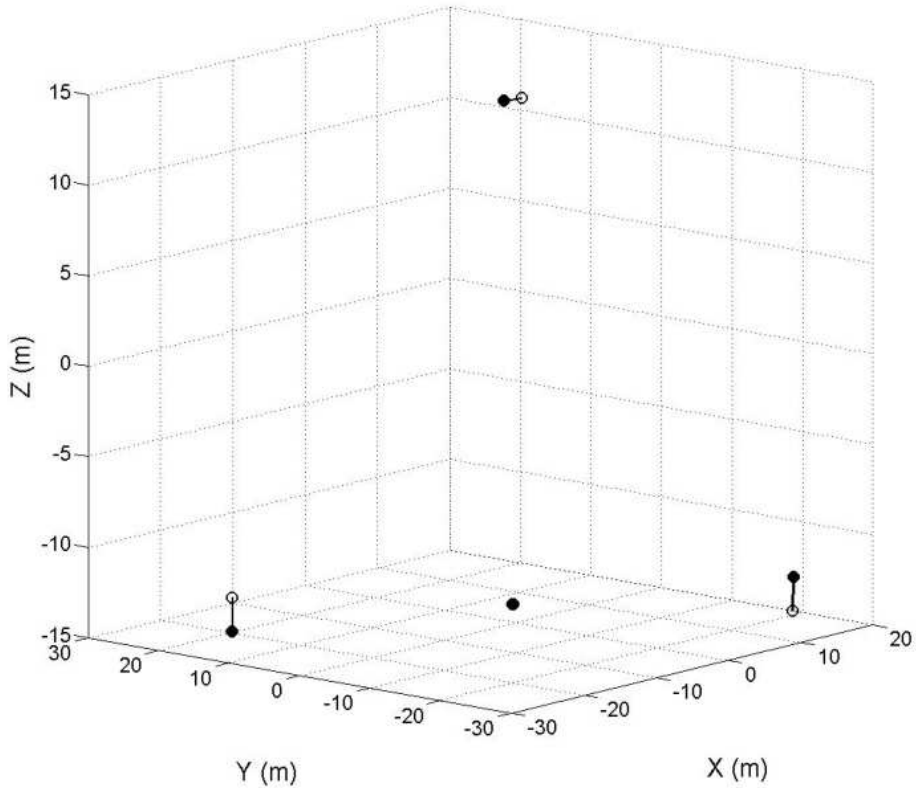

Figure 8-22 Example manoeuvre solution of the large single manoeuvre fuel consumption anomaly - equal spacecraft mass. In equalling the spacecraft mass the PAM has no reason to displace the entire formation to balance the fuel.

This type of manoeuvre planning anomaly can be mitigated in a number of ways. First, if the spacecraft masses are equal across the formation the anomaly disappears. Figure 8-22 shows the same manoeuvre as in Figure 8-21 but with the spacecraft masses equal at $m_i = \{900, 900, 900, 900\}\, kg$. Figure 8-23 shows the same manoeuvre but with the BCS having less initial fuel than the other spacecraft in the formation, i.e. $f_i(t_0) = \{4.6988, 4.9977, 4.7960, 4.7960\}\, kg$. Figure 8-24 shows the same manoeuvre but with the spacecraft ΔV replacing the fuel terms in the PAM cost function so that

$$J_{PAM} = \min_{X} \left\{ \sum_{i}^{n} \left(\Delta V_i(t_0) - \Delta V_i(t_f) \right) + \mu_{PAM} \sum_{i}^{n} \sum_{j:j\neq i}^{n} \left| \Delta V_i(t_f) - \Delta V_j(t_f) \right| \right\} \qquad (8.29)$$

where ΔV_i is the ΔV capability of spacecraft i. In all three manoeuvres the translation of the individual spacecraft is much smaller than the manoeuvre in Figure 8-21 and in-line with the average type of manoeuvre planned. This demonstrates that the large single fuel consumption manoeuvre anomaly is a result of a combination of spacecraft initial conditions, the nature of the cost function and the terms optimised within the cost function.

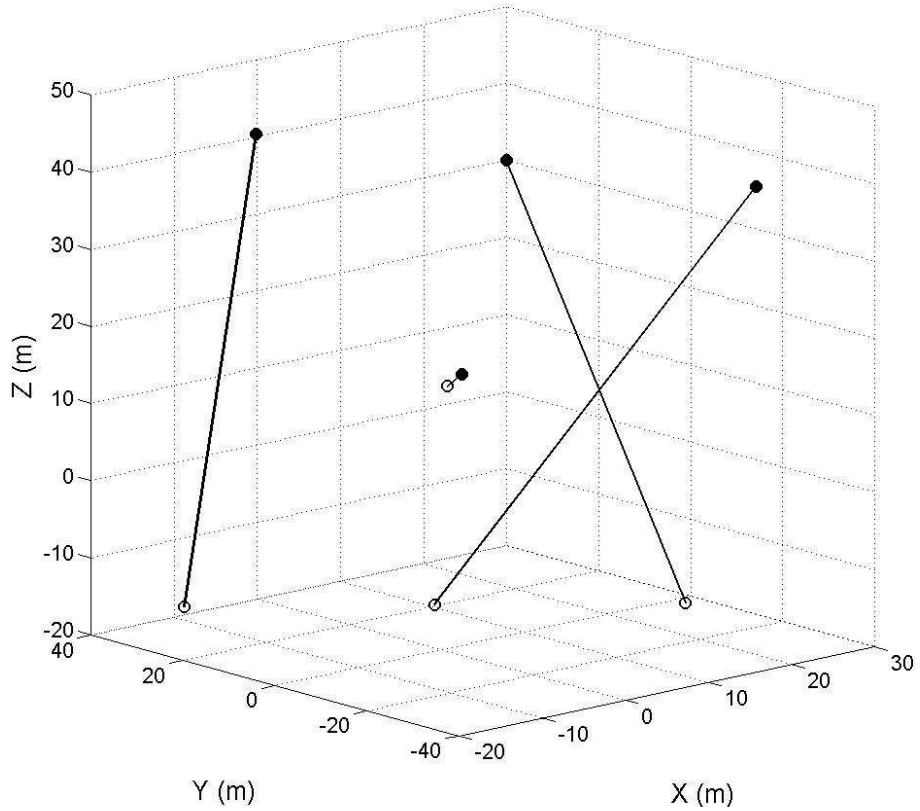

Figure 8-23 Example manoeuvre solution of the large single manoeuvre fuel consumption anomaly - smaller initial BCS fuel. With a smaller initial BCS fuel the other TSs are manoeuvres instead to balance the fuel.

For all the hundreds of manoeuvres analysed in this chapter the author has only observed this one example of the large single fuel consumption manoeuvre anomaly and so assumes that it is a rare occurrence whilst manoeuvre planning with the PAM. This assumption is strengthened by examining Figure 8-25 which shows the end of tour fuel consumption data for 10 difference tours for increasing values of μ_{PAM} for initial fuel difference case 2 (upper) and case 4 (lower). The tours were generated using the BSOM with a target tour duration of 100 days, the Year1 taskflag and 10 different tour starting stars. The markers represent the end of tour fuel consumption for the formation and the solid line represents the mean. The shapes of the mean fuel consumption lines are very similar to the single manoeuvre fuel consumption analysis data in Figure 8-13. For all the tours examined in Figure 8-25 the end of tour fuel consumption is well below the ~2.35kg of fuel used during the tour with the large single manoeuvre fuel consumption anomaly (Figure 8-20) and much more similar to the tours documented in Figure 8-17 where the anomaly is not present. This indicates that the tours in Figure 8-25 are unlikely to have exhibited the large single fuel consumption manoeuvre anomaly and emphasises the rarity of such an event whilst manoeuvre planning using the PAM.

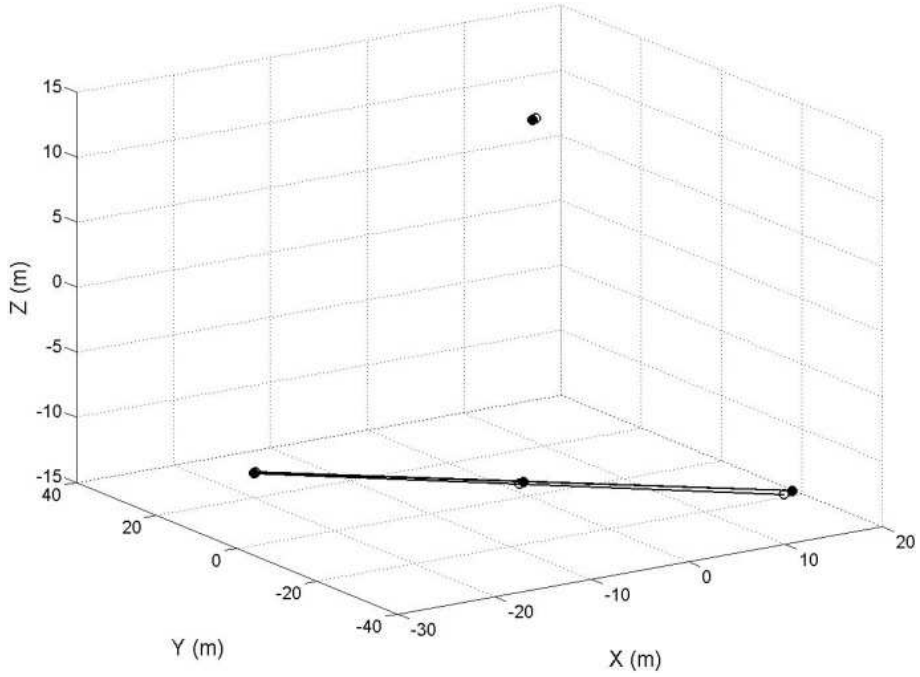

Figure 8-24 Example manoeuvre solution of the large single manoeuvre fuel consumption anomaly - ΔV balancing.

8.3.6.2 Fuel Balancing Anomaly

The fuel balancing anomaly is shown in Figure 8-26 which shows the total fuel remaining (upper) and the sum of the fuel differences (lower) for μ_{PAM} comparisons of a 133 manoeuvre tour. The tour is exactly the same to that in section 8.3.5 however due to the stochastic nature of the PAM optimisation routine the results are slightly different. The initial fuel difference is case 2, $f_i(t_0) = \{5.0, 4.7, 5.0, 5.0\} kg$. The anomaly is observed when comparing the sum of the fuel differences for the $\mu_{PAM} = 0.6$ and the $\mu_{PAM} = 10$ tours.

Contrary to the evidence in section 8.3.4 and section 8.3.5.1 where superior fuel balancing is observed for $\mu_{PAM} = 0.6$ manoeuvres and tours, Figure 8-26 (lower) shows the $\mu_{PAM} = 10$ tour with better fuel balancing than the $\mu_{PAM} = 0.6$ tour. Like the large single fuel consumption manoeuvre anomaly however this fuel balancing anomaly appears to be uncommon and easily explained.

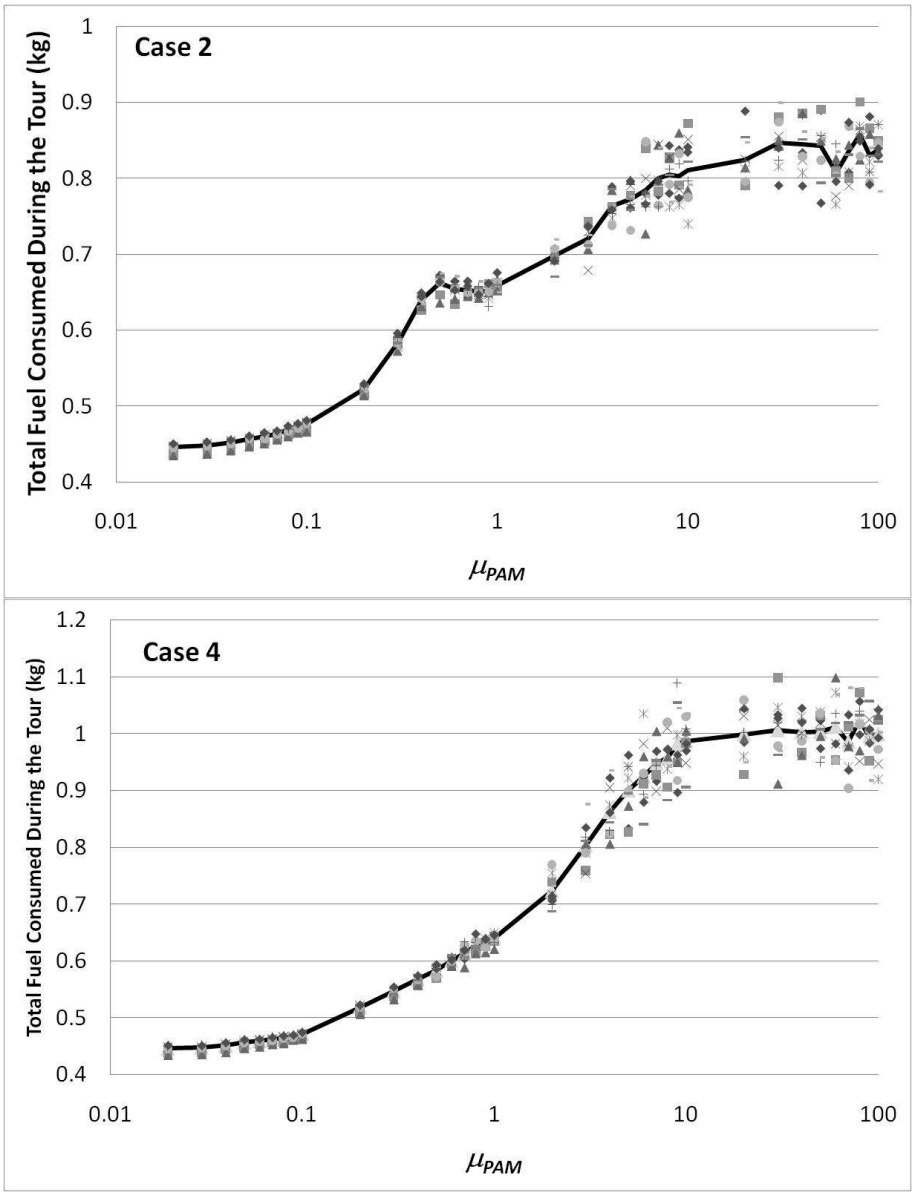

Figure 8-25 End of tour fuel consumption data for 10 separate tours for initial fuel case 2 (upper) and case 4 (lower)

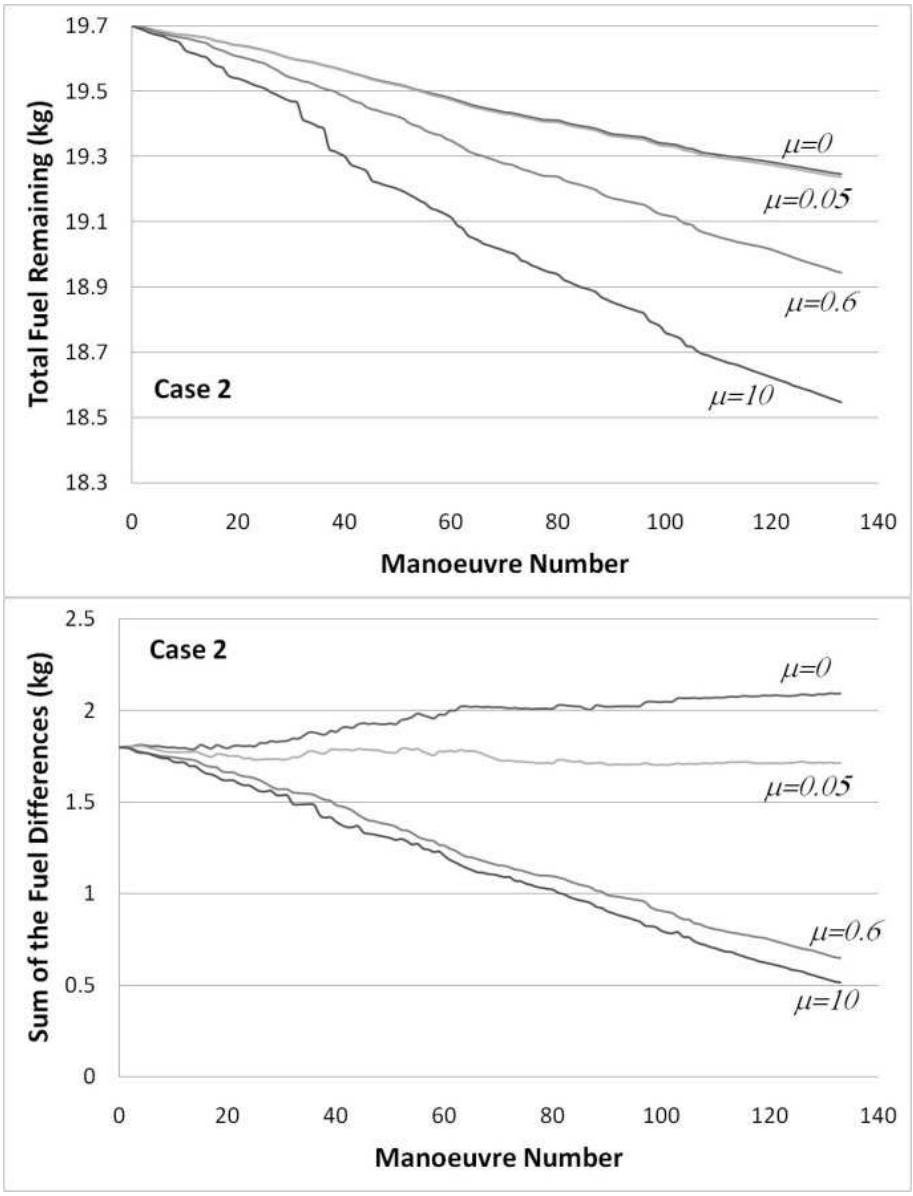

Figure 8-26 Fuel balancing anomaly. Unlike the tours in Figure 8-16 and Figure 8-17 the fuel balancing is better for when $\mu_{PAM} = 10$ than when $\mu_{PAM} = 0.6$.

Figure 8-27 shows the end of tour sum of the fuel differences for 10 different tours for increasing values of μ_{PAM} for initial fuel difference case 2 and case 4. The tours are the same tours used to generate the fuel consumption data in Figure 8-25. The markers represent the individual tours whilst the solid line represents the mean. In both charts $0.5 \leq \mu_{PAM} \leq 0.7$ represents the mean minimum fuel balancing observed. Examining the spread of values about the mean however shows that for some tours a $\mu_{PAM} > 0.7$ does provide better end of tour fuel balancing than when $0.5 \leq \mu_{PAM} \leq 0.7$. Therefore, although the mean data show that for $0.5 \leq \mu_{PAM} \leq 0.7$ superior fuel balancing can be achieved, there is a likelihood that for certain tours (defined by starting star, target tour time and taskflag selection) better fuel balancing can be found for $\mu_{PAM} \geq 0.7$. For the data set in Figure 8-27 this amounts to ~7% for case 2 and ~23% for case 4 of all tours where $\mu_{PAM} \geq 0.7$ compared with the mean $\mu_{PAM} = 0.5$ value

(for this data set $\mu_{PAM} = 0.5$ represents the mean value with the best fuel balancing). Although these percentages are significant fuel balancing should not be optimised to the detriment of fuel usage and Figure 8-25 clearly shows the fuel consumption penalties involved using high μ_{PAM} values.

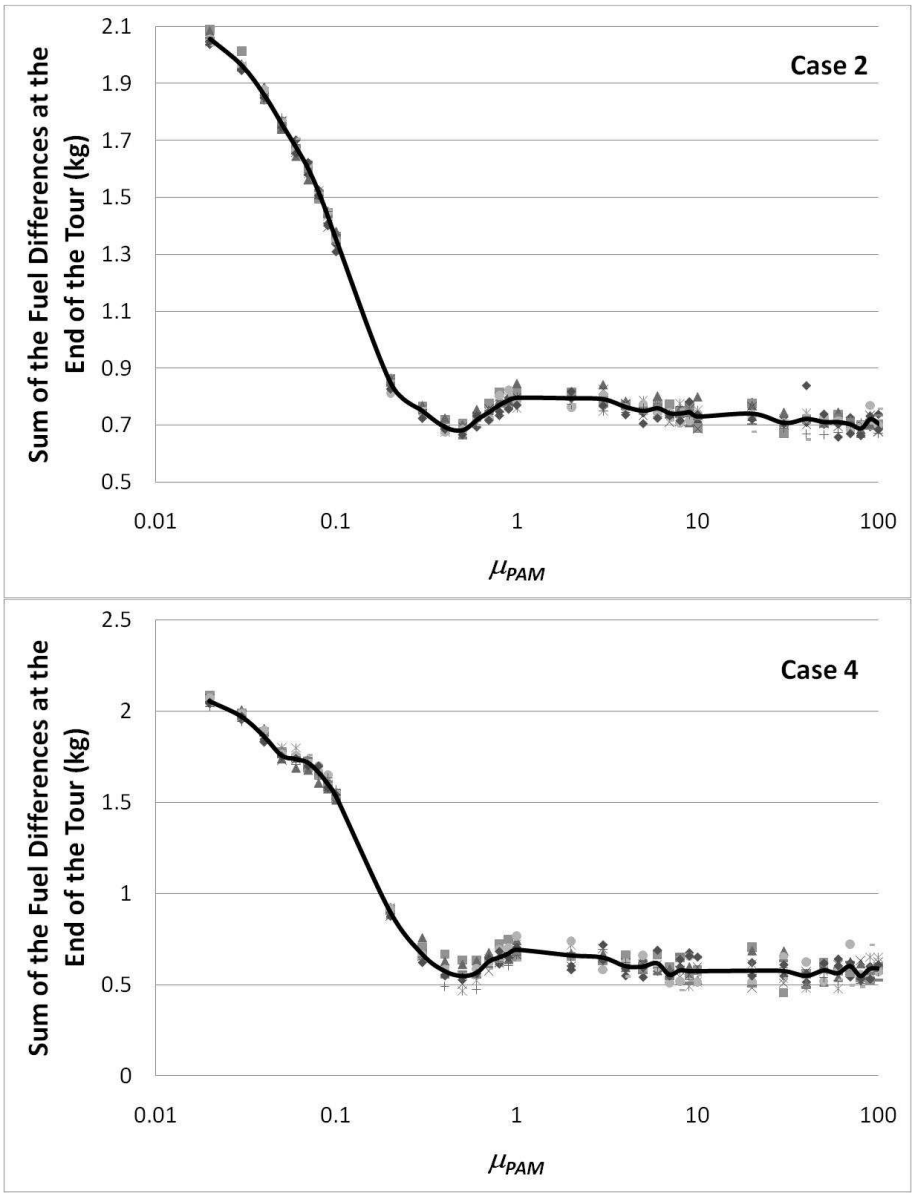

Figure 8-27 End of tour sum of the fuel differences for 10 separate tours for initial fuel balancing case 2 (left) and case 4 (right). The charts shows that the fuel difference data is very similar to that found in Figure 8-13 indicating that the fuel balancing anomaly is a rare event.

8.3.7 Position Assignment Module Analysis Conclusions

The analysis of the PAM has provided a wealth of data which is summarised here.

8.3.7.1 Single Manoeuvre Conclusions

The conclusions drawn from the analysis of single manoeuvres are as follows:

- The PAM can always achieve better performance than the benchmark algorithm whether it is for fuel minimising or for fuel balancing manoeuvres
- From the mean data, $0.5 \leq \mu_{PAM} \leq 0.7$ gives the best reduction in the sum of the fuel differences (i.e. best fuel balancing)

8.3.7.2 Tour Analysis Conclusions

The conclusions drawn from analysis if the DARWIN tour data are as follows:

- From the mean data $0.5 \leq \mu_{PAM} \leq 0.7$ gives the best reduction in the sum of the fuel differences (i.e. best fuel balancing)
- For initially balanced fuel formations $\mu_{PAM} = 0.03$ provides the best fuel balancing performance (where performance is the percentage increase in fuel balancing achieved per percentage increase in fuel consumption relative to the $\mu_{PAM} = 0$ tour).
- For initially unbalanced fuel formations $\mu_{PAM} = 0.1 - 0.2$ provides the best fuel balancing performance (where performance is the percentage increase in fuel balancing achieved per percentage increase in fuel consumption relative to the initial sum of the fuel differences).
- The PAM can sometimes plan manoeuvres that use significant amounts of fuel to achieve better fuel balancing. A check function needs to be used to reject such plans and recalculate if they occur.
- For $\mu_{PAM} \geq 0.7$ better fuel balancing can be achieved than when $\mu_{PAM} = 0.6$ however the fuel balancing performance at these high μ_{PAM} values is significantly poorer than for $\mu_{PAM} = 0.6$ and so high μ_{PAM} values should be avoided.

The analysis of the PAM has revealed a number of μ_{PAM} values that maximise the performance of the algorithm depending on the requirements of the user. These μ_{PAM} values are important in the design of a 'μ_{PAM}-selection algorithm' that is required for the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) to plan the many manoeuvres required during the mission. The 'μ_{PAM}-selection algorithm' is required for each manoeuvre plan to select the appropriate value of μ_{PAM} that helps maintain the formation fuel differences within some defined limit whilst optimising the fuel consumption to do so.

8.4 Further Work

The Position Assignment Module (PAM) introduced and analysed in this chapter uses a number of assumptions and simplifications to perform its calculations (as introduced in chapter 6). The removal of these constraints would allow the PAM to calculate manoeuvres

much more accurately and therefore would aid optimisation of the spacecraft positions but this would necessitate a change in the Manoeuvre Planning Architecture designed. Within the scope of the SepM-MPA however further improvements could be made. These are discussed in this section.

8.4.1 ΔV Optimisation

The large single manoeuvre fuel anomaly is a rare symptom of the fuel optimisation nature of the PAM. One of the solutions to this problem is identified in sub-section 8.3.6.1 and uses the individual spacecraft ΔV capability in the optimisation cost function instead of the individual spacecraft fuel. This is shown in Equation (8.30) below.

$$J_{PAM} = \min_{X} \left\{ \sum_{i}^{n} \left(\Delta V_i \left(t_0 \right) - \Delta V_i \left(t_f \right) \right) + \mu_{PAM} \sum_{i}^{n} \sum_{j:j \neq i}^{n} \left| \Delta V_i \left(t_f \right) - \Delta V_j \left(t_f \right) \right| \right\} \tag{8.30}$$

where ΔV capability is a function of thruster specific impulse, *Isp*, spacecraft mass, M_i, and fuel mass, mf_i, for spacecraft *i*:

$$\Delta V_i = g I_{sp} \ln \left(\frac{M_i}{M_i - mf_i} \right) \tag{8.31}$$

where *g* is the acceleration due to gravity at sea level.

Using the ΔV capability is a much more robust method for balancing the spacecraft manoeuvre resources across the fleet since ΔV capability is a true measure of each spacecraft's future capacity to perform manoeuvres. Comparative ΔV capability is independent of spacecraft mass and fuel remaining. This means that ΔV capability balancing manoeuvres will not incur the large single spacecraft manoeuvre fuel anomaly. Analysis of the PAM using Equation (8.30) as the cost function is an undertaking for future work as it may reveal different PAM performance data than that found for the fuel optimisation analysis.

8.4.2 Optimisation Algorithm Constraints

Another reason for the large single manoeuvre fuel anomaly was the lack of constraints given to the optimisation algorithm through the independent variable, *X*. This allowed the algorithm to position of the BCS very far away from the starting position in order to reduce the sum of the fuel differences of the formation. Adding a position constraint to X would solve this problem and could be implemented as

$$\sqrt{x_1^{\,2} + x_2^{\,2} + x_3^{\,2}} \leq 100m \tag{8.32}$$

8.4.3 Adopting a Multi-objective Approach

The selection of a single objective optimisation approach for the PAM was driven by the desire for analysis transparency and ability to tune the level of fuel balancing required by the manoeuvre. This led to what is essentially a multi-objective problem (fuel/time minimisation and fuel balancing) being reduced to a single-objective problem (the minimisation of J_{PAM} using a weighted sum objective function). The weighted sum approach will always favour either fuel consumption or fuel balancing but never a happy medium between the two. This simplification of the problem may have negatively affected the optimality of the PAM and certainly lead to the large single manoeuvre fuel consumption anomaly. It is suggested therefore to re-write the optimisation problem as a multi-objective one to examine whether any performance gains achievable can offset the increased calculation time required. Although control will be lost in specifying the level of fuel balancing to employ for each manoeuvre, the multi-objective approach solutions will incorporate some component of fuel balancing in all manoeuvres.

8.5 Chapter Summary

This chapter dealt with the requirements for the Position Assignment Module (PAM), its implementation, its output and its comparison with a benchmark algorithm (BPAM). The chapter began by introducing the requirements for formation reconfiguration manoeuvres during the reconfiguration GNC mode of the DARWIN mission. Also identified were the three optimisation goals and two constraints associated with these manoeuvres. A review of the literature related to spacecraft position assignment for formation flying reconfiguration followed highlighting the differences between the manoeuvres previously modelled and those required for the DARWIN mission.

The PAM algorithm was developed taking into account the nature of the dynamic environment and the envisaged properties of the DARWIN spacecraft. Crucially the model employs free space dynamics with point mass spacecraft and time invariant mass. The PAM's goal is to find the post-manoeuvre spacecraft positions that satisfy the configuration requirements whilst optimising the manoeuvre duration and fuel management. Simulation of these spacecraft positions is governed by a 4-element independent vector and the formation geometry guidelines. The optimisation algorithm modifies the post-manoeuvre positions to minimise a cost function. This cost function measures fuel consumption and fuel balancing for the manoeuvre. A fuel balancing weight, μ_{PAM}, is used to trade-off fuel balancing against fuel consumption for the optimisation process. In addition, straight-line trajectories between initial and final spacecraft positions using bang-bang and bang-coast-bang thrust profiles ensure a minimum formation manoeuvre duration for any set of post-manoeuvre positions. For comparison a benchmark algorithm (BPAM) was also designed. The post-manoeuvre spacecraft positions are governed by a rigid body rotation of the formation from its initial configuration and the formation geometry guidelines. Fuel usage and manoeuvre duration for the BPAM are calculated in a similar fashion to the PAM.

The comparison underlines the importance of optimally assigning end-of-manoeuvre spacecraft positions as the PAM was on average able to save ~6.5% fuel whilst fuel balancing

and increase fuel balancing by 0.35% whilst in fuel balancing mode on a manoeuvre-by-manoeuvre basis. The performance of the PAM varied depending on the manoeuvre with greater fuel savings and fuel balancing found for manoeuvres involving a formation reconfiguration (i.e. linTTN to triTTN and vice versa). The fuel balancing weight, μ_{PAM}, was extensively analysed to gain insight into a 'μ_{PAM} -selection algorithm' that would be required for the autonomous Manoeuvre Planning Architecture (MPA). This analysis revealed usable trends in the μ_{PAM} data for different initial spacecraft fuel distribution cases. In particular, optimal fuel balancing was observed for $0.5 \leq \mu_{PAM} \leq 0.7$. Whilst no explanation for this was found the phenomena remained persistent with changes in spacecraft thruster parameters and initial fuel distribution.

As well as analysing single manoeuvres the PAM was analysed for tours of manoeuvres. This analysis was performed for different values of μ_{PAM} and different initial fuel distribution cases and revealed many of the same trends found in the single manoeuvre data. Tour performance was also analysed detailing how fuel efficient the manoeuvres are with respect to increasing fuel balancing. Again, a number of usable trends were found for different initial fuel distribution cases.

During the course of the analysis a few anomalies in the data were discovered and reported. Both the large single manoeuvre fuel anomaly and the fuel balancing anomaly appeared to be isolated cases in a very large data set. Both anomalies were explained with corroborating evidence and were deemed worthy of note but were not frequent enough to warrant a re-design of the PAM.

The final part of this chapter introduced some future work concepts that could be applied to the PAM. ΔV capability optimisation is a more robust way to perform manoeuvre optimisation since ΔV capability comparison is independent of spacecraft mass and fuel remaining. The PAM analysis should be repeated using ΔV capability optimisation to see if any differences are observed with the PAMs manoeuvre planning performance.

The PAM represents one way optimising the post-manoeuvre spacecraft positions for formation flying reconfiguration manoeuvres within the calculation time restricted environment imposed on the simulation. The PAM integrates well into the Separate Modular Manoeuvre Planning Architecture (SepM-MAP) from sub-section 6.4. The following chapter describes the next optimisation module within the SepM-MPA data flow, the Trajectory Design Module (TDM).

9. TRAJECTORY DESIGN MODULE

Manoeuvre planning within the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) is a two-stage optimisation process. In the first stage the post-manoeuvre spacecraft positions are found by the Position Assignment Module (PAM) and in the second stage the trajectories for each spacecraft are found by the Trajectory Design Module (TDM). Figure 9-1 shows the position of these modules within the SepM-MPA. This chapter concerns the TDM only.

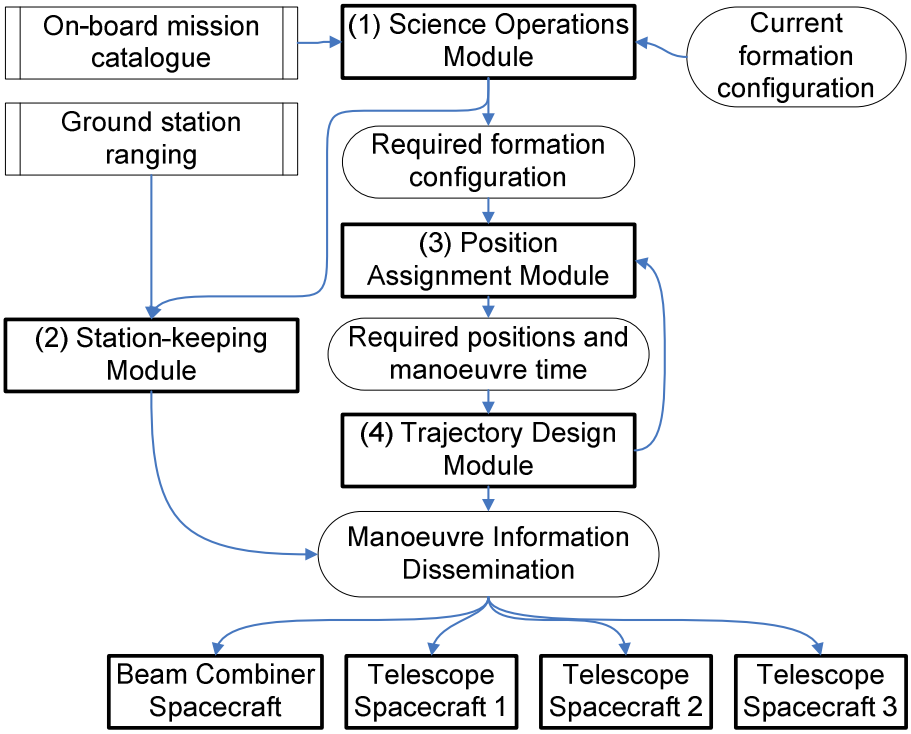

Figure 9-1 Separate Modular Manoeuvre Planning Architecture (reproduced from Figure 6-10)

As introduced in section 6.4.1 the goal of this module is to find the spacecraft trajectories that satisfy the boundary positions returned from the PAM and ensure there are no collisions or thruster plume impingement issues during the manoeuvre.

9.1 Previous Contributions

There have been many methods developed to plan trajectories for spacecraft formation flying missions. Sultan, Seereeram and Raman (2004a, 2004b, and 2004d) describe a way to calculate collision free, energy optimal reconfiguration trajectories. The spacecraft in the formation are modelled as point-masses in free space and the initial and final positions and velocities of each spacecraft are given. Spacecraft trajectories are achieved by the introduction of a sequence of way-points through which the spacecraft must travel through. The position and velocity of the way-points are the independent variables for the method and the energy optimal trajectory through the way-points is a piecewise cubic polynomial. The

solution of the optimisation problem involves a two-part ‘sequential algorithm’ using gradient-based optimisation techniques. In the first part collision-free trajectories are found and in the second part these trajectories are modified to minimise/equalise the total energy of the formation manoeuvre. The results show the method is able to quickly solve complex reconfiguration problems with as many as 20 spacecraft in the formation.

A similar method is introduced by the same authors but using a different optimisation technique to generate energy sub-optimal reconfiguration manoeuvres (Sultan, Seereeram and Raman, 2004c). The method involves first calculating the energy optimal trajectory for each spacecraft from their initial to final states (in time, position and velocity). Then for each pair of spacecraft the conflict scenario of those trajectories is analysed. If a conflict is present then each trajectory is modified by the introduction of a waypoint that one spacecraft must travel through that resolves the conflict. The trajectory that passes through the waypoint is energy optimal by parameterising it using a piecewise cubic polynomial. The position (and time) of the waypoint is approximated using linear matrix inequalities (LMI) or quadratic matrix inequalities (QMI). The introduction of the waypoint allows a new piecewise cubic polynomial trajectory to be calculated for each spacecraft. This is repeated for all spacecraft pairs to generate a set of trajectories. If the resulting trajectories contain conflicts the process can be iterated (by changing the affected spacecraft’s waypoint and/or the position of the waypoint) until a set of non-conflicting trajectories has been found. The introduction of LMI and QMI simplifies the optimisation somewhat over the method in Sultan, Seereeram and Raman (2004a).

Singh and Hadaegh (2001) present a similar method to Sultan, Seereeram and Raman (2004a). They designed ‘an optimal formation path planning approach which is suitable for implementation onboard a single spacecraft.’ Again the spacecraft are modelled as point masses in free space and the trajectories the spacecraft follow parameterised as cubic polynomials. The optimisable parameters in this method however are the coefficients of the polynomial describing the energy-minimal trajectory between the initial and final boundary conditions. A multi-step gradient-based numerical algorithm is implemented to obtain the trajectory solutions.

Richards et al. (2002) show how fuel optimal manoeuvres with collision avoidance constraints can be planned by formulating the reconfiguration problem as a simple linear program (LP). Linearised relative vehicle dynamics (Hill’s equations) are used for the spacecraft dynamics and the spacecraft initial and final state vectors (in position and velocity) are given as the boundary conditions to the optimisation to formulate a simple minimum-fuel path-planning LP. Obstacle avoidance, collision avoidance and plume avoidance are added to the simple LP through the introduction of additional logical constraints as binary variables that act as extra decision variables in the optimisation problem. The resulting mixed integer linear program (MILP) is solved by a commercial off-the-shelf (COTS) software package (CPLEX, 2008). Examples show the ability of the method to path-plan for formation flying and for manoeuvring around larger space structures. Richards, et al. (2002) also introduces a method to include the final formation state as an optimisable parameter and in Richards, et al. (2001) a similar method is outlined but with more details regarding the implementation of the plume avoidance criteria.

Optimal control theory is the tool employed by Kim, Mesbahi and Hadaegh (2003) to solve the optimal collision-free reconfiguration problem. The free-space spacecraft dynamic environment is used with the initial and final relative positions of the spacecraft given to the optimiser for a two-spacecraft formation. The reconfiguration problem is formulated as a state constrained optimal control problem with a goal to find the control forces that manoeuvre the spacecraft to their desired relative positions in a collision-free and energy minimal way. Simulations show the method works well to achieve the desired trajectories but the authors admit that increasing the number of spacecraft greatly increases the complexity of the equations to be solved. Though this method does not explicitly find the required trajectories for each spacecraft, the control forces can be used with a free-space model to generate the energy optimal, collision-free trajectories.

McQuade, Ward and McInnes (2003) use potential function theory for formation guidance. The Potential Function Method involves describing the formation state using a potential function (PF), the solution of which is zero when the formation is in the desired configuration. The PF describes the 'correctness' of the formation state. Using the PF the formation state can be modified progressively towards lower potentials until the desired final state is obtained. This is achieved by ensuring that the rate of change of the potential between two subsequent states is always negative. Potential increasing terms are added to the PF to incorporate collision avoidance, formation geometry and spacecraft attitude constraints. The method is used to demonstrate a DARWIN-like formation deployment manoeuvre at L_2 using a linearised dynamics model. The Potential Function Method presented here, though simple and effective, does not incorporate any fuel or manoeuvre time optimisation. In addition its reliance on the behavioural control co-ordination approach makes it unsuitable for the TDM.

In a comparable technique to the Potential Function Method, Equilibrium Shaping (ES) has also been applied to formation flying reconfiguration manoeuvres (Pettazzi, Izzo and Theil, 2006). ES is a behavioural-based approach in which each spacecraft follows a desired velocity vector that is a sum of three velocity vectors that represent different behaviours the spacecraft is to exhibit; gather (towards a desired target), dock (at a desired target) and avoid (collisions with other satellites). A feedback control law is used to ensure the spacecraft follow the desired motion. This method allows the formation to autonomously perform reconfiguration manoeuvres in a fully decentralised way with no requirement for an inter-satellite communications link. Pettazzi, Izzo and Theil (2006) use Equilibrium Shaping to execute formation reconfiguration manoeuvres on a swarm of coulomb satellites in a relative dynamic environment characterised by the Clohessy-Wiltshire equations. Though effective and simple for generating collision-free trajectories Pettazzi, Izzo and Theil (2006) show no incorporation of energy or manoeuvre time optimisation in their work. In addition its reliance on the behavioural control co-ordination approach makes it unsuitable for the TDM.

Genetic algorithm techniques are used by Seereeram, et al. (2000) to calculate fuel-optimal collision free trajectories for multi-spacecraft reconfiguration manoeuvres. The spacecraft are governed by free-space dynamics with initial and final spacecraft positions given to the optimiser. The trajectories are parameterised by line-of-sight (LOS) and collision avoidance (CA) velocity components. The LOS is the straight-line trajectory from initial to final position and the velocity component is implemented using a bang-coast-bang thrust profile (this is analogous to the method in Beard and Hadaegh (1999)). The CA component is

added to the LOS component, perpendicular to the LOS, using a bang-bang thrust profile to accommodate collision avoidance. The objective function contains a collision avoidance, path length, execution time, fuel minimisation and fuel balancing terms all scaled by variable objective weights. The optimisation goal is to find values for the LOS and CA velocity components that minimise the objective function. A genetic algorithm is used to solve this optimisation problem and the results show the effectiveness of the method for reconfiguration manoeuvres for up to five spacecraft. The paper also introduces the concept of using genetic algorithms and Pareto-optimality to combine multi-objective manoeuvre planning and task planning.

9.2 Trajectory Design Module

The approaches introduced in the previous sub-section could all be applied to the problem faced by the TDM. The approach presented in this chapter however is inspired by that in Seereeram, et al. (2000) and adapted to include the complexities of the DARWIN mission concept. Major differences of this approach to that published by Seereeram, et al. (2000) are highlighted in sub-section 9.2.4.

9.2.1 Model Definition

9.2.1.1 Spatial Geometry

Let F_0 be an co-ordinate frame with orthonormal basis vectors $\{\mathbf{i_0}, \mathbf{j_0}, \mathbf{k_0}\}$, where $\mathbf{i_0}$ lies in the ecliptic plane and points towards the J2000.0 vernal equinox, $\mathbf{j_0}$ lies in the ecliptic plane normal to $\mathbf{i_0}$ such that $\mathbf{j_0} = \mathbf{k_0} \times \mathbf{i_0}$ and $\mathbf{k_0}$ lies normal to the ecliptic plane in the direction of the angular momentum vector of the orbit of the Earth around the Sun. F_0 is designated the inertial reference frame. Additionally let $\mathbf{r_{0,i}}$ and $\mathbf{r_{f,i}}$ denote the initial and final position vectors of spacecraft i in F_0 respectively, where $i=\{BCS,TS1,TS2,TS3\}$ and the spacecraft trajectory defined as

$$\mathbf{r_{traj,i}} = \mathbf{r_{f,i}} - \mathbf{r_{0,i}} \tag{9.1}$$

This can be seen more clearly in Figure 9-2.

9.2.1.2 Model Assumptions

The following assumptions have been made to aid model simplification:

- The formation is in free space.
- Each spacecraft is modelled as a point mass
- Each spacecraft has mass that is time-invariant.
- The position of each spacecraft can be determined.

These assumptions are the same ones used for the Position Assignment Module (PAM) and the reasoning behind them has been covered in sub-section 8.2.1.3.

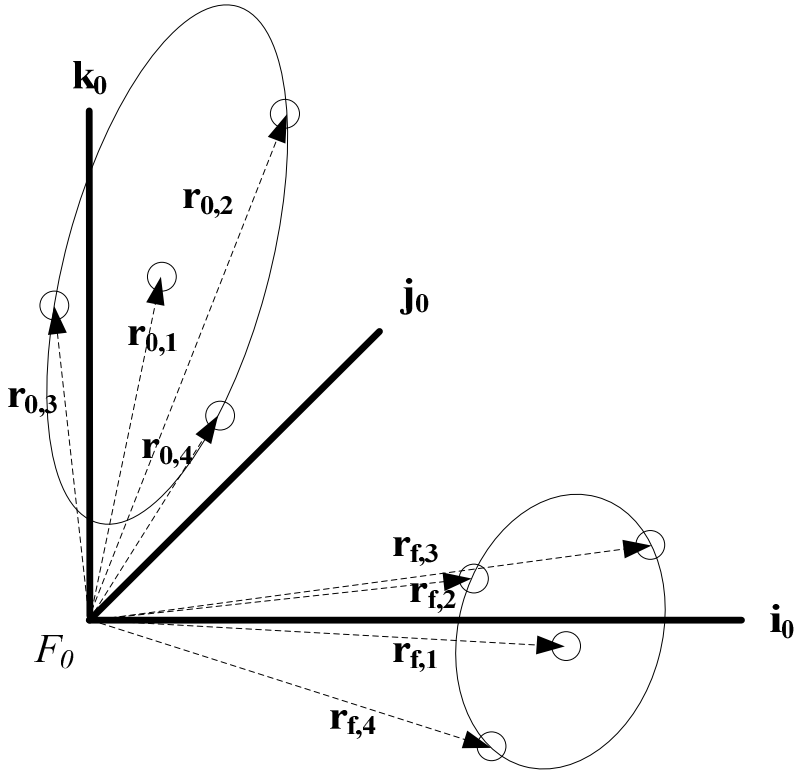

Figure 9-2 Spatial geometry of the TDM model with initial formation in the triTTN configuration and final formation in the linTTN

9.2.1.3 Translational and Fuel Dynamics

From these assumptions the translational and fuel dynamics for each spacecraft are given by:

$$m_i \ddot{\mathbf{r}}_i = \begin{cases} \alpha_i \mathbf{u_i}; & mf_i(t) > 0 \\ 0; & otherwise \end{cases}$$
$$\dot{mf_i} = \begin{cases} -\gamma \alpha_i; & mf_i(t) > 0, |\mathbf{u_i}| \neq 0 \\ 0; & otherwise \end{cases} \tag{9.2}$$

where m_i is the mass (in kg) of the ith spacecraft, $\ddot{\mathbf{r}}_\mathbf{i}$ is the acceleration vector in F_0 (ms^{-1}), α_i is the thrust saturation limit (N), $\mathbf{u_i}$ is a unit force vector (N), $\dot{f}_i$ is the fuel consumption rate (kgs^{-1}), mf_i is the amount of fuel (kg) and γ is a proportionality constant

$$\gamma = 1/I_{sp}g \tag{9.3}$$

where I_{sp} is the specific impulse of the thruster and g is the acceleration due to gravity at sea level.

9.2.1.4 Spacecraft Geometry

The spacecraft geometry is taken from Ankersen (2003), D'Arcio (2005) and Karlsson, et al. (2004). Let F_c be an co-ordinate frame with orthonormal basis vectors $\{\mathbf{i_c}, \mathbf{j_c}, \mathbf{k_c}\}$, where $\mathbf{i_c}$ points along the bore sight of the TS towards the BCS and along the bore sight of the BCS towards TS1, $\mathbf{j_c}$ lies normal to $\mathbf{i_c}$ such that $\mathbf{j_c} = \mathbf{k_c} \times \mathbf{i_c}$ and $\mathbf{k_c}$ points along the bore sight of the telescope for the TS and normal to the solar panel of the BCS. F_c is designated the spacecraft reference frame as shown in Figure 9-3.

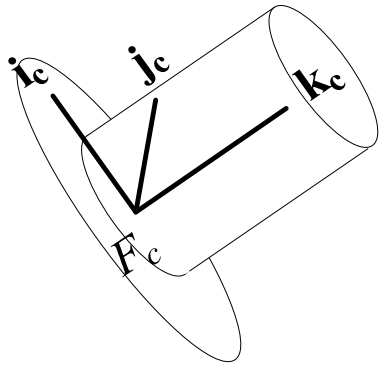

Figure 9-3 Spacecraft geometry used in the TDM

9.2.2 Trajectory Design Module Optimisation

9.2.2.1 TDM Method

The positions generated by the PAM are based on an optimisation routine that assumes straight-line bang-coast-bang (or bang-bang) trajectories. Although these trajectories are optimal they may break collision avoidance or thruster plume avoidance constraints. If they do then they will need to be modified to ensure the safety of the formation manoeuvre. The goal of the TDM is to find the optimal safe trajectories from each spacecraft's initial to final positions. This is achieved through the method that follows.

The TDM uses the time-optimal straight-line trajectory calculated by the PAM as a nominal trajectory that each spacecraft has to follow. An assessment of each trajectory is made to check for collision or thruster plume constraint violations. If there are none then the TDM simply passes this nominal trajectory for Manoeuvre Information Dissemination in the Separate Modular Manoeuvre Planning Architecture (SepM-MPA). If proximity violations are present, the TDM calculates and optimises new trajectories that are collision and thruster plume impingement free. To achieve these modified trajectories thrust components are added so that the spacecraft depart from their nominal trajectories to follow the modified ones.

The PAM uses a bang-coast-bang (or bang-bang) thrust profile to generate the nominal trajectory and this thrust occurs parallel to the direction of the trajectory. This is the parallel thrust component, $\mathbf{T_{par}}$. To achieve a departure from the nominal trajectory the TDM

implements a thrust component normal to the nominal trajectory. This is the perpendicular thrust component, $\mathbf{T_{per}}$. The perpendicular thrust follows two complementary bang-bang thrust profiles so that each spacecraft temporarily leaves the nominal trajectory for a set period of time before returning to it when the constraint violations will no longer occur. The manoeuvre time remains fixed, as optimised by the PAM, so that each spacecraft reaches their final positions at the same time.

The thrust profiles used by the TDM can be seen in Figure 9-4 where $|\mathbf{T_{per}}|$ is the magnitude of the perpendicular thrust component, $t_{\omega per}$ is the pulse width of the perpendicular thrust component and t_{per} is the execution time of the perpendicular thrust. The spacecraft and thrust geometry can be seen in Figure 9-5 where $\mathbf{T_{par}}$ is the parallel thrust component, $\mathbf{T_{per}}$ is the perpendicular thrust component and θ_T is an angle from an arbitrary reference axis to $\mathbf{T_{per}}$. The magnitude of θ_T determines the direction of $\mathbf{T_{per}}$ whilst the magnitude of $\mathbf{T_{per}}$ determines how far from the nominal trajectory the spacecraft travels.

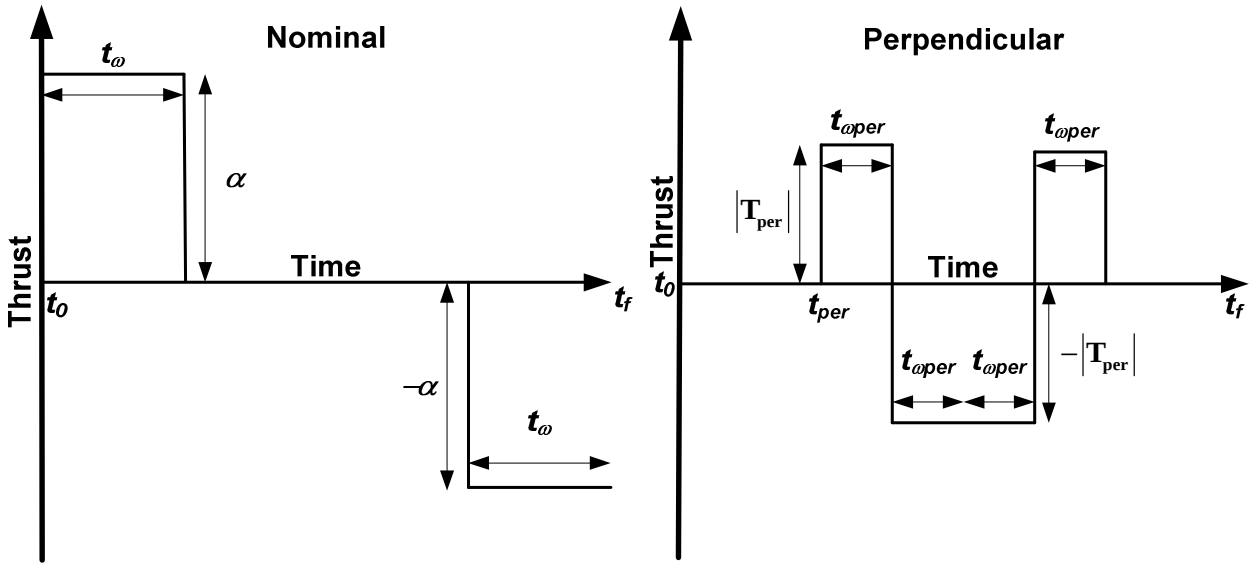

Figure 9-4 Nominal (left) and perpendicular (right) thrust profiles used by the TDM for calculating each spacecraft's trajectory

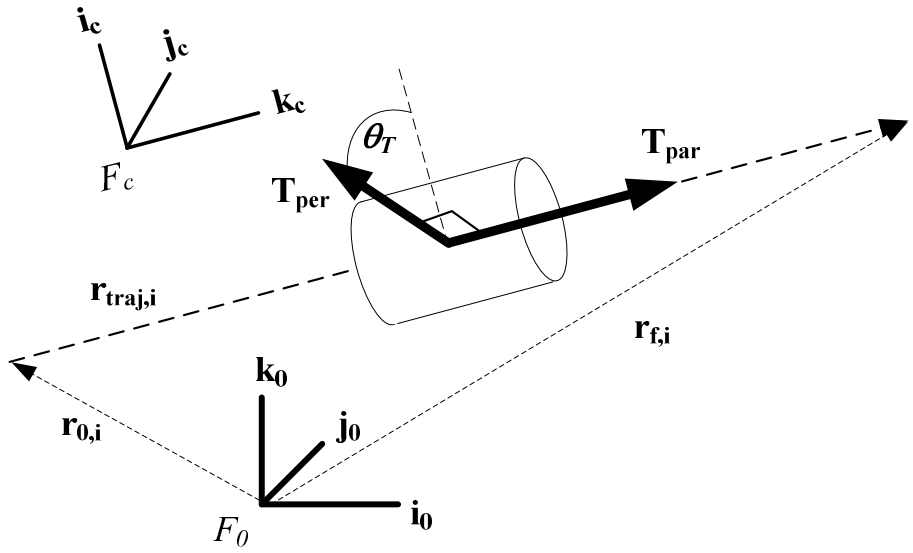

Figure 9-5 Spacecraft and thrust geometry used in the TDM

There is no need to include time optimisation in the TDM since the manoeuvre time remains fixed by the PAM. Fuel minimising/balancing is the only concern after the avoidance criteria are met. The independent variable contains 4 elements per spacecraft:

- $\theta_{T,i}$ - the direction of the perpendicular thrust component. This allows the spacecraft to depart from the nominal trajectory in any direction normal to it. $\theta_{T,i}$ can therefore have a range $-\pi \le \theta_{T,i} \le \pi$.
- $\left|\mathbf{T}_{\mathbf{per,i}}\right|$ - the magnitude of the perpendicular thrust component. This affects how far from the nominal trajectory the spacecraft travels from. $\left|\mathbf{T}_{\mathbf{per,i}}\right|$ can have the range $0 \le \left|\mathbf{T}_{\mathbf{per,i}}\right| \le \alpha_{per,i}$, where $\alpha_{per,i}$ is the perpendicular thrust saturation limit for the *i*th spacecraft.
- $t_{\omega per,i}$ - the pulse width of perpendicular thrust component. This determines how long the spacecraft departs from the nominal trajectory. Since four thrust pulses are required by the perpendicular thrust profile the range of values for $t_{\omega per,i}$ is $0 \le t_{\omega per,i} \le \frac{t_f}{4}$, where t_f is the manoeuvre duration.
- $t_{per,i}$ - the execution time of perpendicular thrust component. This is when the first perpendicular thrust pulse fires and affects where along the nominal trajectory the avoidance manoeuvre takes place. Since four thrust pulses are required by the perpendicular thrust profile the range of values for $t_{per,i}$ is $t_0 \le t_{per,i} \le \left(t_f - 4t_{\omega per,i}\right)$, where t_0 is the manoeuvre start time.

These four components allow the avoidance manoeuvre to occur in any direction (perpendicular to the nominal trajectory) and almost anywhere along the nominal trajectory maximising the optimisation routine's ability to optimise the fuel minimisation/balancing requirements.

The cost of the manoeuvre is calculated using the following cost function

$$J_2 = \min_{Y}\left\{\sum_{i}^{n}\sum_{j:j\neq i}^{n} J_c(i,j) + \sum_{i}^{n}\sum_{j:j\neq i}^{n} J_p(i,j) + J_1\right\} \tag{9.4}$$

where $J_c(i,j)$ is the collision cost, Equation (9.6), $J_p(i,j)$ is the plume impingement cost between the ith and jth spacecraft, Equation (9.7), J_1 is the fuel cost function used by the PAM, Equation (9.8), and Y is the independent variable, Equation (9.5).

$$Y = \left\{\theta_{T,i}, \left|\mathbf{T_{per,i}}\right|, t_{\omega per,i}, t_{per,i}\right\} : i = \{1 \ldots N\} \tag{9.5}$$

$$J_C(i,j) = \begin{cases} \infty; & \left|\mathbf{r_i}(t) - \mathbf{r_j}(t)\right| \leq 10m \\ 0; & otherwise \end{cases} \tag{9.6}$$

$$J_P(i,j) = \begin{cases} \infty; & \sin^{-1}\left(\dfrac{\left|\mathbf{r_i}(t)\times\mathbf{r_j}^{nom}(t)\right|}{\left|\mathbf{r}_i(t)\right|\left|\mathbf{r_j}^{nom}(t)\right|}\right) \leq 5^\circ \\ \infty; & \sin^{-1}\left(\dfrac{\left|\mathbf{r_i}(t)\times\mathbf{r_j}^{per}(t)\right|}{\left|\mathbf{r}_i(t)\right|\left|\mathbf{r_j}^{per}(t)\right|}\right) \leq 5^\circ \\ 0; & otherwise \end{cases} \tag{9.7}$$

$$J_1 = \left\{\sum_{i}^{n}\left(mf_i(t_0) - mf_i(t_f)\right) + \mu_{TDM}\sum_{i}^{n}\sum_{j:j\neq i}^{n}\left|mf_i(t_f) - mf_j(t_f)\right|\right\} \tag{9.8}$$

For Equations (9.5)-(9.8), $\mathbf{r_i}(t)$ is the position vector of spacecraft i in F_0 at time t, $\mathbf{r_j}^{\mathbf{nom}}(t)$ is the nominal thrust plume cone direction vector for spacecraft j in F_0 at time t, $\mathbf{r_j}^{\mathbf{per}}(t)$ is the perpendicular thrust plume cone direction vector for spacecraft j in F_0 at time t and N is the number of spacecraft in the formation. Fuel minimising/balancing optimisation is achieved through the modification of μ_{TDM} as occurs in the PAM. In Equation (9.8) the final fuel amount $mf_i(t_f)$ can be found using

$$mf_i(t_f) = mf_i(t_0) - mF_{par,i} - mF_{per,i} \tag{9.9}$$

where $mf_i(t_0)$ is the initial fuel amount of the ith spacecraft, $mF_{par,i}$ is the fuel consumed by the parallel thrust profile (calculated initially by the PAM, Equation (8.20)) and $mF_{per,i}$ is the fuel consumed by the perpendicular thrust profile, Equation (9.11).

$$mF_{par,i} = 2t_{\omega i}\gamma\alpha_i \tag{9.10}$$

$$mF_{per,i} = 4t_{\omega per,i}\gamma\alpha_{per,i} \tag{9.11}$$

Minimising Equation (9.4) with respect to independent variable, *Y*, provides the collision free plume avoiding trajectories for time minimal and fuel optimal/balanced formation reconfiguration manoeuvres for the model defined.

9.2.3 Trajectory Design Module Algorithm

The optimisation procedure is summarised in the flow chart in Figure 9-6. The optimisation routine starts with the output spacecraft initial and final positions and manoeuvre duration from the PAM. Also an initial estimate for the independent vector, Y_0, is set. Next, the positions are tested for proximity violations. This is achieved through generating the spacecraft trajectories and then assessing these trajectories segment-by-segment to check for proximity violations using Equations (9.6) and (9.7). One hundred segments per trajectory are examined. If proximity violations are detected at this point then the independent variable, *Y*, needs to be rejected. A new *Y* is generated and the resulting trajectories checked for proximity violations. When no proximity violations are detected for *Y* then the cost for the manoeuvre is calculated using Equation (9.8). This cost, and it's associated *Y*, is stored in the optimisation algorithm's database and, assuming the algorithm termination conditions have not been met, *Y* is updated and the optimisation process iterates. When the termination conditions for the optimisation algorithm have been met the routine outputs the independent variable *Y* with the lowest manoeuvre cost associated to it. Using this *Y* the spacecraft trajectories are re-calculated and then passed to the MIDM.

9.2.4 Differences of the TDM to Seereeram's Method

Although the paper by Seereeram, et al. (2000), hereafter abbreviated to SEA, has been the inspiration for the development of the TDM, there are a number of significant differences that have been made to specifically accommodate the constraints of the DARWIN mission and improve the optimisation flexibility. The main similarity between the TDM method and the SEA method is the use of the perpendicular thrust component to execute an avoidance manoeuvre. However it is the implementation of this thrust component that differs between the two methods.

9.2.4.1 Perpendicular Thrust/Velocity Components

In SEA the avoidance manoeuvre is given by an optimisable velocity component normal to the nominal (line-of-sight) trajectory. This velocity component is a 3D vector for each spacecraft defining the direction and magnitude of the spacecraft's deviation from the nominal trajectory. The velocity component is executed at t_0 using two complementary bang-bang velocity profiles so that the spacecraft deviates from the nominal trajectory immediately and does not return until t_f. Optimising the perpendicular velocity component for four spacecraft requires twelve optimisable parameters.

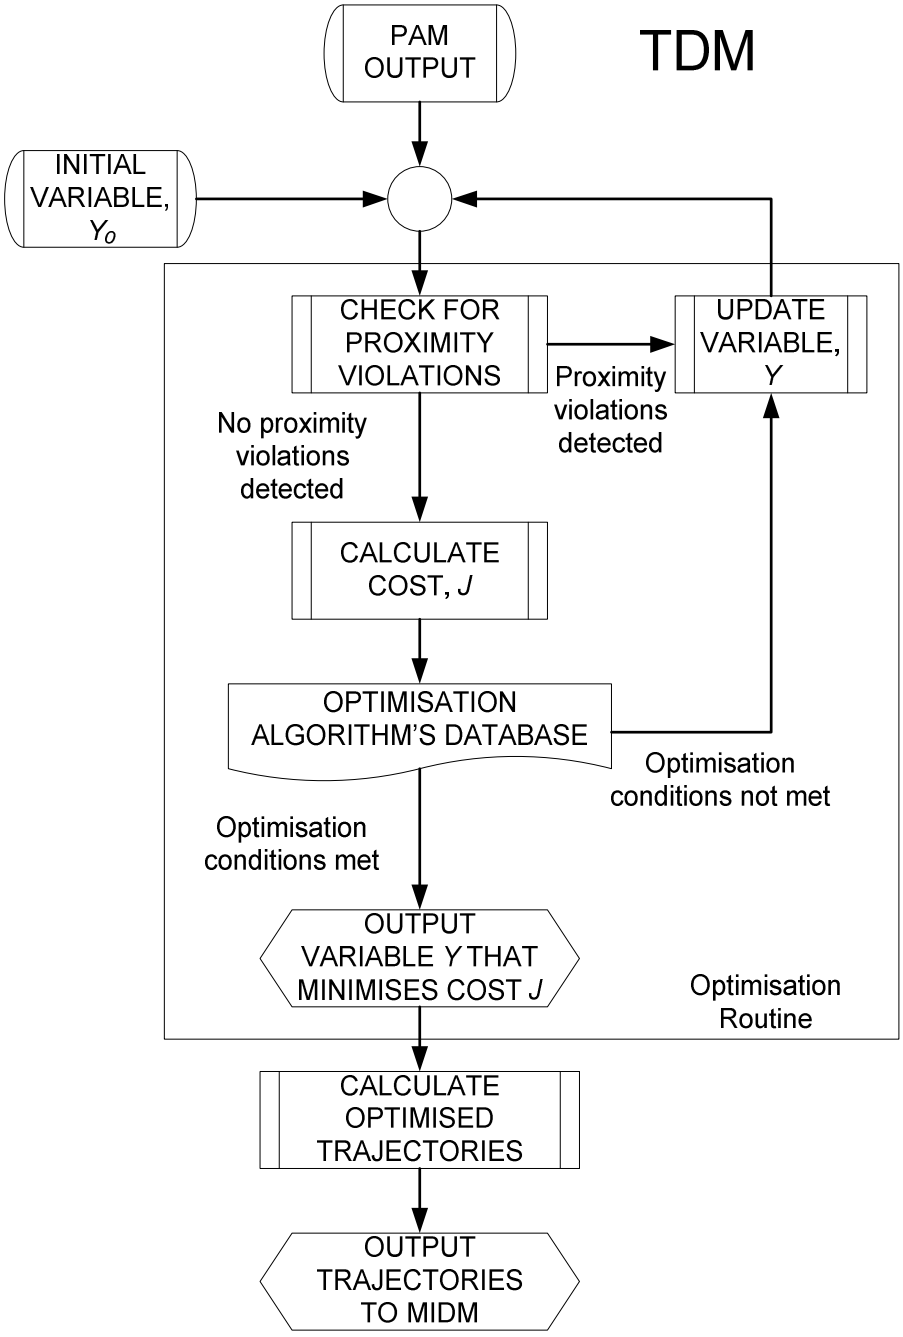

Figure 9-6 Data flow schematic for the TDM

The TDM method allows a higher degree of flexibility to be given to the optimisation algorithm than the SEA method. The direction and magnitude of the perpendicular thrust components are given by an angle, $\theta_{T,i}$ and a scalar, $\left|\mathbf{T}_{\mathbf{per,i}}\right|$ using only eight optimisable parameters for a four spacecraft formation. In addition, the timing of the perpendicular thrust is not restricted to an execution time of t_0 and a pulse width of $t_f/4$. This addition gives the optimisation algorithm more flexibility to optimise the fuel consumption/balancing but at a cost of requiring the number of optimisable parameters to increase from 12 in SEA to sixteen for the TDM.

9.2.4.2 Nominal Thrust/Velocity Components

In SEA the nominal velocity component is also an optimisable parameter. This not only increases the number of optimisable parameters of the objective function by twelve but also requires an addition term in the cost function itself, an 'execution time penalty'. Since the nominal velocity component determines the duration time of each spacecraft's manoeuvre this additional term is required to minimise this manoeuvre time. In addition, the individual spacecraft manoeuvre times are independent of each other and so the total time for the formation to perform the manoeuvre is much more dynamic.

In the TDM method, the nominal thrust component is fixed by the PAM and so the total formation manoeuvre duration is fixed, forcing the manoeuvre to be executed as quickly as possible. This is the most desirable situation since manoeuvre duration optimisation is a key requirement for the manoeuvre planning algorithm. Due to this requirement no extra optimisable parameters are required leaving the number of parameters for the TDM constant at sixteen. Though no direct comparison has been made the SEA method uses 24 optimisable parameters for an undesirably more time flexible manoeuvre and will therefore require more time and processing power to reach a solution.

9.3 Analysis

9.3.1 Analysis Setup

For all the manoeuvres generated in this section the following initial parameters were used:

$$
\begin{aligned}
i &= \{BCS, TS1, TS2, TS3\} \\
m_i &= \{1100, 900, 900, 900\}\, kg \\
\alpha_i &= \{6.0,\ 6.0,\ 6.0,\ 6.0\}\, mN \\
I_{sp} &= 3300s \\
\alpha_{per,i} &= \{10, 10, 10, 10\}\, mN
\end{aligned}
$$

The mass data is representative of the estimated masses of each spacecraft type in the DARWIN mission (Karlsson, et al. 2004) and the thruster saturation data and thruster specific impulse describe a thruster configuration using the RIT-10 FEEP thruster (D'Arcio, 2005). The TDM routine was written and executed within MATLAB® and the initial optimisation performed using the Mesh Adaptive Direct Search (MADS) – Pattern Search (PS) algorithm found within the Genetic Algorithm and Direct Search Toolbox (GADS).

9.3.2 Example Manoeuvre

The example TDM manoeuvre uses the following initial and final positions in F_0:

$$\mathbf{r_0} = \begin{bmatrix} 50 & 0 & 50 \\ 45 & 45 & 0 \\ 0 & -35 & -35 \\ 0 & 0 & -100 \end{bmatrix}, \quad \mathbf{r_F} = \begin{bmatrix} -50 & 0 & -50 \\ -45 & -45 & 0 \\ 0 & 35 & 35 \\ 0 & 0 & 100 \end{bmatrix} \tag{9.12}$$

The nominal manoeuvre to attain these final positions from these initial positions involve all the spacecraft colliding at point (0,0,0) halfway through the manoeuvre (as in Figure 9-7 (upper)). The resulting manoeuvre calculated by the TDM to include collision avoidance of 10m and plume avoidance of 5º is shown in Figure 9-7 (lower). This manoeuvre is performed with $\mu_{TDM} = 0$ and $mf_i(t_0) = \{4.8, 4.7, 4.9, 4.85\}\, kg$. In the trajectory plots the spacecraft initial positions are denoted by open circles (○) and the final positions by closed circles (●). The manoeuvre shows the flexibility the input parameter *Y* gives to the optimisation algorithm. During the manoeuvre all the spacecraft perform different sized avoidance manoeuvres of differing lengths and at different times with the net result that all the spacecraft achieve their desired positions without violating any of the avoidance constraints.

9.3.2.1 Spacecraft Separations in the Example Manoeuvre

The spacecraft separations for this manoeuvre are shown in Figure 9-8 plotted using 100 time steps. The upper chart shows the separations when no avoidance is activated (Figure 9-7 (upper) trajectories) and the lower chart show the separations with avoidance activated (Figure 9-7 (lower) trajectories). The lower chart clearly shows that none of the spacecraft break the 10m collision violation. The chart also emphasises the fuel optimisation at work since all the spacecraft separations become less than 15m sometime during the manoeuvre with the closest separation just tenths of a millimetre greater than the 10m limit. This shows the optimiser is trying to reduce fuel consumption by getting the spacecraft as close as possible without violating the collision rules.

9.3.2.2 Thrust Plume Separations in the Example Manoeuvre

The separation angles of each spacecraft's nominal (Nom) and perpendicular (Per) thrust plumes and the other spacecraft in the formation are given in Figure 9-9. The individual charts show the separation angles between the spacecraft's thrust plumes and the other spacecraft in the formation for the BCS (upper left), TS1 (upper right), TS2 (lower left) and TS3 (lower right). The four charts in Figure 9-9 clearly show the bang-coast-bang nature for the nominal thrust. For the BCS, TS1 and TS2 the nominal thrust is not activated during the middle of the manoeuvre and so the nominal thrust separation angles are not present. For TS3 however the nominal thrust is following a bang-bang profile since this spacecraft is the one that takes longest to complete its manoeuvre. For TS3 the nominal thrust separation angles remain for the entire manoeuvre duration.

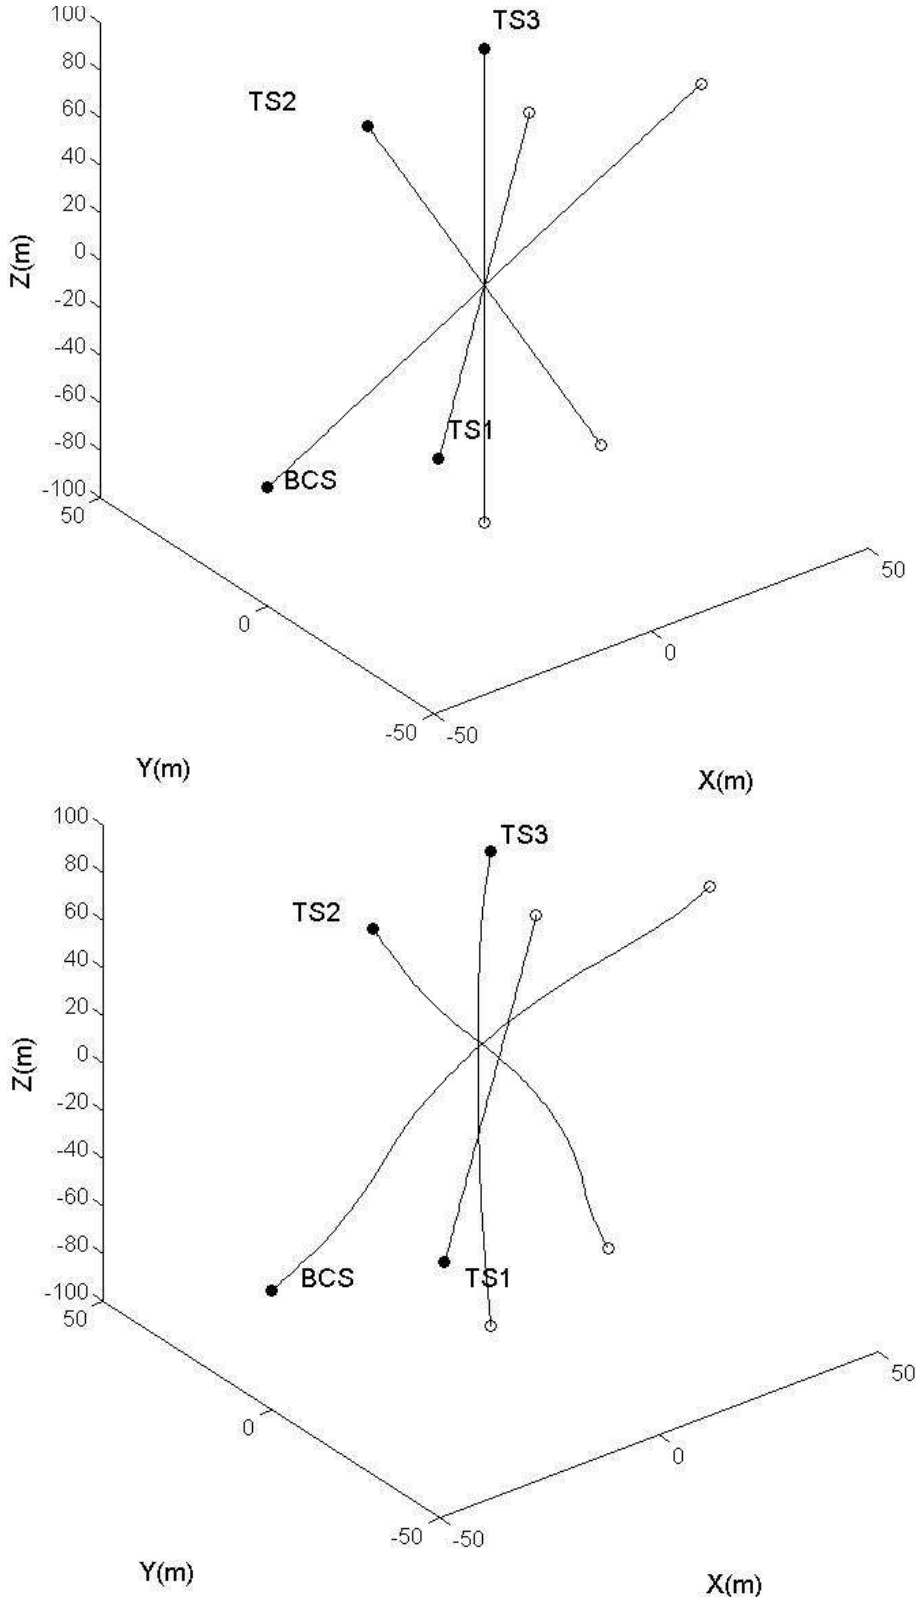

Figure 9-7 TDM example manoeuvre. No avoidance activated (upper) and avoidance activated (lower). The lower chart shows the avoidance trajectories taken by each spacecraft to avoid the planned collision demonstrated in the upper chart.

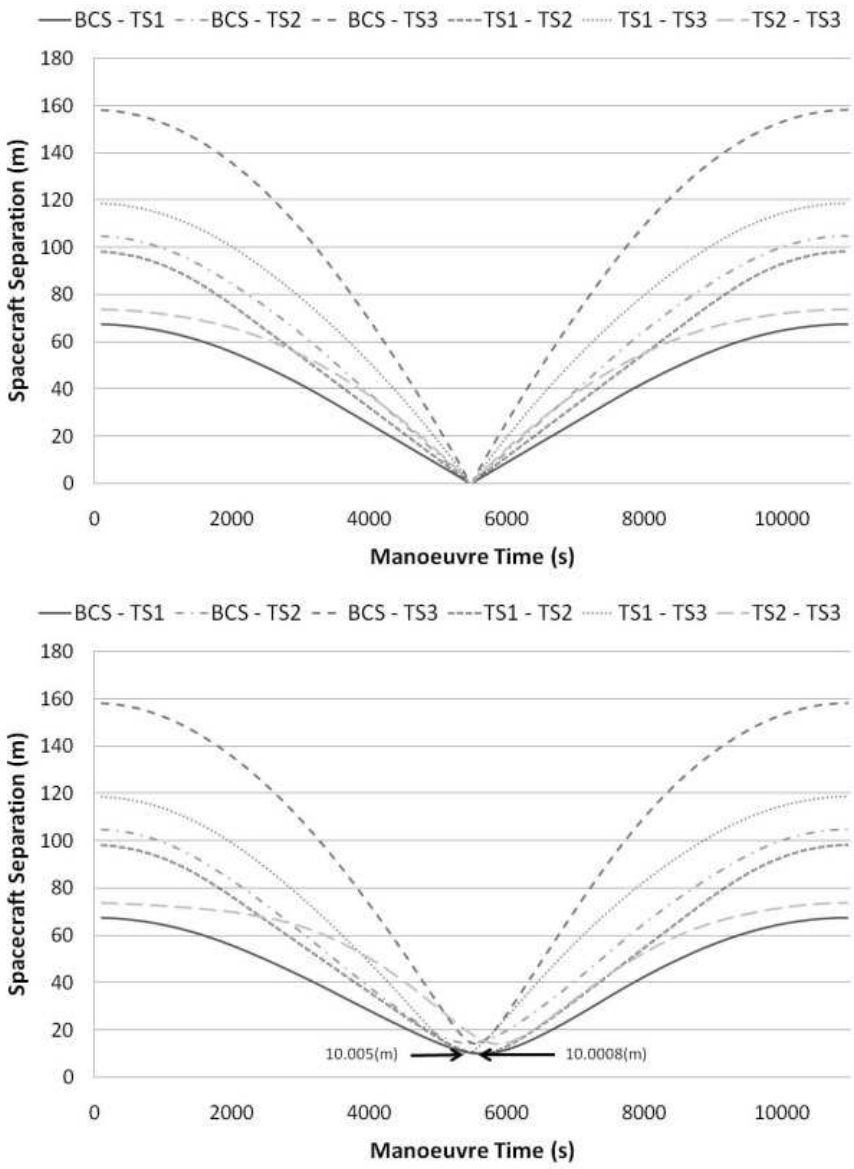

Figure 9-8 Spacecraft separations for the TDM example manoeuvre. No avoidance activated (upper) and avoidance activated (lower). The lower chart shows how close the spacecraft separations come to violating the 10m limit indicating that the TDM is good at optimising for fuel minimal trajectories.

The perpendicular thrusts that generate the avoidance manoeuvres can also be seen in Figure 9-9. Here the thrusts are following two complimentary bang-bang profiles. Large changes in the perpendicular thrust separation angles can be observed on all the charts. This is due to the perpendicular thrust direction changing by 180° at times $t_{\omega per}$ and 3 $t_{\omega per}$. For all the 24 possible plume impingement violations, Figure 9-9 shows that none of them occur and the nearest any thrust plume gets to hitting another spacecraft is ~11° (TS3-BCS nominal).

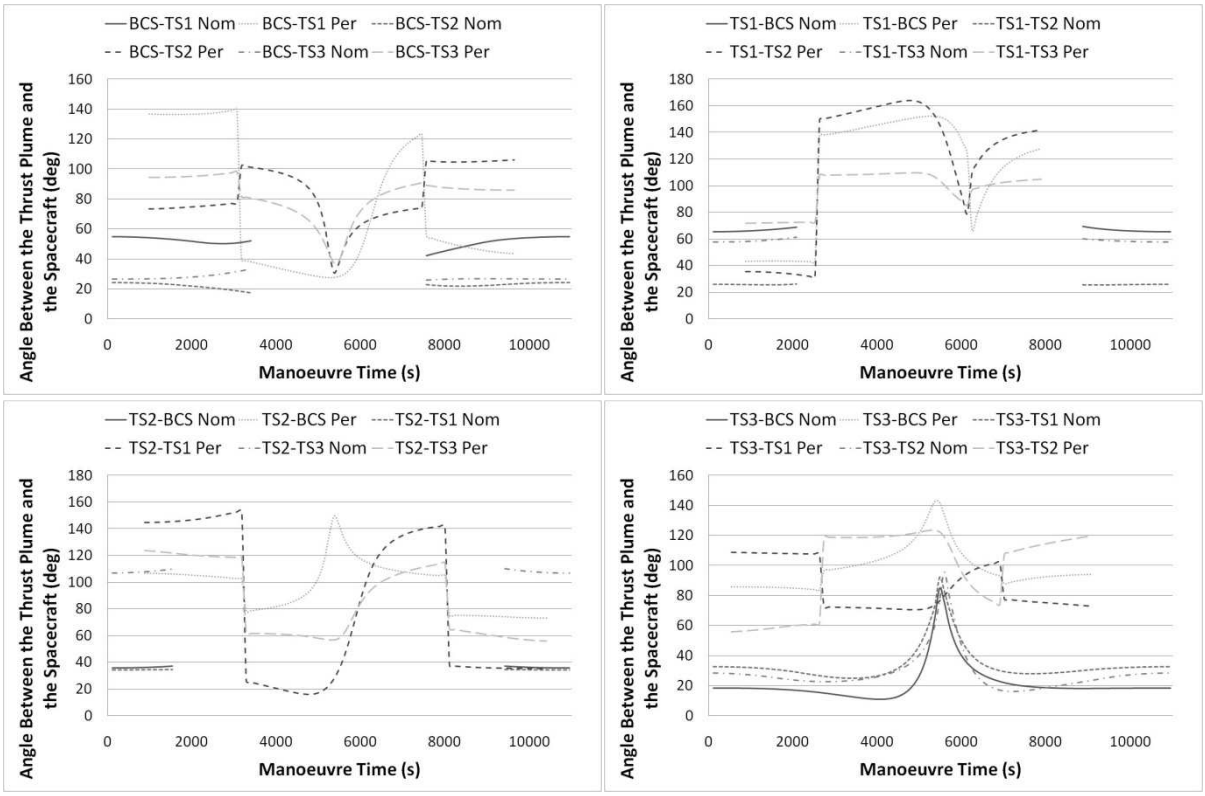

Figure 9-9 Thrust separation angles for the example TDM manoeuvre from BCS thrust plumes (upper left), TS1 thrust plumes (upper right), TS2 thrust plumes (lower left), TS3 thrust plumes (lower right)

9.3.2.3 Manoeuvre Data for the Example Manoeuvre

The input parameter data for the example manoeuvre is given in Table 9-1. The thrust directions, magnitudes and timings for each spacecraft all fall within the ranges specified in section 3.2.1.

Table 9-1 Optimised input data to produce the example manoeuvre

	θ_T (°)	$\lvert\mathbf{T}_{\mathbf{per}}\rvert$ (mN)	$t_{\omega per}$ (s)	t_{per} (s)
BCS	171	2.4	2173	966
TS1	53	4.0	1775	851
TS2	104	1.0	2406	799
TS3	84	2.0	2148	547

9.3.3 Optimisation Issues within the Trajectory Design Module

The example manoeuvre given in the previous section is the result of a number of refinements and modifications that had to be imposed on the optimisation procedure. These changes and their consequences to the optimality of the TDM are detailed in this sub-section.

9.3.3.1 Limitations of the Pattern Search Optimisation Algorithm

As detailed in sub-section 4.4.2.2 the Pattern Search (PS) algorithm generates a mesh in the solution space from a vector of initial input parameters. If any of the mesh points contain a solution lower than the initial solution then the next iteration generates a mesh using the input parameters that generated the lower solution. If however the initial solution remains the lowest solution then the next iteration generates the mesh from the same initial input parameter but using a smaller mesh. For the TDM the easiest initial input parameter is for all the components of the vector Y to be zero, i.e. $Y_0 = \{0,\ldots,0\}$. This will return the cost as infinity ($J_2 = \infty$) since the manoeuvre constraints (Equation (9.6) and/or Equation (9.7)) will have been breached. Due to the complexity of the cost function and the number of elements in the independent variable it is highly unlikely that any of the mesh points in the PS algorithm will find solutions other than infinity when this value of Y_0 is used. Since no smaller solution is found, the PS algorithm iterates with the same Y_0 but using a smaller mesh. This iteration likewise returns a value of infinity for all the mesh points and the process is repeated in an 'infinity-loop'. The algorithm eventually terminates when the mesh size reaches its tolerance value. The net result of this method is no change in the solution given by Y_0 with all the manoeuvre violations still in place. In tests with the example manoeuvre this is the case 100% of the time.

The solution for this problem is simple. The TDM using the PS algorithm needs a viable input parameter (i.e. one where no manoeuvre violations exist) to start correctly. This viable input parameter is found by including an algorithm to find a viable independent variable. The function, dubbed '*TDM_findx0*', does this by randomly generating Y within the range constraints for each component and testing for manoeuvre violations. Figure 9-10 shows a histogram detailing the number of iterations of the TDM objective function using random input variables is required for a viable solution to be found for 1000 repetitions of '*TDM_findx0*'. The mean number of iterations required is ~60 and Figure 9-10 shows that ~80% of all the viable solutions are found within 100 iterations of the objective function. Every iteration of the TDM objective function takes ~44 ms to complete so '*TDM_findx0*' is capable of finding viable solutions to the manoeuvre planning problem within a few seconds. Using the function '*TDM_findx0*' to find the initial input parameter Y_0 for the PS algorithm allows the optimisation algorithm to avoid the 'infinity-loop.'

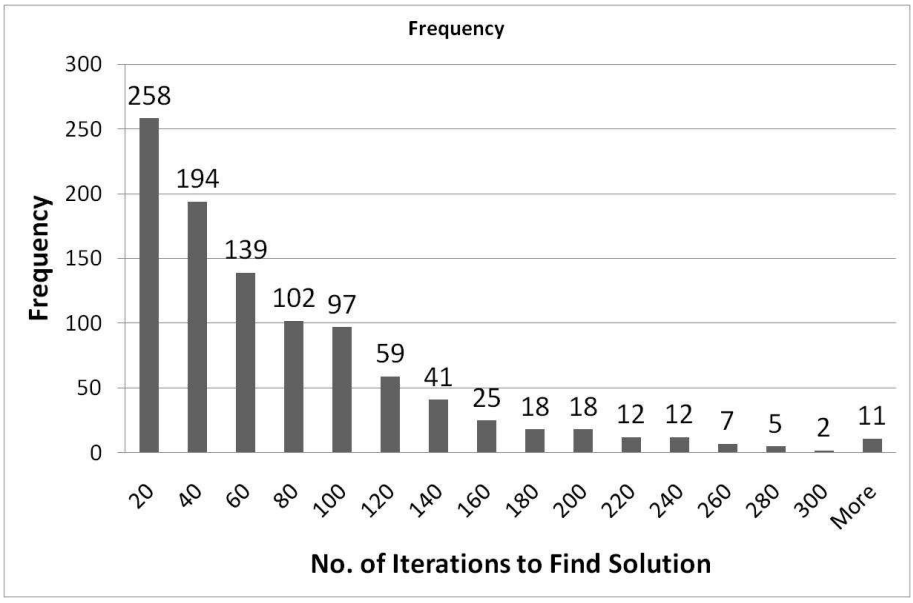

Figure 9-10 Histogram showing the number of iterations required by function '*TDM_findx0*' to find a solution to the manoeuvre planning problem

Using the '*TDM_findx0*' function to find a viable Y_0 is a successful way to start the PS algorithm however it raises further issues with the nature of the TDM solution-space that the PS algorithm cannot resolve. Analysis shows that the quality of the results of the PS algorithm is highly dependent on the values of the initial independent variable Y_0.

Figure 9-11 shows the optimised input parameters for four separate PS optimisations using the example manoeuvre initial conditions and a random initial input parameter for each optimisation. Each diagram is a radar plot of the sixteen input parameters as the vertices. The shape of the plots can be used to compare the components of one input parameter with another. The vertices have been group normalised to enhance the visual aspect of the plot but the numerical values shown are the actual values for the input parameters. The vertices are read clockwise starting at the '12 o'clock' position with $\theta_{T,1}$, then $\theta_{T,2}$ etc. in the order of components of independent variable *Y*, Equation (9.5). The associated cost for each of the four optimisations is given below each radar plot.

From the shape of the plots it is clear that each different optimisation of the TDM with a random Y_0 using the PS algorithm gives very different solutions with very different associated costs. This indicates that the solution of the TDM using the PS algorithm is highly dependent on the initial independent variable used. Using the function '*TDM_findx0*' to randomly find Y_0 suggests that every time the TDM is run with the same initial conditions the PS algorithm will be unlikely to replicate results from a previous iteration.

Figure 9-12 shows radar plots of the optimised input parameters for three PS optimisations using the example manoeuvre initial conditions and a fixed initial input parameter. The upper left plot in Figure 9-12 shows the fixed Y_0 used and its associated cost. Costs for the other three optimisations are given below each of the plots. The three optimisations yield results that are not identical to each other due to the tolerances and randomisations present within the PS optimisation algorithm. However the similarity of the three solutions is evident from the shape of the plots and the numerical values associated with

each individual element of *Y*. This indicates that the PS algorithm can replicate results if a fixed Y_0 is used.

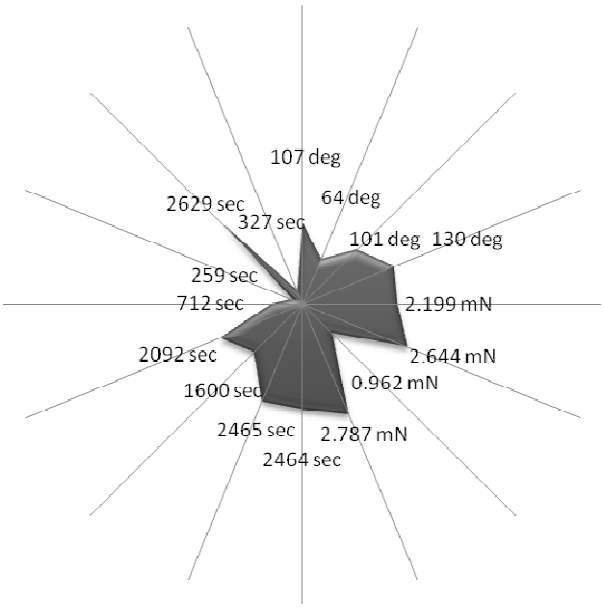

Optimisation 1 Cost = 0.007102

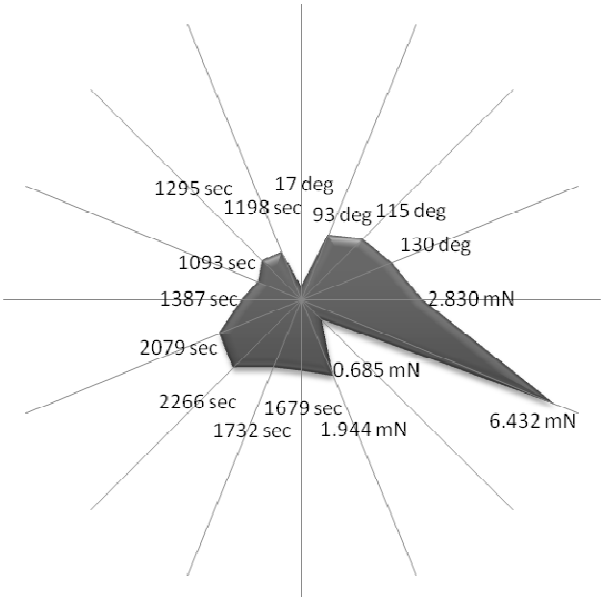

Optimisation 2 Cost = 0.007453

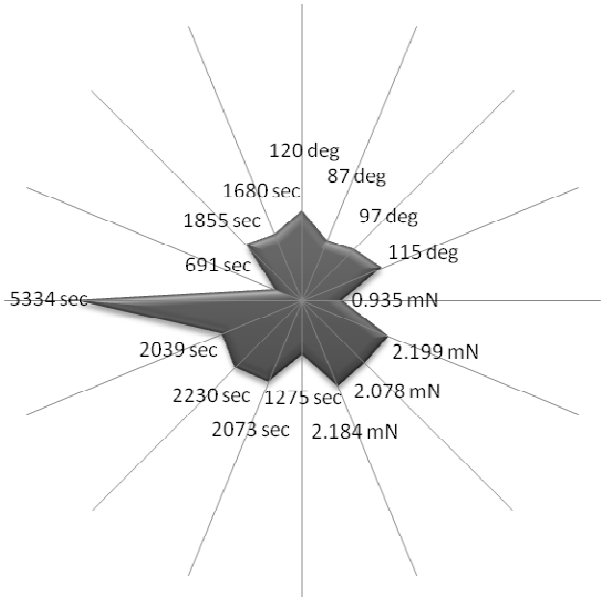

Optimisation 3 Cost = 0.006553

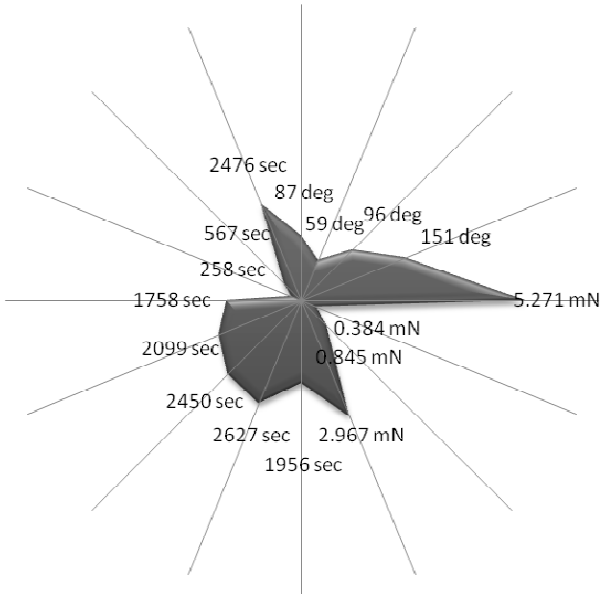

Optimisation 4 Cost = 0.007380

Figure 9-11 Radar plots of four PS optimised manoeuvres using the example manoeuvre initial conditions but random initial input parameter Y_0. The differences in the shapes of the plots indicate that the PS is unable to replicate results from consecutive iterations.

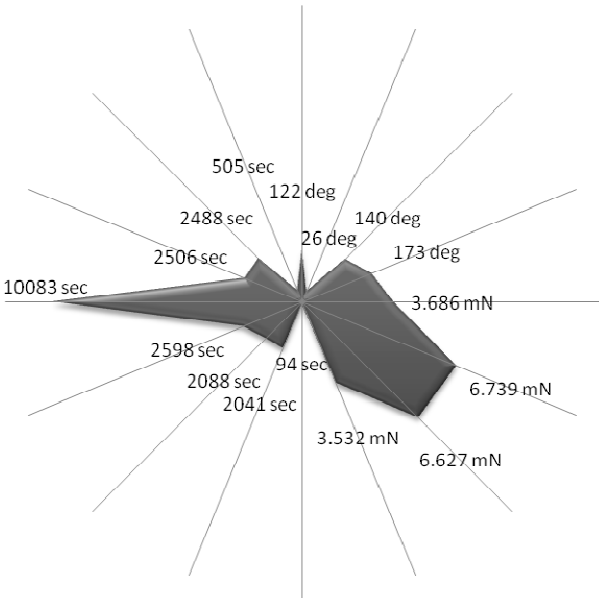

Y_0 Manoeuvre Cost = 0.009279

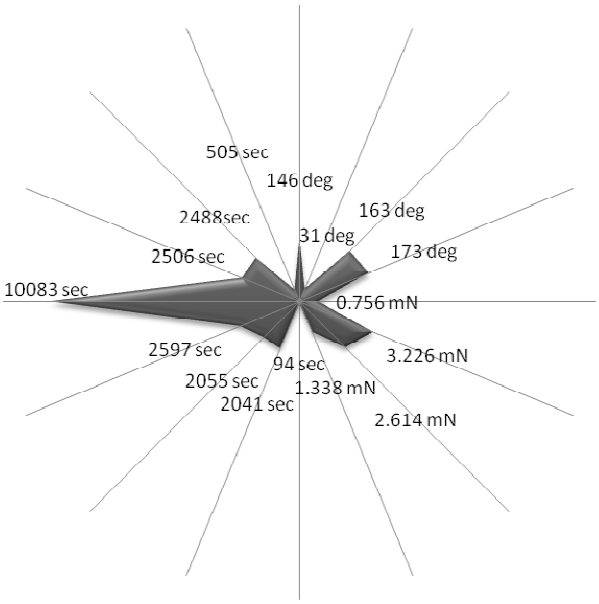

Optimisation 1 Cost = 0.006627

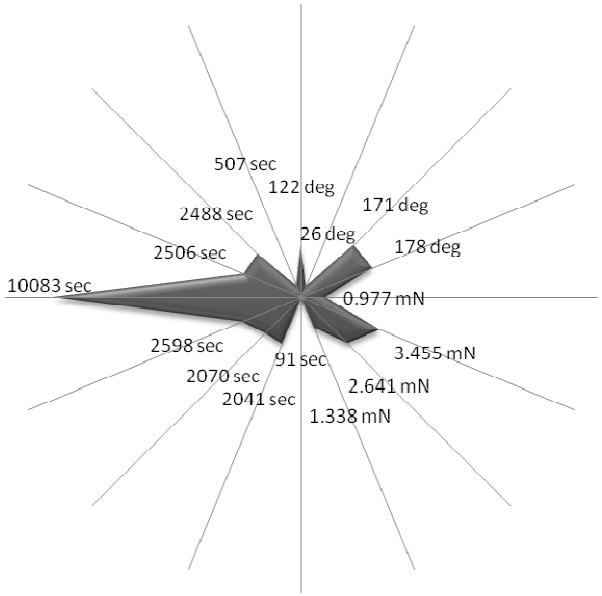

Optimisation 2 Cost = 0.006696

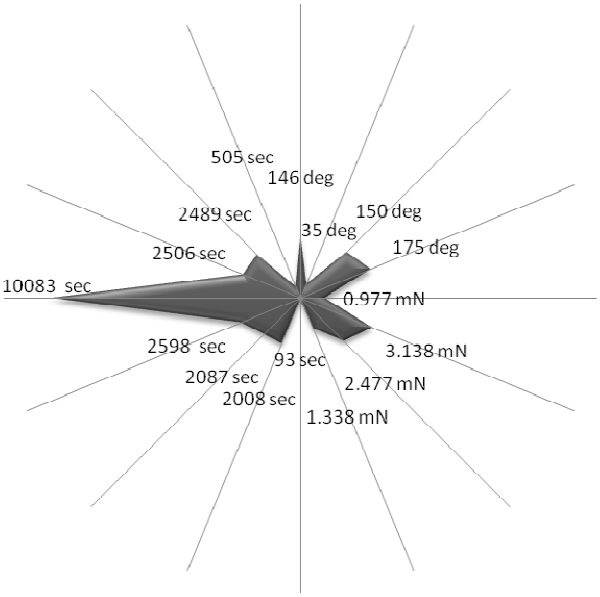

Optimisation 3 Cost = 0.006567

Figure 9-12 Radar plots of three PS optimised manoeuvres using the example manoeuvre initial conditions with fixed initial input parameter Y_0. The similarities in the shape of the plots indicate that the PA algorithm with the TDM cost function is greatly affected by the initial input variables and unable to escape from local minima in the solution-space.

The necessity for the '*TDM_findx0*' function and the data shown in Figure 9-11 and Figure 9-12 lead to some interesting conclusions about the characteristics of the solution-space of the TDM objective function and the PS optimisation algorithms effectiveness to optimise within that solution-space. These are as follows:

1. The solution-space of the TDM objective function resembles a uniform 'sheet' of solutions that violate the manoeuvre avoidance criteria punctuated by multiple local minima representing viable solutions.

2. The '*TDM_findx0*' function randomly finds a solution within the vicinity of one of these local minima and the PS algorithm optimises that solution to approach the actual value of the local minimum.
3. Due to the size of the solution-space the PS algorithm is highly unlikely to be able to escape from the local minimum initially defined by Y_0 and thus is unable to find better local minima solutions or the global minimum within the solution-space

The inability of the PS algorithm to escape the local minimum defined by Y_0 affects its usefulness as the optimisation algorithm to be used for the TDM. There is of course always a random chance that '*TDM_findx0*' places Y_0 within the vicinity of the global minimum but with the likely thousands of local minima available the chances are very low.

9.3.3.2 Using a Genetic Algorithm to Optimise the Trajectory Design Module

Since the PS algorithm is unable to effectively optimise the TDM objective function another optimisation method needs to be considered. The Genetic Algorithm (GA) is better suited to optimising the more complex types of objective function typified by the TDM as it uses multiple, distributed, simultaneous solutions to explore the solution-space. Specifically, where the PS algorithm gets trapped in a local minimum defined by the initial input parameter Y_0, the GA starts with multiple, distributed Y_0s, dubbed the 'initial population'. The GA is able to search much more extensively around the solution-space and has a higher probability of finding the global minimum of the objective function. There is no way of knowing whether the GA has found the global minimum but as long as it consistently finds better solutions than the PS algorithm it is worth considering as an alternative.

Sub-section 4.4.4.1 gives detailed explanation as to how the GA works. For this analysis, two parameters specifically were chosen to be tuned:

- Initial Population Size – the number of individuals used to search the solution space with. The larger the initial population size the more processor intensive and time consuming the optimisation but the higher likelihood of finding the globally optimal solution. The initial population itself is calculated using successive iterations of the function '*TDM_findx0*', thus all initial individuals are viable solutions of the objective function and do not violate any manoeuvre constraints.
- Crossover Fraction – the fraction of the population (not including elite children) that evolve by crossover (as opposed to mutation). The optimal value of this depends on the properties of the solution-space.

Population size analysis for the GA algorithm using the TDM-objective function can be seen in Figure 9-13. The data is produced by running the GA with variable population size using the initial conditions of the example manoeuvre and a crossover fraction of 0.2. The filled diamonds represent the data points while the solid line represents the mean over 10 iterations. Figure 9-13 shows that, as expected, the optimised cost decreases with increasing population size however for population sizes greater than ~100 there is little gain compared with the extra time required to achieve optimised solutions.

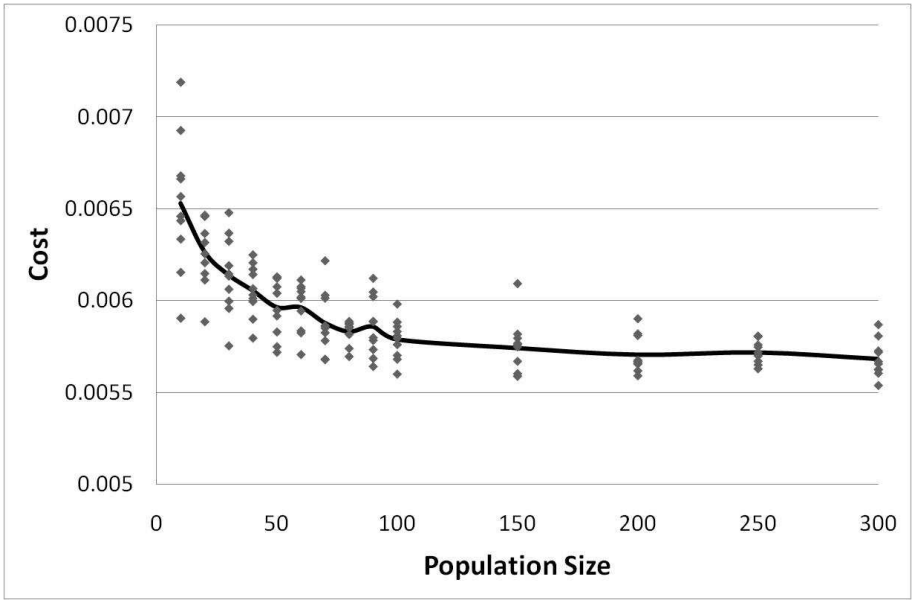

Figure 9-13 Population size analysis for the GA with the TDM objective function. The chart shows that the cost decreases with increasing population size but that decrease slows as population size increases.

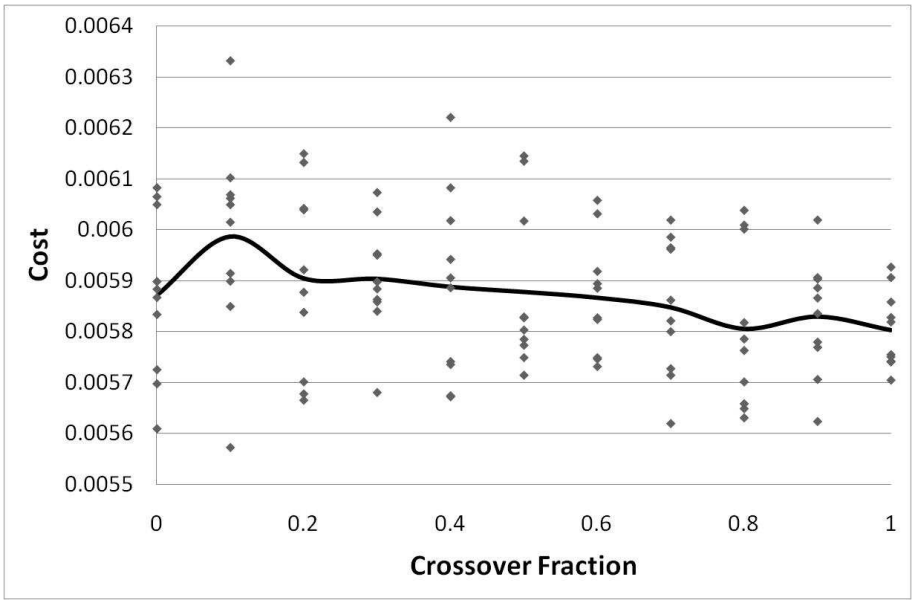

Figure 9-14 Crossover fraction analysis for the GA with the TDM objective function. The chart shows a slight decrease of cost with increasing crossover fraction with a minimum observed at 0.8.

Crossover fraction analysis for the GA using the TDM-objective function can be seen in Figure 9-14. The data is produced by running the GA with variable crossover fraction using the initial conditions of the example manoeuvre and an initial population size of 100. The filled diamonds represent the data points while the solid line represents the mean over 10 iterations. The data in Figure 9-14 is highly spread out over all crossover fraction values however the mean data does show a slight decrease in cost as the crossover fraction increases.

Due to the spread of the values it is reasonable to presume that for the TDM objective function the value of the crossover fraction makes no real difference to the minimum cost found by the GA.

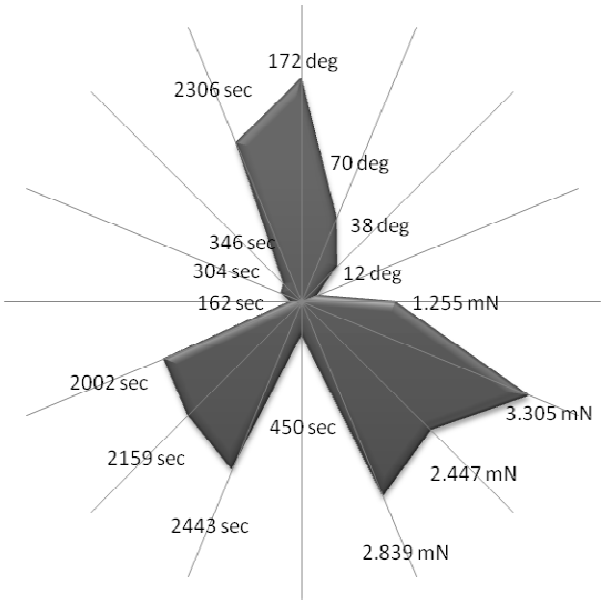

Optimisation 1 Cost = 0.006954

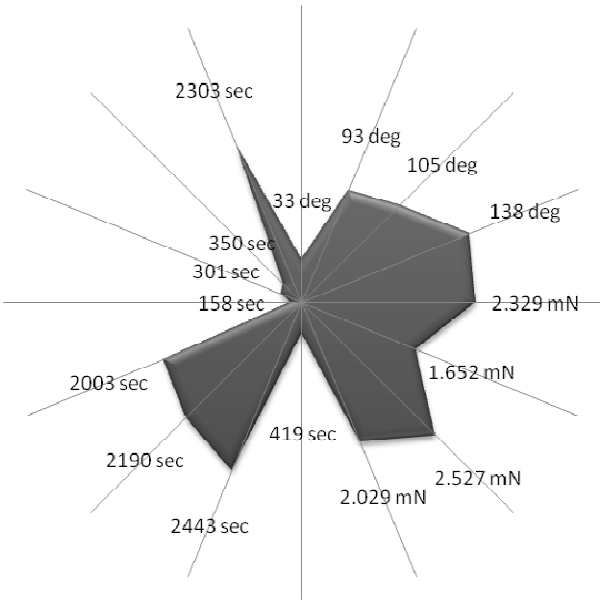

Optimisation 2 Cost = 0.006534

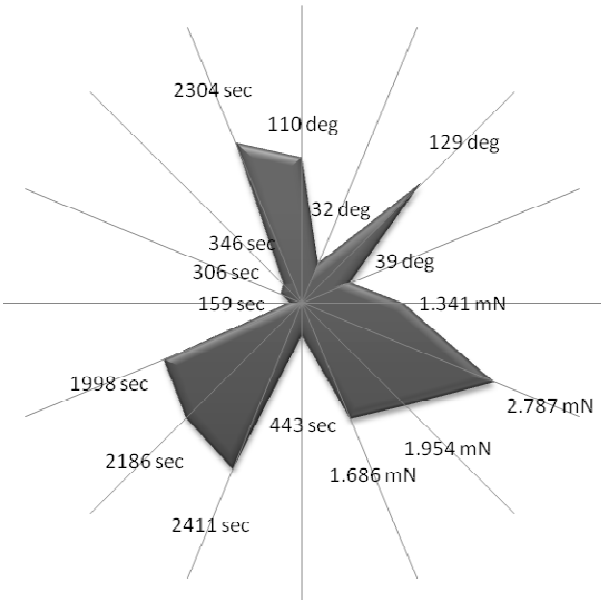

Optimisation 3 Cost = 0.006570

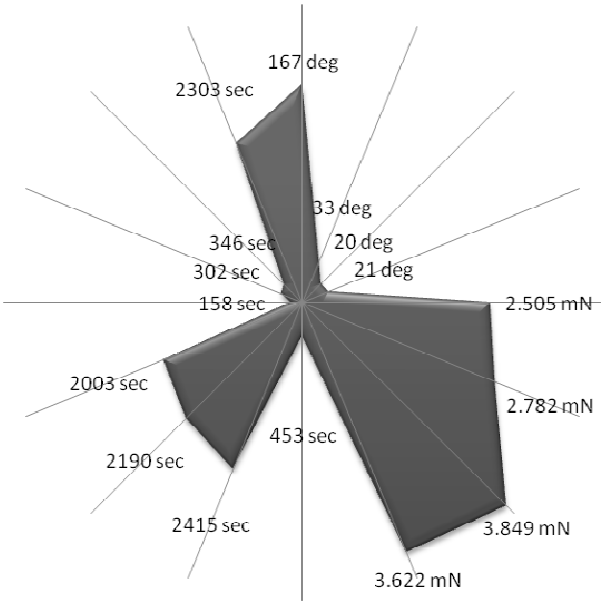

Optimisation 4 Cost = 0.007113

Figure 9-15 Radar plots of four GA optimised manoeuvres using the example manoeuvre initial conditions with random initial population. The differences in the shapes indicate the stochastic nature of the GA but the similarities in the shapes indicate that the GA is able to consistently find good values for some of the input elements.

Radar plots of four GA optimised manoeuvres using the example manoeuvre initial conditions with a random initial population are shown in Figure 9-15. Populations of 100 individuals were used with a crossover fraction of 0.8. The four plots show the optimised independent variable and associated costs. In all four plots these values are different. This

emphasises the random nature of both the function '*TDM_findx0*' and the GA. An interesting point to note however is that the GA solution input parameters have similarities especially for the values of $t_{\omega per,i}$ and $t_{per,i}$. So although the initial populations are different for each plot the GA is finding that certain values of $t_{\omega per,i}$ and $t_{per,i}$ yield solutions with the lowest cost. This result was not observed whilst using the PS algorithm with random Y_0s.

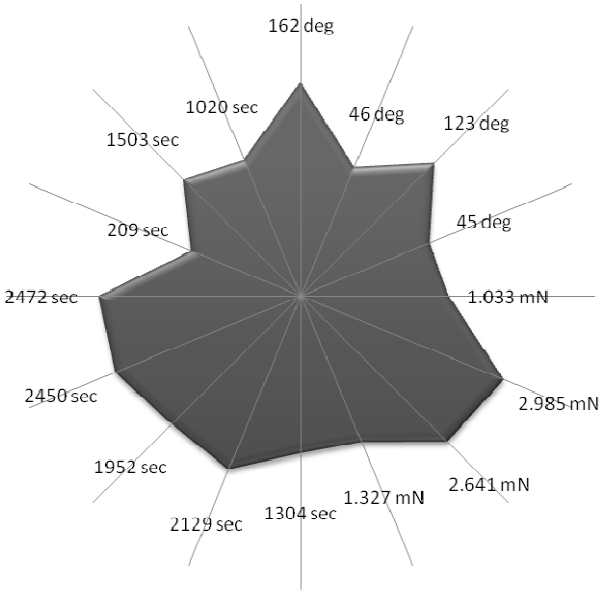

Optimisation 1 Cost = 0.006700

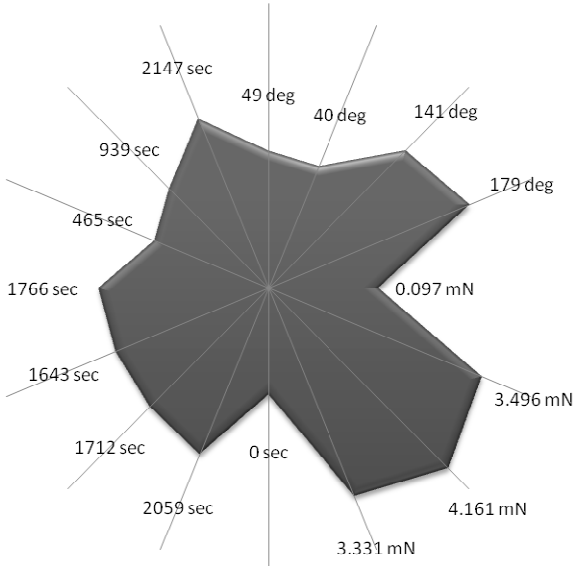

Optimisation 2 Cost = 0.007155

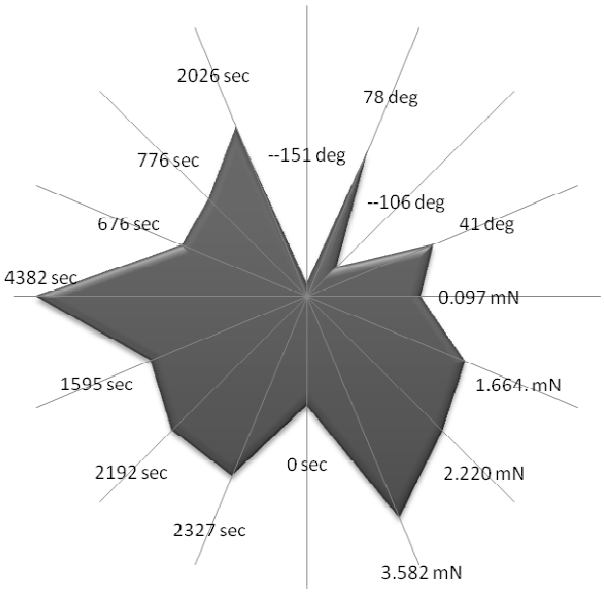

Optimisation 3 Cost = 0.006495

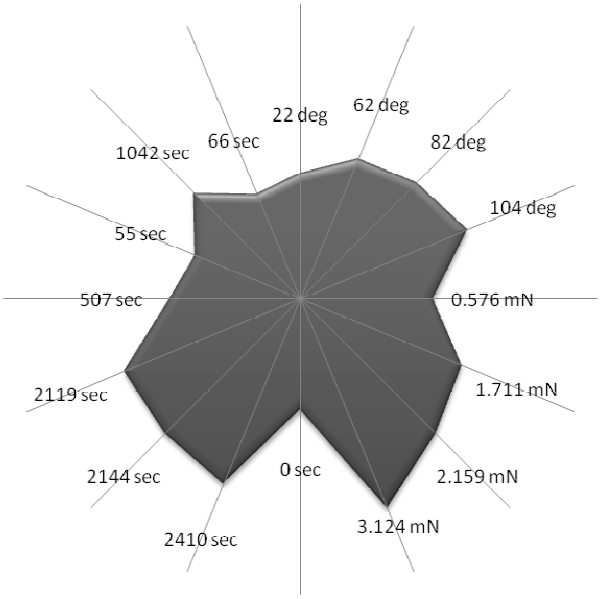

Optimisation 4 Cost = 0.006608

Figure 9-16 Radar plots of four GA optimised manoeuvres using the example manoeuvre initial conditions with a fixed initial population. The differences in the shapes of the plots indicate that the GA is able to escape local minima even when the initial population is the same between iterations.

Figure 9-16 shows radar plots of four GA optimised manoeuvres using the example manoeuvre initial conditions but the same initial population of 100 individuals and a crossover fraction of 0.8. Again the plots all look different emphasising the GAs ability to search though many more solutions and not be affected by the individuals in the initial population. This result is the reverse of that for the PS algorithm with fixed Y_0 where all the results are very similar to each other due to the inability of the PS to escape local minima (Figure 9-12).

9.3.3.3 Comparing the Genetic Algorithm with the Pattern Search Algorithm

The associated costs given in Figure 9-11, Figure 9-12, Figure 9-15 and Figure 9-16 give a small measure of comparison regarding the ability of each optimisation algorithm to find low cost solutions to the example problem. These figures show that both algorithms are comparable in this task however the size of the data set to too small to make any firm judgements. A larger data set can be found in Table 9-2.

The data in Table 9-2 is generated using 100 iterations of TDM using the example manoeuvre for the initial conditions. The column headed 'Random' shows the mean cost and mean calculation time when no optimisation algorithm is used and the manoeuvre input parameters are found simply from the function '*TDM_findx0*'. The columns headed 'PS' and 'GA' show the mean cost and mean calculation time using the Pattern Search algorithm and the Genetic Algorithm respectively whilst the column headed 'GA+Grad' shows the data where the GA is used to find a solution and then a gradient-based optimisation algorithm is used to search the local solution-space to find a more optimal solution[11]. The data show that on average compared with the selection of a random solution all the optimisation algorithm choices are capable of finding more optimal solutions and therefore it is worth spending the extra calculation time required to find these solutions. The optimisation algorithm showing the lowest mean cost is the 'GA+Grad' algorithm. Although this algorithm takes the longest time to find a solution a calculation time less than 10 min using the reference hardware is acceptable for the TDM. Obviously, calculation time is dependent on the system the TDM is run on but if necessary the GA or PS algorithm could be used with only a small loss in optimisation performance. For all analysis on the TDM that follows for the rest of the chapter the 'GA+Grad' algorithm is used.

Table 9-2 Comparison of different optimisation methods using the example manoeuvre with the TDM objective function

Over 100 iterations:	Random	PS	GA	GA + Grad
Mean Cost	0.010377	0.007355	0.006813	0.006788
Mean Calc Time (sec)	1.289	75.715	353.141	433.55

[11] This gradient based algorithm is the '*fmincon*' algorithm from MATLAB's Optimisation Toolbox.

9.3.4 Analysis of DARWIN-like Manoeuvres

The example manoeuvre used in the previous sub-sections is not a manoeuvre that a DARWIN formation would make. In this sub-section the performance of the TDM when optimising DARWIN-like manoeuvres will be analysed.

9.3.4.1 Trajectory Design Module Usage

The TDM is designed to generate collision free and plume impingement free trajectories between the spacecraft initial and final positions optimised by the PAM. Trajectory optimisation by the TDM however may not be necessary if the straight line trajectories used by the PAM are manoeuvre violation free. With avoidance criteria fixed at <10m for collisions and <5° for plume impingement 200 PAM optimised manoeuvres were examined for manoeuvre violations, 100 of these manoeuvres were optimised for fuel minimising ($\mu_{PAM}=0$) and 100 for fuel balancing ($\mu_{PAM}=0.6$). The data shows that the TDM is required to optimise the trajectories for ~10% of all PAM optimised manoeuvres when $\mu_{PAM}=0$ and 44% of all PAM optimised manoeuvres when $\mu_{PAM}=0.6$. The higher occurrence of manoeuvre violations for fuel balancing optimisation is related to the increase in trajectory distances for fuel balancing manoeuvres thus a higher risk of collision and/or plume impingement. This data is beneficial for a number of reasons:

- It justifies the usage of the TDM for manoeuvre planning of DARWIN-like manoeuvres
- The majority of manoeuvre plans will be performed quicker since the TDM is not required
- The majority of manoeuvre plans will remain fuel/balancing optimal since the TDM is not required

9.3.4.2 Example Trajectory Optimisation on a DARWIN-like Manoeuvre

An example of a TDM optimisation of a DARWIN-like manoeuvre can be seen in Figure 9-17. The manoeuvre involves a formation retarget of ~45° but remaining in the linTTN configuration. The final spacecraft positions were optimised by the PAM for fuel balancing ($\mu_{PAM}=0.6$) with initial fuel distribution $mf_i(t_0)=\{4.8,4.7,4.9,4.85\}kg$. For the TDM the same initial parameters are used as in sub-section 9.3.1, the avoidance criteria remain at 10m and 5° and the TDM is set to optimise for fuel minimising ($\mu_{TDM}=0$). The collision violations occur between TS1-TS2 and TS2-TS3.

The optimised independent variable, *Y*, to generate the collision-free formation manoeuvre in Figure 9-17 is shown in Table 9-3. The avoidance manoeuvres are performed by TS1 and TS2. Though the optimisation algorithm has allocated values for the input parameters for the BCS and TS3 their affect on the manoeuvre itself is negligible since the allocated thrust, $|\mathbf{T}_{\mathbf{per}}|$, is below the set thrust resolution (1mN) and so is registered as $|\mathbf{T}_{\mathbf{per}}|=0$ within the algorithm.

Table 9-3 Optimised input data to obtain the TDM optimisation of the PAM optimised manoeuvre

	θ_T (°)	$\lvert\mathbf{T}_{per}\rvert$ (mN)	$t_{\omega per}$ (s)	t_{per} (s)
BCS	-59	0.9	559	948
TS1	97	5.5	1257	660
TS2	171	4.9	1319	338
TS3	177	0.6	3	714

9.3.4.3 Fuel Minimising and Fuel Balancing Performance

As well as generating spacecraft trajectories that are collision and plume impingement free the TDM cost function, Equations (9.4) and (9.8), also includes a term (μ_{TDM}) that can be set to optimise the trajectories for fuel minimising or fuel balancing. This means that although the final spacecraft positions are optimised by the PAM for fuel minimisation or fuel balancing, the avoidance trajectories can also be optimised for the same properties.

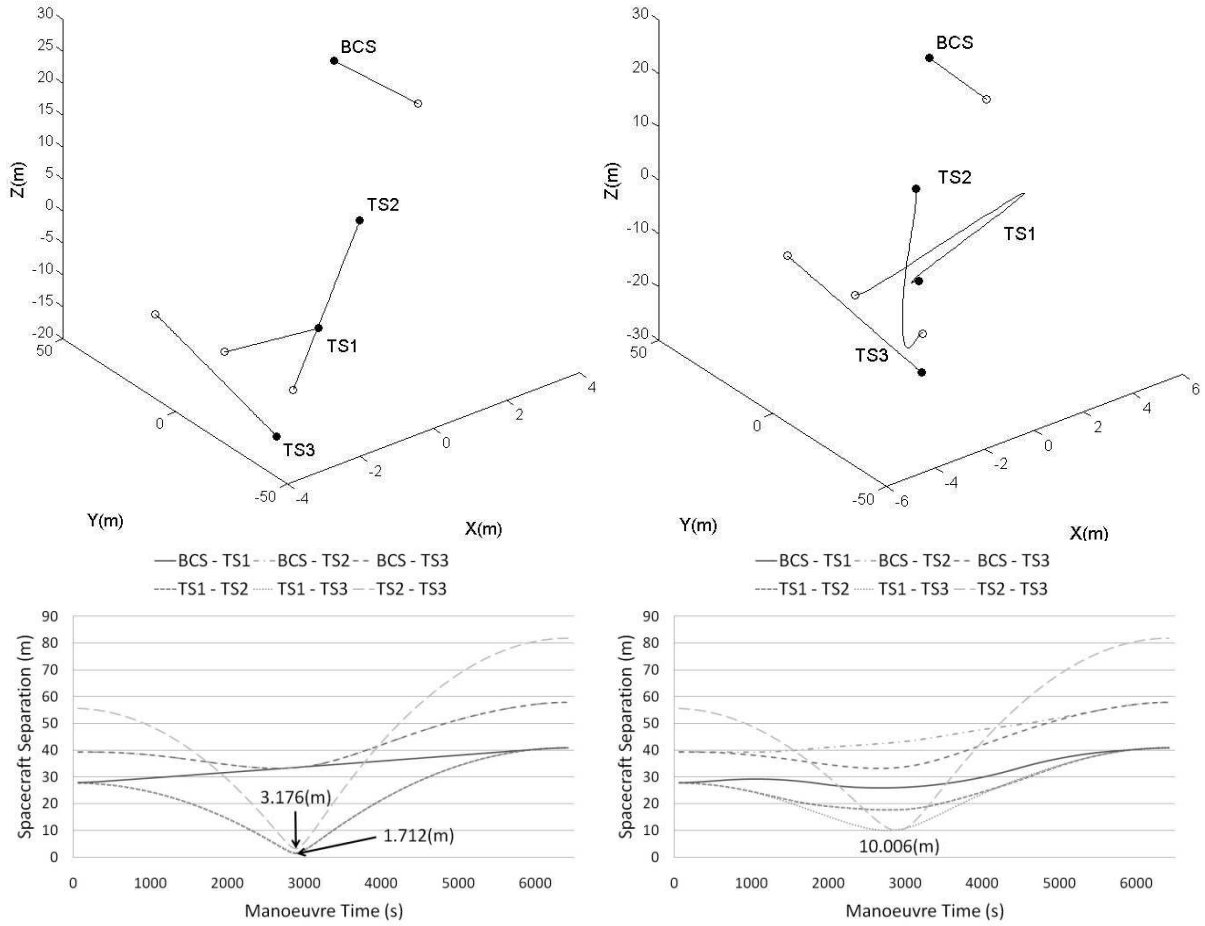

Figure 9-17 Example of TDM optimisation of a PAM optimised manoeuvre. No avoidance (upper left), avoidance activated (upper right), no avoidance spacecraft separations (lower left), avoidance activated spacecraft separations (lower right)

Using the 10 PAM optimised manoeuvres with $\mu_{PAM}=0$ that incurred manoeuvre violations (from sub-section 9.3.4.1) Figure 9-18 compares the fuel usage and fuel balancing for TDM optimised trajectories of those PAM manoeuvres. The charts show the data for the PAM optimisation (i.e. the trajectories with manoeuvre violations), data for TDM optimisation for fuel minimising ($\mu_{TDM}=0$) and data for TDM optimisation for fuel balancing ($\mu_{TDM}=0.6$). A value of μ_{TDM}=0.6 for fuel balancing was chosen since in Chapter 8, 0.6 is identified as giving maximum fuel balancing returns. Since the fuel optimisation part of the TDM cost function is practically identical to the PAM cost function it is likely that $\mu_{TDM}=0.6$ gives the same results as $\mu_{PAM}=0.6$ does for the PAM. A more in-depth analysis of μ_{TDM} follows in sub-section 9.3.4.4.

The upper chart in Figure 9-18 shows the fuel usage of each manoeuvre. As expected the TDM optimisation causes the spacecraft to consume more fuel since they are performing avoidance manoeuvres that remove them from the nominal (PAM optimised) trajectory. When the TDM is set for fuel minimising the average increase in fuel consumption is ~40% whilst when set for fuel balancing the average increase is ~294%. The lower chart in Figure 9-18 shows the sum of the fuel differences for each manoeuvre. When the TDM is set for fuel minimising the sum of the fuel differences increases by, on average, ~0.25% for all the manoeuvres observed but when the TDM is set for fuel balancing the sum of the fuel differences decreases by ~1.25% on average. So although the TDM is able to perform fuel balancing for $\mu_{PAM}=0$ manoeuvres it does this at a very high cost of fuel.

Using 10 of the 44 PAM optimised manoeuvres with $\mu_{PAM}=0.6$ that incurred manoeuvre violations (from sub-section 9.3.4.1) Figure 9-19 compares the fuel usage and fuel balancing for TDM optimised trajectories of those PAM manoeuvres. The charts show the data for the PAM optimisation, data for TDM optimisation for fuel minimising ($\mu_{TDM}=0$) and data for TDM optimisation for fuel balancing ($\mu_{TDM}=0.6$). As in Figure 9-18 the fuel usage data shows an increase in fuel consumption for TDM optimised manoeuvres. When the TDM is set for fuel minimising the fuel consumption increases by, on average, ~50% over all 44 manoeuvres observed. When set for fuel balancing the fuel consumption is ~157% greater. The fuel balancing data is shown in the lower chart in Figure 9-19. The general decreasing trend reflects the fact that the PAM data was generated using a tour of manoeuvres, however it is the comparison of the performance of the TDM optimised trajectories with the PAM optimised ones that is of interest here. When the TDM is set for fuel minimising the data show that the sum of the fuel differences is sometimes greater and sometimes less than the PAM optimised value. Averaged out over the 44 manoeuvres examined gives an average decrease in the sum of the fuel differences of ~0.3%. When set for fuel balancing however the TDM always finds trajectories that decrease the sum of the fuel differences for a ~1.6% increase in fuel balancing. Again the TDM is capable of generating trajectories that fuel balance better than those generated by the TDM but the fuel cost remains high.

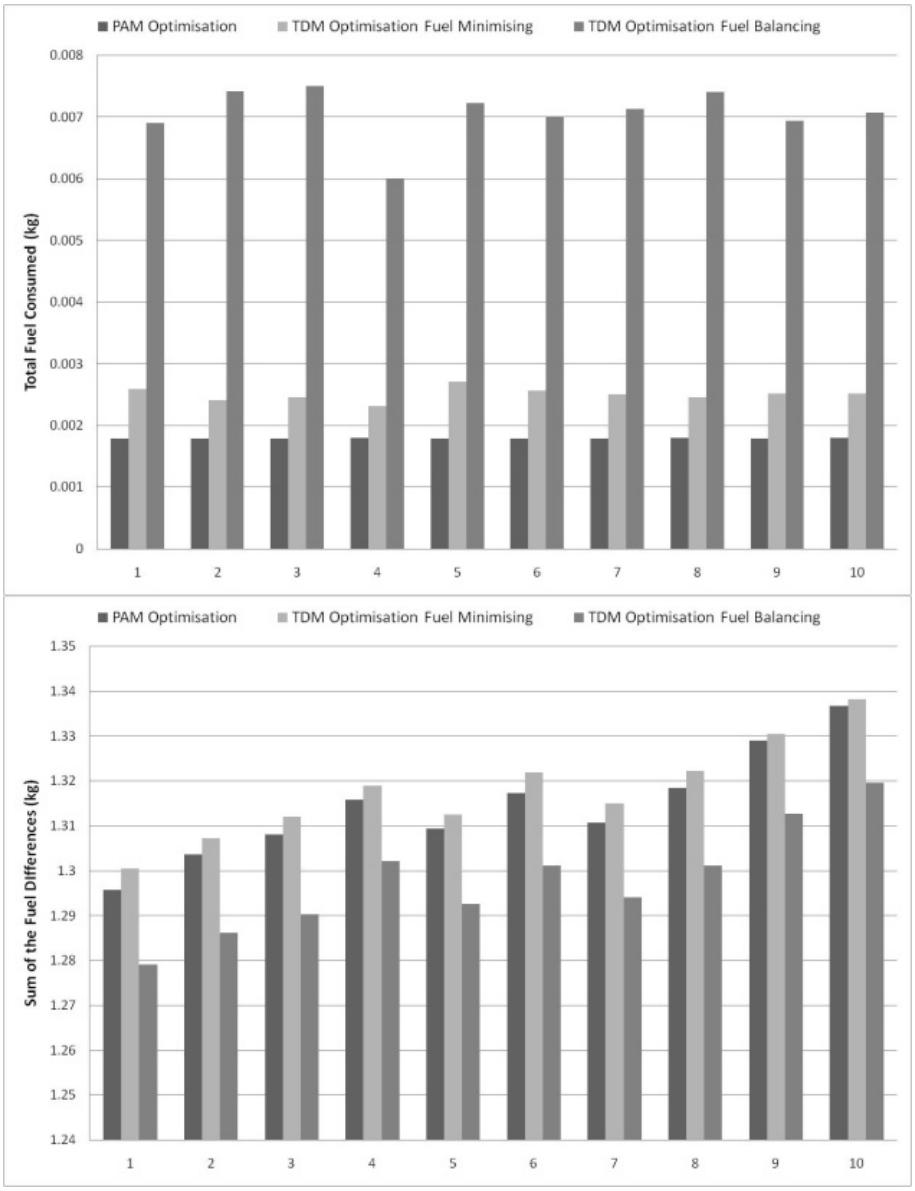

Figure 9-18 Fuel usage data (upper) and fuel balancing data (lower) for 10 TDM optimised trajectories for PAM optimised manoeuvres with $\mu_{PAM}=0$. The charts show how much extra fuel is required by the TDM for avoidance manoeuvres and the effect of the fuel balancing of the formation. For a fuel balancing TDM the fuel balancing achieved is minimal with respect to the fuel consumed.

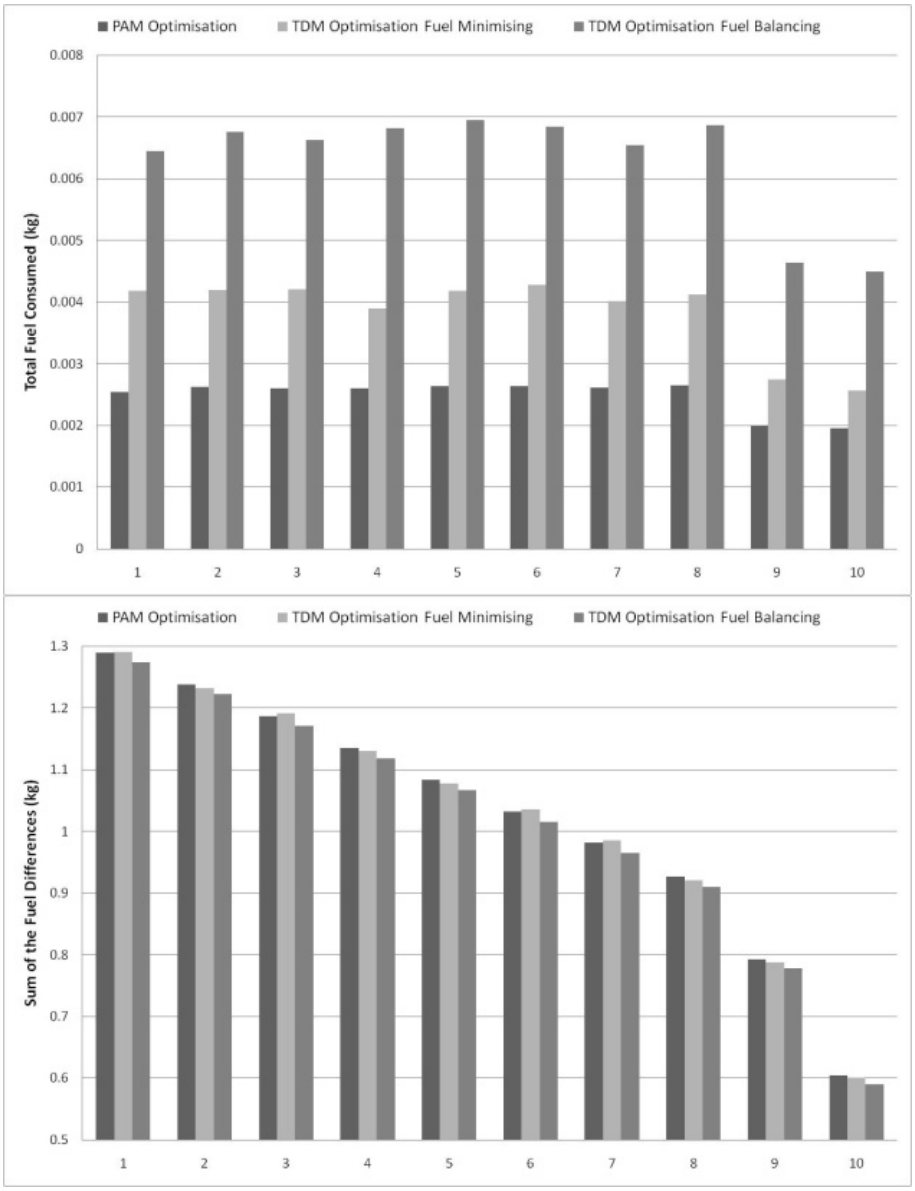

Figure 9-19 Fuel usage data (upper) and fuel balancing data (lower) for 10 TDM optimised trajectories with for PAM optimised manoeuvres with $\mu_{PAM}=0.6$. The charts show how much extra fuel is required by the TDM for avoidance manoeuvres and the effect of the fuel balancing of the formation. For a fuel balancing TDM the fuel balancing achieved is minimal with respect to the fuel consumed.

9.3.4.4 Analysing μ_{TDM}

The performance analysis from the previous sub-section used a value of $\mu_{TDM}=0.6$ for the fuel balancing element. As for the PAM however, this fuel balancing term can take any value and the fuel balancing performance of the algorithm is affected. Figure 9-20 shows how the total fuel consumption and the sum of the fuel differences varies for a TDM optimised manoeuvre with different values of μ_{TDM}. The top two charts are for one $\mu_{PAM}=0.0$ manoeuvre whilst the lower charts are for one $\mu_{PAM}=0.6$ manoeuvre. For both manoeuvres the initial conditions remain the same as in previous sections and the initial fuel distribution is un-balanced with $mf_i(t_0)=\{4.8, 4.7, 4.9, 4.85\}kg$.

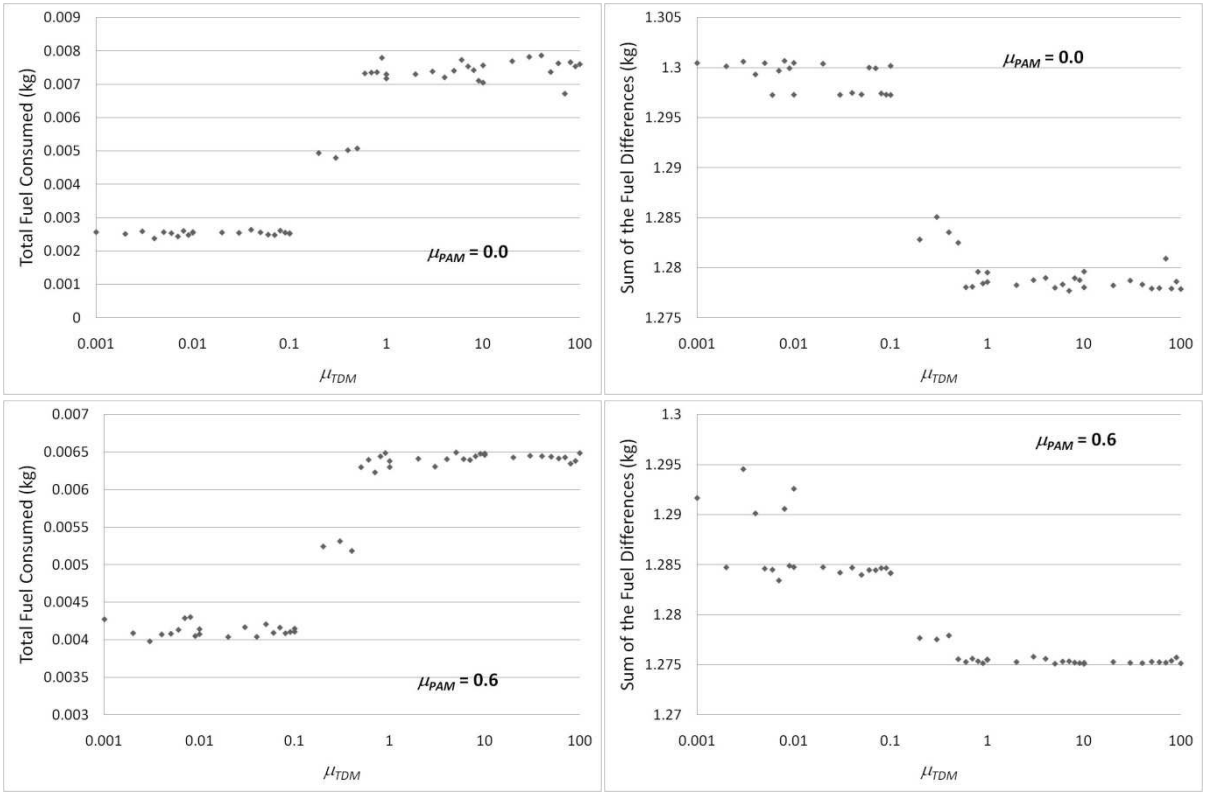

Figure 9-20 μ_{TDM} analysis. Total fuel consumed and sum of the fuel differences for a μ_{PAM}=*0.0* manoeuvre (top left and top right) and total fuel consumed and sum of the fuel differences for a μ_{PAM}=*0.6* manoeuvre (lower left and lower right). The charts show great similarities to those on Chapter 8 (Figure 8-13) suggesting the TDM cost function returns similar μ_{TDM} performance as μ_{PAM} does for the PAM cost function.

Although the charts in Figure 9-20 only represent two separate manoeuvres they show a direct analogy with the charts generated for the analysis of μ_{PAM} in Chapter 8. The effective range of μ_{TDM} appears to be $0.1 \leq \mu_{TDM} \leq 1$ where fuel consumption changes and sum of the fuel differences changes are at their greatest and very little change is observed outside this range. These similarities with the μ_{PAM} analysis indicate that the TDM fuel cost equation, Equation (9.8), creates a similar cost topography within the local minima in the TDM solution-space to the cost topography for the PAM solution-space. This is to be expected since the equations are in essence identical. With this assumption in place it is inferred that for maximum fuel balancing within the TDM, $0.5 \leq \mu_{TDM} \leq 0.7$.

9.3.5 Trajectory Design Module Analysis Conclusions

The analysis of the operation and performance of the TDM has revealed a number of important issues which are summarised below:

- The TDM is not required to find the trajectories for all PAM optimised manoeuvres, however it is required for ~10% of all fuel minimising PAM manoeuvres and ~44% of all fuel balancing PAM manoeuvres.

- Due to the complexity of the cost function, the TDM needs to utilise a population-based optimisation algorithm (like a genetic algorithm) that uses multiple distributed solutions whilst searching the solution-space. This inevitably increases the computing resources required to run the algorithm. The best performance is found by using the Genetic Algorithm followed by a gradient-based optimisation algorithm.
- Maximum fuel minimisation is achieved when $\mu_{TDM}=0.0$ (as expected) and maximum fuel balancing when $\mu_{TDM}=0.6$.
- With $\mu_{TDM}=0.0$ for fuel minimising the TDM typically finds trajectories that on average use ~40% more fuel for fuel minimising PAM manoeuvres and ~50% more fuel for fuel balancing PAM manoeuvres. The increase in the sum of the fuel differences for both types of PAM optimised manoeuvre is in the region of ~0.3% on average.
- With $\mu_{TDM}=0.6$ for fuel balancing the TDM typically finds trajectories that on average use ~150% more fuel for fuel minimising PAM manoeuvres and ~300% more fuel for fuel balancing PAM manoeuvres. The decrease in the sum of the fuel differences for both types of PAM optimised manoeuvre is in the region of ~1.5% on average.
- For fuel balancing PAM optimised manoeuvres the total fuel consumption is on average >70% more than the equivalent fuel minimised manoeuvre (see Chapter 8). Using $\mu_{TDM}=0.6$ to perform further fuel balancing on the formation is impractical due to the considerable amount of fuel required to do so and the relatively small increase in fuel balancing that will be achieved. Therefore, to minimise the extra amount of fuel the avoidance manoeuvres require (and accept the small fuel balancing penalty) it is suggested that for all TDM optimised manoeuvres μ_{TDM} should be set to zero. This modifies (and simplifies) Equation (9.8) to:

$$J_1 = \sum_i^n \left(f_i(t_0) - f_i(t_f) \right) \tag{9.13}$$

9.4 Future Work

The Trajectory Design Module (TDM) introduced in this chapter presents a good way to include collision and thrusters plume impingement avoidance into the manoeuvre planning for the Separate Modular Manoeuvre Planning Architecture (SepM-MPA). The TDM is not without its limitations however. This sub-section details these limitations and suggests a few ways to overcome them that can be included as future work.

9.4.1 Solution-Space Limitations

As detailed in sub-sections 9.2.2.1 and 9.3.3.1 the solution-space of the TDM cost function, Equations (9.4) to (9.8), resembles a uniform ‘sheet’ of unviable solutions (with a cost $J_2 = \infty$) punctuated with viable solutions that conform to the proximity constraints of the manoeuvre. The nature of this solution-space limits the effectiveness of the Pattern Search (PS) optimisation algorithm as it is unable to escape local minima. The solution for this is to adopt a Genetic Algorithm (GA), with an initial population of 100 viable solutions, as it is

able to search the solution-space more effectively. The initial population of solutions is found by a function '*TDM_findx0*' that used random values for the independent variables elements to build up the required population size.

During the course of the analysis of the TDM algorithm it was found that in certain circumstances the function '*TDM_findx0*' was unable to find any viable solutions to the optimisation problem. Without an initial population the GA was unable to execute. This was observed most keenly for manoeuvres where a formation reconfiguration is required (i.e. from linTTN to triTTN and vice-versa) with only a very small angular retarget. In this manoeuvre the Beam Combiner Spacecraft (BCS) and Telescope Spacecraft 1 (TS1) effectively swap positions within the formation with a collision proximity violation triggered. In this situation there are relatively few independent variables, *Y*, that provide the required avoidance manoeuvre since the proximity violation is so severe. The random nature of the '*TDM_findx0*' function has very little chance of finding them in a time acceptable for the optimisation.

One way to avoid this is to change the topology of the solution space so that the infinities are removed. This can be achieved by removing the collision avoidance and thrusters plume impingement terms from the cost function in instead including them as constraints to the optimisation algorithm. This can be realised by changing the TDM cost function equation (9.4) to

$$J_2 = \min_{Y}\{J_1\} \tag{9.14}$$

effectively making it the same cost function as used by the PAM. The avoidance parameters can then be realised as optimisation constraints

$$\left|\mathbf{r}_\mathbf{i}(t) - \mathbf{r}_\mathbf{j}(t)\right| \le 10m \tag{9.15}$$

for the collision avoidance cost and

$$\sin^{-1}\left(\frac{\left|\mathbf{r}_\mathbf{i}(t)\times\mathbf{r}_\mathbf{j}^\mathbf{P}(t)\right|}{\left|\mathbf{r}_i(t)\right|\left|\mathbf{r}_\mathbf{j}^\mathbf{P}(t)\right|}\right) \le 5^\circ \tag{9.16}$$

for the thruster plume avoidance cost. Not only would this remove the infinities from the solution-space, thus smoothing it out, it would also allow for simpler optimisation algorithm to be employed, thus speeding up the optimisation time.

9.4.2 Alternative Trajectory Generation Methods

There are a large number of other ways in which the avoidance trajectories could be planned each with their advantages and disadvantages over the TDM method described in sub-section 9.2. One of these is simply to change the timing of the nominal thrust pulses for each spacecraft so that the avoidance criteria are met. This would be a very simple task to

achieve and removes the complexity of adding additional thrust components to each spacecraft's manoeuvre. The disadvantage of this approach however would be the increase of the formation manoeuvre time. As time optimality is a key requirement for the manoeuvre planning this trajectory generation method is less suitable.

Assuming the manoeuvre time calculated by the PAM is to remain constant other ways to generate the avoidance manoeuvres involve variations on the thrust profile chosen. One way could be to use two complimentary bang-coast-bang thrust profiles as in Figure 9-21 (left). Another could be to use multiple complimentary bang-bang avoidance thrust profiles as in Figure 9-21 (right). The alternative thrust profiles shown in Figure 9-21 require additional optimisable parameters to utilise the increase in manoeuvre flexibility that they provide. The TDM method uses four parameters and with four spacecraft in the DARWIN fleet this gives a total of sixteen optimisable parameters for the optimisation routine to use. This already poses a complex combinatorial optimisation problem. Using either of the alternative thrust profiles increases the number of optimisable parameters (per spacecraft) and therefore increases the complexity of the optimisation problem. To solve this more complex problem would require increased processing capacity and more calculation time and these requirements could be disadvantageous for an autonomous onboard manoeuvre planning algorithm.

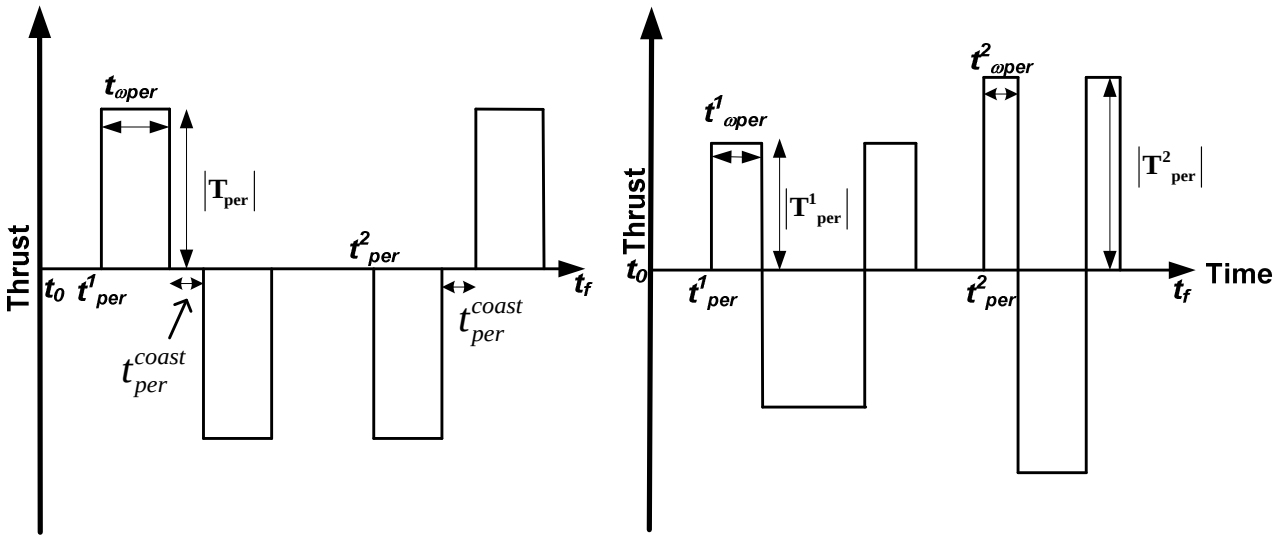

Figure 9-21 Alternative thrust profiles: complimentary bang-coast-bang (left) and multiple complimentary bang-bang (right)

A final avoidance trajectory generation method to consider would be removing the constraint that the avoidance thrust be normal to the nominal trajectory. This would allow the optimisation algorithm to choose any direction with which to point the avoidance thrust. Whilst this adds a great deal of flexibility to the possible avoidance manoeuvres it increases the complexity of the knowledge required to execute the manoeuvre. A much more accurate model of the spacecraft would be required. Every thrust direction chosen would have a different magnitude of thrust available due to the placement of thrusters around the body of the spacecraft. Furthermore, this is coupled with the reduction in the availability of thrust when the nominal thruster firings are occurring. Due to these additional spatial and temporal

complexities, the constraint of only using an avoidance thrust normal to the nominal trajectory has been placed.

9.5 Chapter Summary

This chapter dealt with the requirements for the Trajectory Assignment Module (TDM), its implementation, its output and limitations. The chapter began by introducing the TDM's position within the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) and the proximity constraints that the TDM is designed to eliminate. This was followed by a review of the literature pertaining to this problem. The review introduced a number of novel ways of incorporating avoidance criteria into manoeuvre planning problems and highlighted the complexity of the methods required to solve these issues.

The TDM algorithm was developed using the initial and PAM optimised final spacecraft positions as boundary conditions for the spacecraft trajectories. It was shown how the TDM assesses the PAM trajectories for proximity violations and executes a trajectory design algorithm should violations exist. This algorithm optimises the size and timings of avoidance manoeuvres for each spacecraft. In addition the trajectory design algorithm was able to optimise for fuel management in a similar way to the PAM. An example manoeuvre demonstrated how the TDM was able to find optimal and safe formation reconfiguration trajectories for extreme initial conditions simulating multiple collisions.

Analysis of the use of the PatternSearch (PS) optimisation algorithm within the TDM revealed that its output was not reliably optimised and depended highly on the initial conditions input to the algorithm. This was due to the complex nature of the TDM cost function. It was shown that an algorithm such as the Genetic Algorithm (GA) was better at finding more optimal solutions than the PS and was less constrained by the initial conditions imposed on the algorithm. This was due to the population-based nature of the GA that allows it to search the solution-space more effectively than the PS. The only drawback to the GA was the increased calculation time required to find a solution. This calculation time was deemed suitable however for use within the SepM-MPA.

Analysis of the TDM on DARWIN-like manoeuvres showed that there was an up to 44% chance that a PAM optimised manoeuvre would require further optimisation by the TDM depending on the nature of the PAM optimisation. On average the TDM required the use of less than 50% extra fuel (over the PAM optimised value) to execute collision and thruster plume avoidance manoeuvres. Analysis of the affect of μ_{TDM} on the algorithm performance showed a similar response as μ_{PAM} for the PAM however the fuel penalty involved when $\mu_{TDM} \neq 0$ discouraged its use as a way to perform fuel balancing for the trajectories.

The final part of this chapter introduced and discussed a few future work concepts that could be applied to the TDM. A limitation with the nature of the TDM cost function was raised that was shown to prevent an initial population for the GA from being generated for certain initial and final spacecraft position combinations. It was suggested that removing the infinity cost for proximity violations and representing the control of those violation as constraints within the optimisation problem would greatly simplify the cost function and 'smooth-out' the solution space. A number of alternative trajectory generation methods were also introduced but many involved a more complex design than that of the TDM.

The TDM represents one way optimising the spacecraft trajectories between PAM optimised spacecraft positions. The TDM integrates well into the SepM-MAP from sub-section 6.4 and represents the final optimisation module within the reconfiguration manoeuvre planning section of the Architecture. The following chapter describes the station-keeping manoeuvre planning section within the SepM-MPA data flow, the Station-keeping Module (SKM).

10. STATION-KEEPING MODULE

The inclusion of the L_2 dynamic environment within the Separate Modular Manoeuver Planning Architecture (SepM-MPA) is provided by the Station-keeping Module (SKM). The goal of the SKM is to decide whether it is necessary to perform a station-keeping manoeuvre as part of the upcoming reconfiguration phase and if so optimally plan the required manoeuvre. The position of the SKM within the SepM-MPA is given in Figure 10-1.

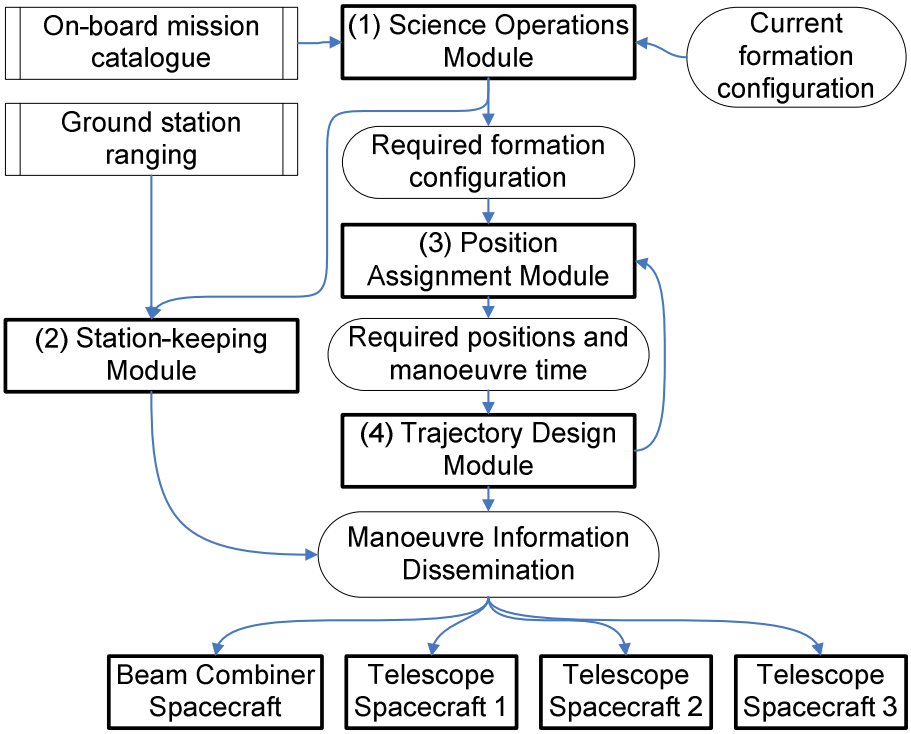

Figure 10-1 Separate Modular Manoeuvre Planning Architecture (SepM-MPA) reproduced from Figure 6-10

There is a requirement for all spacecraft wherever they operate to perform periodic station-keeping manoeuvres to ensure they follow the trajectories designed for the mission. This is due to the differences of the dynamic environment as modelled to generate the reference trajectories and the real dynamic environment experienced by the spacecraft. These differences manifest themselves as disturbances that range from natural disturbances like varying gravitational effects, varying solar radiation pressure and random micrometeor impacts, to spacecraft-made disturbances like translational and attitude manoeuvres. If periodic corrections are not made to counter these disturbances then the spacecraft may deviate so far from their reference trajectory (RT) that it becomes impossible to perform the mission or successfully return to the RT with the remaining amount of fuel on board. In this chapter the station-keeping requirements for the DARWIN mission will be introduced and a method introduced to include this station-keeping as part of a unified manoeuvre planning architecture.

10.1 Station-Keeping at Libration Points

As mentioned in Chapter 2, the DARWIN mission design requires the formation to be placed at the L_2 libration point of the Sun/Earth system for a number of operational reasons. In this sub-section a description of the libration points will be given and a number of methods for station-keeping at the libration points will be introduced.

10.1.1 The Libration Points

The libration points (or Lagrange points) are manifestations of the simplified dynamics of the three-body problem. The three-body problem describes the dynamics of an infinitesimally small object within the gravitation field of two much larger gravitating objects and is most often viewed as a spacecraft within the gravitational fields of the Sun and Earth or the Earth and the Moon. If the two larger bodies are constrained to orbit each other in circular orbits then the problem is called the Circular Restricted Three-Body Problem (CR3BP) whilst the Elliptical Restricted Three-Body Problem (ER3BP) describes the dynamics when the two massive bodies orbit in elliptical orbits. A complete derivation of the CR3BP and the ER3BP can be found in Wie (1998) and will not be reproduced here. However, salient equations will be used to illustrate the phenomena of the libration points.

Within the CR3BP two massive primaries orbit each other with a constant angular velocity

$$n = \sqrt{G\left(M_1 + M_2\right)/D^3} \tag{10.1}$$

with M_1 and M_2 representing the masses of the primaries, G the gravitational constant and D the distance between the primaries. The co-ordinate system for the CR3BP can be found in Figure 10-2.

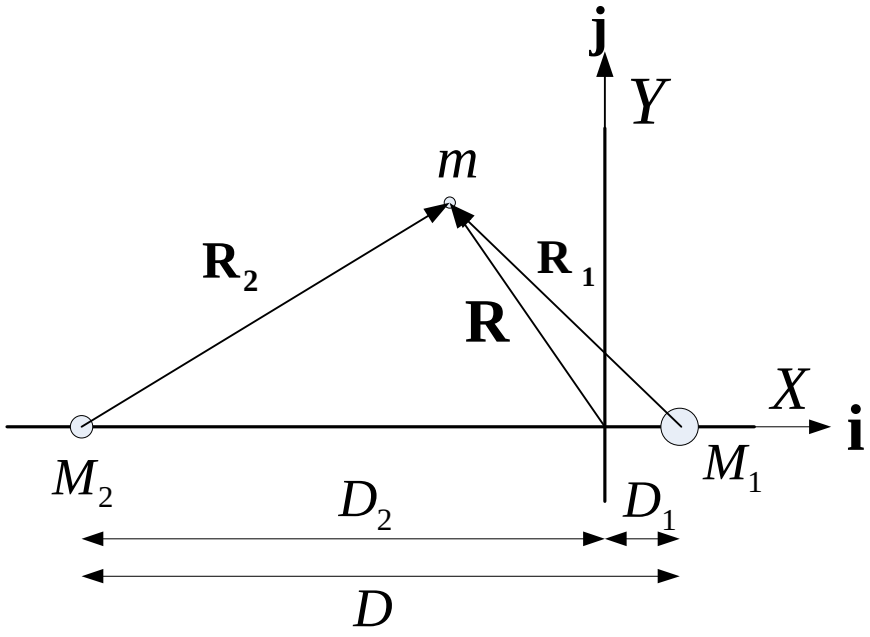

Figure 10-2 The Circular Restricted Three-Body Problem (reproduced from Wie, 1998) where M_1 and M_2 are the primary and secondary masses and the spacecraft is m.

The non-linear dynamics of the spacecraft can be written as

$$\begin{aligned} \ddot{X} - 2\dot{Y} - X &= -\frac{(1-\rho)(X-\rho)}{r_1^3} - \frac{\rho(X+1-\rho)}{r_2^3} \\ \ddot{Y} + 2\dot{X} - Y &= -\frac{(1-\rho)Y}{r_1^3} - \frac{\rho Y}{r_2^3} \\ Z &= -\frac{(1-\rho)Z}{r_1^3} - \frac{\rho Z}{r_2^3} \end{aligned} \tag{10.2}$$

where $\mathbf{R} = X\mathbf{i} + Y\mathbf{j} + Z\mathbf{k}$ is the position vector of the spacecraft from the system origin with the system rotating with angular velocity $n\mathbf{k}$,

$$\rho = M_2 / (M_1 + M_2) \tag{10.3}$$

$$r_1 = \sqrt{(X - D_1)^2 + Y^2 + Z^2} \tag{10.4}$$

$$r_2 = \sqrt{(X + D_2)^2 + Y^2 + Z^2} \tag{10.5}$$

These equations of motion are non-dimensionalised so that time is in units of $1/n$ and distances are in units of *D*.

Equilibrium between the gravitational forces and the centrifugal forces acting on the spacecraft is found when the derivatives in Equations (10.2) are equal to zero. This leads to a set of quintic equations, the solutions of which give the locations of the equilibrium positions within the CR3BP (Roberts, 2005). These locations are shown in the rotating frame in Figure 10-3. L_1, L_2 and L_3 are called the collinear libration points and share co-ordinates within the rotation frame of $Y = 0$ and $Z = 0$. L4 and L5 are dubbed the equilateral libration points and share the property $|\mathbf{R_1}| = |\mathbf{R_2}| = 1$ (in units of *D*).

Linearising the equations of motion about any of the libration points and performing a linear stability analysis shows that the equilateral L4 and L5 points are stable whilst the collinear, L_1, L_2 and L_3 libration points are unstable. This is also illustrated in Figure 10-3 with the equipotential contours providing a two-dimensional 'map' of the potential field in the CR3BP. The blue arrows represent directions of increasing potential whilst the red arrows represent directions of decreasing potential. From this equipotential contour plot it is clear to see the potential wells surrounding the equilateral libration points allowing uncontrolled objects to remain near those libration points. For the collinear libration points however Figure 10-3 shows saddle points preventing uncontrolled objects from maintaining positions near these points. Though the collinear libration points are unstable it has been shown that there exist semi-stable trajectories within the CR3BP that allow objects to remain within the vicinity of a libration point for long periods of time (Wie, 1998). Expanding the complexity of the problem to the ER3BP still reveals the existence of the libration points however in this model the locations of these points are dynamic within the rotating frame whereas in the CR3BP they are static (Wie, 1998).

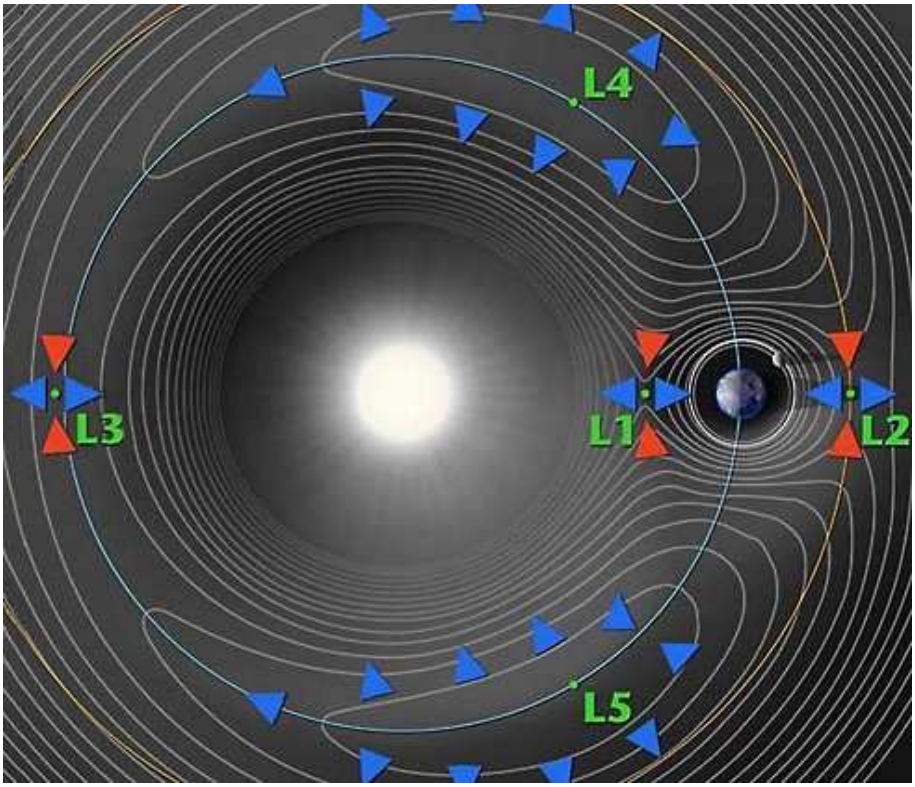

Figure 10-3 The libration points of the Sun – Earth/Moon system in the rotating frame (WMAP website, 2009). The collinear points, L_1, L_2 and L_3 are dynamically unstable whilst the triangular points, L_4 and L_5, are dynamically stable.

10.1.2 Motion about the Collinear Libration Points

Continuing the linear analysis of the CR3BP about a collinear libration point it is possible to find a 1st order solution to the equations of motion that produce a quasi-periodic trajectory (called a Lissajous trajectory) about the libration point. The solution (Wie, 1998) found is

$$\begin{aligned} x(t) &= -A_x \sin \omega_{xy} t \\ y(t) &= -A_y \cos \omega_{xy} t \\ z(t) &= A_z \sin \omega_z t \end{aligned} \tag{10.6}$$

where x, y and z are the spacecraft position components relative to the libration point at time t, A_x, A_y and A_z are the amplitudes of the trajectory, ω_{xy} is the in-plane frequency and ω_z the out-of-plane frequency. Equations (10.6) give a trajectory about L_2 in the Sun/Earth-Moon system as shown in Figure 10-4. The star represents the position of L_2 in the rotating system.

10.1.3 Reference Trajectory Generation

The libration points only exist within the mathematical confines of the CR3BP and the ER3BP in which natural bound motion about the collinear libration points can be calculated. In reality, any spacecraft in these regions of space are subject to gravitation forces from not only the Sun and Earth but also the moon and other planets as well. Spacecraft will also be

perturbed by solar radiation pressure (SRP) and other more random variations such as micro-meteor impacts (Grun, et al. 1985 and Smith, et al, 2004). The dynamic environment in the vicinity of where the libration points are in the CR3BP is however unique and can be utilised to create real halo and Lissajous trajectories for spacecraft using low-thrust station-keeping techniques.

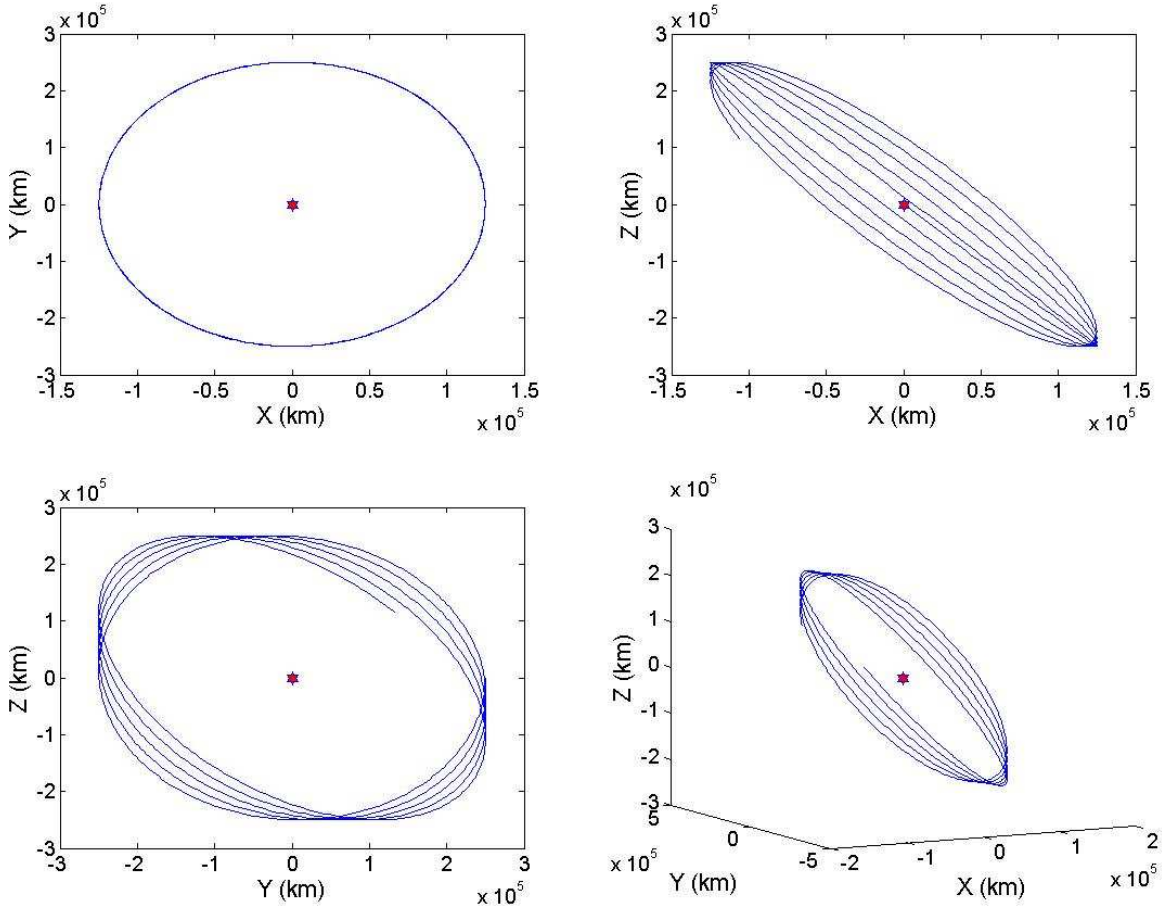

Figure 10-4 Lissajous trajectory around L_2 (star) in the Sun/Earth-Moon system calculated using the equations of motion of the CR3BP linearised about the collinear libration point, Equation (10.6), for a Z-amplitude trajectory of 25000km.

The reference trajectory (RT) for libration point missions is a pre-planned trajectory that is used as a guide for the positioning of the spacecraft throughout the mission. Perturbations due to gravitational effects, SRP etc... will displace the spacecraft from this trajectory. Depending on the requirements of the mission the spacecraft may have to track the RT tightly or the spacecraft may be allowed to drift away from the RT with loose tracking. Regardless of the type of station-keeping required the RT is an important part of the mission planning and station-keeping strategy. Consequently over the years there have been a number of methods developed to generate reference trajectories that are as close as possible to the actual motion the spacecraft will experience at the libration point. Increased accuracy leads to a reduction in the total amount of ΔV required to maintain the trajectory and therefore a reduction in mission cost.

The easiest reference trajectories to generate are those using the first order solutions of the CR3BP linearised at the libration point, Equations (10.6). From these solutions it is possible to create Lissajous and halo-type orbits around a libration point. These solutions though do not present very practical reference orbits for real-world applications as the assumptions made to generate them are numerous resulting in a relatively large ΔV required to track the trajectory. Much work therefore has gone into generating more accurate nominal trajectories.

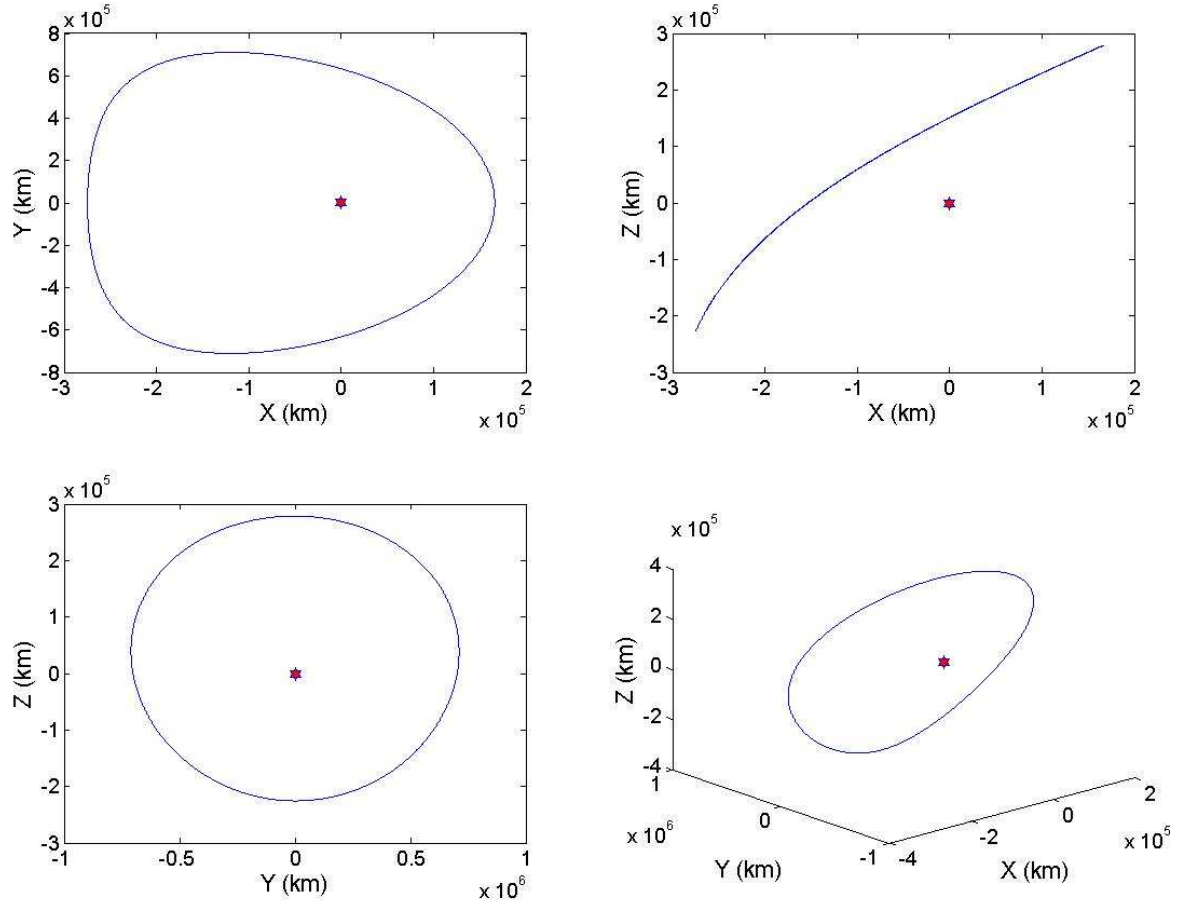

Figure 10-5 Halo trajectory around L_2 in the Sun/Earth-Moon system (star) calculated using the method in Richardson (1980a) for a Z-amplitude trajectory of 25000km.

It is possible to get a more precise analytical representation of periodic orbits using equations derived by Richardson (1980a). From the Lagrangian formulation of the non-linear equations of motion of the CR3BP at the libration point a Lindstedt-Poincare method is used to find a third-order analytical solution that produces closed 'halo'-type orbits. An example of this type of RT can be seen in Figure 10-5. The trajectory is generated around L_2 in the Sun/Earth-Moon system for a Z-amplitude of 250,000 km and one orbit of the trajectory takes ~6 months. The Richardson method is used to derive reference trajectories by Campbell, Zanon and Kulkarni (2004) and Kulkarni, Campbell and Dullerud (2006) and provides the starting point for the derivation of an analytical solution to the relative spacecraft dynamics at L_2 (Segerman and Zedd, 2003). The results Richardson published however have

been found to contain minor errors (Thurman and Wolfolk, 1996). Another semi-analytical method for computing halo orbits can be found in Simo et al (1987).

Owing to the lack of accuracy available using analytical methods, a number of techniques for generating libration point orbits using numerical methods have emerged. A popular way of achieving this was devised by Howell and Pernicka (1988). Target points along the orbit (every half orbit) are selected from a third order analytical solution. Trajectory segments between target points are found by integrating the non-linear equations of the CR3BP for half an orbit and then differentially correcting the initial velocity components until the integration and target points spatially meet. When put together, these segments create a trajectory that is continuous in position but discontinuous in velocity at the target points. To reduce the ΔV at each target point the target positions and segment times are corrected iteratively until the ΔVs are below some threshold value. The new trajectory segments are patched together to essentially give a trajectory that is continuous in both position and velocity. This method is used by Xin, Balakrishnan and Pernicka (2004), Carlson, Pernicka and Balakrishnan (2004) and Pernicka, Carlson and Balakrishnan (2006) to compute nominal trajectories in the CR3BP and modified by Sengupta and Vadali (2005) to compute a nominal trajectory in the ER3BP. A similar method by Thurman and Wolfolk (1996) is used by Junge, et al. (2002) to search for halo orbits that induce only small perturbations on satellite formations to facilitate a low-energy formation-keeping strategy. This method computes closed 'halo'-type orbits around the libration point.

Further accuracy in the calculation of nominal trajectories about the libration point has been found using ephemeris files. Hamilton, Folta and Carpenter (2002) describe a system that uses ephemeris files to generate Lissajous trajectories taking into account gravitational perturbations caused by the moon and other planets, the effects of eccentricity and solar radiation pressure. Similarly, Folta et al. (2004) use a comparable system to generate an L_2 reference halo orbit for formation control design.

10.1.4 Station-keeping Techniques

The RT for a libration point mission is an essential requirement to ensure the spacecraft perform their observations over a long period of time using a minimal amount of fuel. In order to remain on the RT however libration point spacecraft are required to periodically perform manoeuvres to compensate for disturbances due to the space environment and spacecraft operations. These manoeuvres are called station-keeping manoeuvres. For formation flying missions it is essential to differentiate between station-keeping and formation-keeping. Station-keeping involves manoeuvres to correct the spacecraft position relative to a RT (i.e. a position outside the formation). Formation-keeping involves manoeuvres to correct the spacecraft position relative to a position within the formation. This distinction is required to avoid confusion between different manoeuvre types in formation flying missions.

10.1.4.1 Past and Present Missions

In 1978 the first libration point mission was the third International Sun-Earth Explorer (ISEE-3). Its mission at the Sun-Earth L1 point kept it in a halo orbit for nearly 4-years

before it was re-tasked to examine the tail of comet Giacobini-Zinner. The RT was designed using a semi-analytical technique by Richardson (1980b) to create a halo orbit with an out-of-plane amplitude of 120,000km. Station-keeping manoeuvres were performed every 2-3 months to ensure the spacecraft remained in the vicinity of the RT. Factors affecting the selection of the date of the manoeuvre included fuel cost, duration of coast periods and required proximity of the trajectory with respect to the reference orbit (Farquhar et al, 1980).

The Solar Heliospheric Observatory (SOHO) became the second L1 mission in 1996. The nominal orbit is the same as that for ISEE-3 but in the opposite direction. SOHO uses an orbital energy balancing technique. Station-keeping manoeuvres are performed along the direction of Sun-Earth line to either increase or decrease the spacecraft's orbital energy. This allows SOHO to maintain a Lissajous trajectory close to (but not tracking) the RT designed (Dunham and Roberts, 2001). At the time of writing, SOHO is still operational and maintaining its orbit at the L1 point.

Subsequent missions, the Advanced Composition Explorer (ACE) in 1997 (Dunham and Roberts, 2001), the Microwave Anisotropy Probe (MAP) in 2001 (Rohrbaugh and Schiff, 2002) and Genesis in 2001 (Williams et al, 2000) have used similar approaches to that of SOHO in station-keeping planning, design and execution. The nominal trajectory is designed prior to the mission using a number different analytical or numerical methods as outlined in section 10.1.3. During the mission, the station-keeping manoeuvres are planned using a targeting method. The two most common targeting methods are the Target Point strategy (Howell and Pernicka, 1993) and the Floquet Mode approach (Gomez et al, 1998) and both provide discrete impulsive manoeuvre strategies. The manoeuvre is executed at the calculated time, in the calculated direction and for the calculated duration. The spacecraft orbit is then tracked and the results used to plan the next station-keeping manoeuvre.

10.1.4.2 Proposed Station-keeping Methods for Future Missions

With the on-going success of the libration point missions described above there has been an increase in the number of mission concepts based at the collinear libration points in the Sun/Earth-Moon system, most notably L_2, e.g. Planck (ESA Planck website, 2009), Herschel (ESA Herschel website, 2009) and the James Webb Space Telescope (JWST) (ESA JWST website, 2009). This has lead to increased activity in the design of control systems for libration point station-keeping tasks.

Linear Quadratic Regulator (LQR) control for station-keeping at the Earth-Moon L_2 point is derived in Wie (1998). The first-order solution to the equations of motion of the CR3BP linearised at the L_2 point are used to calculate a reference Lissajous trajectory and a continuous linear time-invariant (LTI) LQR controller based on the same dynamics is successfully implemented. A similar controller for a Sun-Earth L_2 mission is described in Hamilton and Folta (2002). Here the RT is produced using ephemeris files but the controller is designed using the dynamics of the CR3BP linearised at the L_2 point. The designed controller is of the Linear Quadratic Gaussian-type (LQG) and is successfully demonstrated using a discrete linear time-varying dynamics matrix, output from the same program that computes the reference orbit, in a full ephemeris model. Roberts (2005) also derives LQR control for L_2 station-keeping.

In their 2004 paper, Açikmese et al. (2004) present a combined control architecture for station-keeping, formation-keeping and attitude control for the L_2 Earth Atmosphere Observatory. The separate control tasks are developed independently from each other using ephemeris file dynamics that include the Sun, Earth, moon and solar radiation pressure. For translational motion feed-forward control is used to offset known perturbations e.g. gravitational, coriolis and solar radiation pressure forces. Then precise station-keeping is provided by a Proportional-Integral-Derivative (PID) controller. The controller is successfully demonstrated in an ephemeris model following an unspecified trajectory at the Sun-Earth L_2 point with zero mean white Gaussian noise added to the states and state estimation achieved through linear Kalman filters.

A non-linear station-keeping control algorithm is presented by Wong and Kapila (2003). They use the non-linear dynamics of the CR3BP expressed relative to the Sun-Earth L_2 point to design a trajectory-tracking controller. The trajectory derived is a Lyapunov orbit (a periodic orbit in the Sun-Earth orbital plane about a libration point) using the technique presented by Thurman and Wolfolk (1996). The full-state feedback controller is derived using a "Lyapunov-type design" and successfully implemented in the non-linear dynamics model of the Sun-Earth system.

Kulkarni, Campbell and Dullerud (2006) use an H_∞ control approach to tackle the station-keeping problem around halo orbits in the CR3BP. The controller is designed by considering the halo orbit tracking problem as one of stabilizing a discrete linear time-varying (LTV) system. Using techniques developed in another paper, the controller is designed using the equations of motion linearised and discretised about the reference halo orbit. Here the third-order analytical method of Richardson (1980a) at the L1 point of the Sun-Earth system is used as the RT. The controller is successfully demonstrated in a non-linear model of the CR3BP and includes perturbations due to eccentricity and the moon's gravity, sensor noise and thruster limitations.

Optimal control methods are used by Rahmani, Jalai and Pourtakdoust (2003) to derive a station-keeping controller in the CR3BP. The reference orbit is found using the method by Gomez et al. (2001). The control acceleration is calculated using an "iterative numerical technique" called the variation of extremals. The designed controller is successfully implemented in a non-linear dynamics model of the CR3BP around the Sun-Earth L1 point.

10.2 Reference Trajectory for the Station-keeping Module

Mission requirements for DARWIN state that the formation must remain within the vicinity of the L_2 point in the Sun/Earth-Moon system. The mission itself does not require any special reference trajectory (RT) and so one must be generated that can be used within the Station-keeping Module (SKM). A number of different techniques were used to generate the RT for the SKM. Figure 10-4 shows Lissajous trajectory generated using the method described in Wie (1998) and Figure 10-5 shows a Halo trajectory (of the same Z-amplitude) generated using the method in Richardson (1980a). Many of the station-keeping methods described in section 10.1.4 however require more accurate reference trajectories to ensure they can generate station-keeping manoeuvres that use the least amount of fuel as possible.

One of the more robust and popular is the numerical method devised by Howell and Pernicka (1988) and it is this method that was adopted for the SKM RT.

10.2.1 Numerical Reference Trajectory Generation

The method for generating accurate reference trajectories by Howell and Pernicka (1998) is reproduced in this sub-section for clarity. Only the necessary equations and techniques are given but the full derivation is given in the paper (Howell and Pernicka, 1998).

The equations of motion used are equivalent to the non-linear equations of motion within the CR3BP, as in Equations (10.2), but written as partial derivatives of the pseudo-potential, U:

$$\begin{aligned} \ddot{X} - 2\dot{Y} &= \frac{\partial U}{\partial X} \\ \ddot{Y} + 2\dot{X} &= \frac{\partial U}{\partial Y} \\ \ddot{Z} &= \frac{\partial U}{\partial Z} \end{aligned} \tag{10.7}$$

where $\mathbf{R} = X\mathbf{i} + Y\mathbf{j} + Z\mathbf{k}$ is the position vector of the spacecraft from the system origin with the system is rotating with angular velocity $n\mathbf{k}$

$$U = \frac{\left(X^2 + Y^2\right)}{2} + \frac{(1-\rho)}{d} + \frac{\rho}{r} \tag{10.8}$$

$$d = \sqrt{(X+\rho)^2 + Y^2 + Z^2} \tag{10.9}$$

$$r = \sqrt{(X-1+\rho)^2 + Y^2 + Z^2} \tag{10.10}$$

and ρ is the non-dimensionalised mass of the smaller primary, Equation (10.3). Using the column matrix

$$\mathbf{x} = \left[X, Y, Z, \dot{X}, \dot{Y}, \dot{Z} \right]^T \tag{10.11}$$

the method also utilises the state transition matrix $\Phi(t, t_0)$ where

$$\Phi(t_0, t_0) = I \tag{10.12}$$

and

$$\frac{d\Phi(t, t_0)}{dt} = A(t)\Phi(t, t_0) \tag{10.13}$$

Here *A(t)* is defined as

$$A(t) = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ U_{XX} & U_{XY} & U_{XZ} & 0 & 2 & 0 \\ U_{YX} & U_{YY} & U_{YZ} & -2 & 0 & 0 \\ U_{ZX} & U_{ZY} & U_{ZZ} & 0 & 0 & 0 \end{bmatrix} \tag{10.14}$$

where U_{XX} represents the second partial derivative of U with respect to X, etc…

In order to generate the RT using the numerical method it is necessary to find target points that act as initial guesses of the six dimensional state, **x**, at various points along the trajectory. These target points are extracted from analytical solutions to the CR3BP. Howell and Pernicka (1998) describe a number of different analytical approaches that can be used, including the first-order approach (Wie, 1997), but opt for a more accurate third-order analytical method (Richardson and Cary, 1975) to generate target points from a 200,000km Lissajous trajectory at L1 in the Sun/Earth-Moon system. The method is a two level iterative process and is described in the following sub-sections in more detail.

10.2.1.1 Numerical Reference Trajectory Generation – Stage 1

In this first stage trajectory segments between the target points are generated by integrating the state vector between the target points. These trajectory segments are then patched together using a process of differential corrections to create a trajectory that is continuous in position but, subsequently, discontinuous in velocity at the patched points.

Figure 10-6 shows an illustrative example of target points and their respective trajectory segments. The first trajectory segment starts a point ‘o’ and ends at point ‘p’ whilst the second trajectory segment starts at point ‘p’ and ends at point ‘f’. The first trajectory segment is found by numerically integrating the initial state vector $\mathbf{x}_o$ from time t_o to t_p to the point ‘p$_*$’. The positions ‘p’ and ‘p$_*$’ do not coincide since the Lissajous trajectory ‘o’ and ‘p’ were extracted from does not actually exist in the non-linear CR3BP. A differential corrections procedure is used to change the velocity at ‘o’ and the segment duration ($t_p - t_o$) so that ‘p$_*$’ and ‘p’ coincide within some small tolerance.

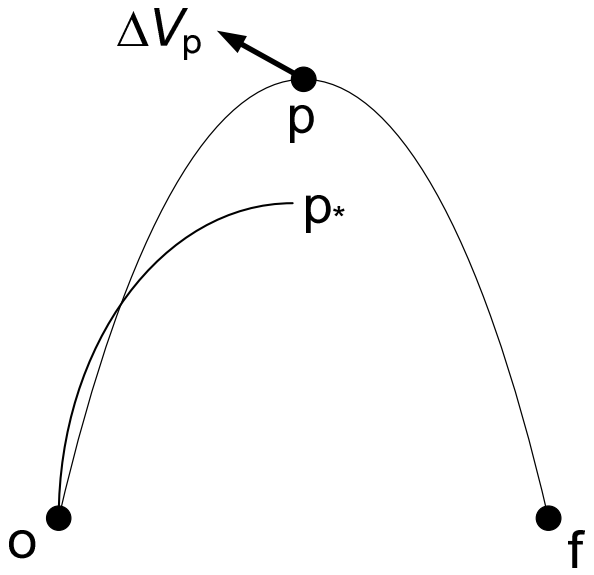

Figure 10-6 Target points (o, p, p* and f) and trajectory segment examples used for the reference trajectory generation by Howell and Pernicka (1998). The ΔV_p represents the velocity difference between segment 1 (from o to p) and segment 2 (from p to f).

This is achieved using the equation

$$\bar{u} = L^T (LL^T)^{-1} \bar{b} \tag{10.15}$$

where $\bar{b}$ is the difference in position between 'p' and 'p*',

$$\bar{b} = \left[\delta X_p, \delta Y_p, \delta Z_p \right]^T \tag{10.16}$$

$$L = \begin{bmatrix} \Phi_{14} & \Phi_{15} & \Phi_{16} & \dot{X} \\ \Phi_{24} & \Phi_{25} & \Phi_{26} & \dot{Y} \\ \Phi_{34} & \Phi_{35} & \Phi_{36} & \dot{Z} \end{bmatrix}_{p*} \tag{10.17}$$

and $\bar{u}$ describes the size of correction in position and segment duration to make

$$\bar{u} = \left[\delta X_o, \delta Y_o, \delta Z_o, \delta \left(t_{p*} - t_o \right) \right]^T \tag{10.18}$$

The segment integration is restarted at t_o with the new velocity and segment duration components and this process continues until the position 'p*' is equal to the position 'p'. The new segment, from $\mathbf{x}_p$ to $\mathbf{x}_f$ is generated in a similar manner with the velocity components $\dot{X}_p, \dot{Y}_p$ and $\dot{Z}_p$ and the segment duration $\left(t_f - t_p \right)$ differentially corrected to ensure the trajectory passes through the position 'f'. The two trajectory segments are patched together to give a trajectory that is continuous in position but discontinuous in velocity at position 'p'.

This position is called the manoeuvre point and a manoeuvre of ΔV_p is required to ensure the trajectory can be accurately followed.

10.2.1.2 Numerical Reference Trajectory Generation – Stage 2

In the second stage of the RT generation method the goal is to reduce the size of the ΔV at the manoeuvre point to below some tolerable value thus giving a resulting trajectory that is continuous in both position and velocity. To achieve this, the positions 'o', 'p', and 'f' and the segment durations $(t_p - t_o)$ and $\left(t_f - t_p\right)$ are differentially corrected. Defining the following vectors

$$\begin{aligned} \mathbf{r} &= \left[X, Y, Z\right]^T \\ \mathbf{v} &= \left[\dot{X}, \dot{Y}, \dot{Z}\right]^T \\ \mathbf{a} &= \left[\ddot{X}, \ddot{Y}, \ddot{Z}\right]^T \end{aligned} \tag{10.19}$$

and the state transition matrix sub-divided into four 3x3 sub-matrices

$$\Phi\left(t_p, t_f\right) = \begin{bmatrix} A_{pf} & B_{pf} \\ C_{pf} & D_{pf} \end{bmatrix} \tag{10.20}$$

the necessary changes in target point position and segment duration can be found using

$$\delta\mathbf{h} = -M^{T}\left(MM^{T}\right)^{-1}\Delta V_p \tag{10.21}$$

where $\delta\mathbf{h}$ is a vector describing the necessary changes.

$$\delta\mathbf{h} = \left[\delta\mathbf{r}_o, \delta t_o, \delta\mathbf{r}_p, \delta t_p, \delta\mathbf{r}_f, \delta t_f\right] \tag{10.22}$$

and

$$M = \left[M_o, M_{t_o}, M_p, M_{t_p}, M_f, M_{t_f}\right] \tag{10.23}$$

where

$$
\begin{aligned}
M_o &= D_{po} B_{po}^{-1} A_{po} - C_{po} \\
M_{t_o} &= \mathbf{a}_{p^-} - D_{po} B_{po}^{-1} \mathbf{v}_{p^-} \\
M_p &= D_{pf} B_{pf}^{-1} - D_{po} B_{po}^{-1} \\
M_{t_p} &= D_{po} B_{po}^{-1} \mathbf{v}_{p^-} - D_{pf} B_{pf}^{-1} \mathbf{v}_{p^+} + \mathbf{a}_{p^+} - \mathbf{a}_{p^-} \\
M_f &= C_{pf} - D_{pf} B_{pf}^{-1} A_{pf} \\
M_{t_f} &= D_{pf} B_{pf}^{-1} \mathbf{v}_{p^+} - \mathbf{a}_{p^+}
\end{aligned}
\tag{10.24}
$$

In Equations (10.24) the subscript 'p^+' denotes conditions at 'p' on segment two integrated backwards from 'f' whilst the subscript 'p^-' denotes conditions at 'p' on segment one integrated forwards from 'o' (as calculated during stage one).

Once a new set of target positions and segment durations, $\delta\mathbf{h}$, has been computed the entire process is repeated. Using the new $\delta\mathbf{h}$ stage one is run again to find a trajectory continuous in position. This time however the resulting ΔV_p should be much smaller than that found in the first iteration. This whole procedure is repeated until ΔV_p is reduced to within an acceptable tolerance. At the end of the entire process the resulting trajectory is one that is in essence continuous in both position and velocity.

10.2.2 Computing the Reference Trajectory

The RT for use within the SKM was to be computed using the method introduced above in sub-section 10.2.1. The target points were extracted from a trajectory computed using the third-order analytical method by Richardson (1980a). This trajectory is generated at L_2 of the Sun/Earth-Moon system with a *Z*-amplitude of 300,000 km. The trajectory is shown in Figure 10-7 along with the positions of the four target points (●). The origin of the co-ordinate system is the Sun. The positions and times of the target points are given in Table 10-1. Two of the target points (1 and 3) represent the point where the trajectory crosses the X-Y plane and the other two target points (2 and 4) represent the maximum and minimum *Z*-extent of the trajectory (respectively).

Unfortunately this is as far into the RT generation process that can be reported in this thesis. The author did write the software to generate the RT for the SKM however was unable to debug and validate it to produce a suitable RT for further analysis. Furthermore, as an accurate RT is an essential element for the station-keeping method adopted for the SKM no further coding work has been achieved on the SKM. The remaining sub-sections in this chapter describe the proposed station-keeping method and the architecture for how the SKM is envisioned to work.

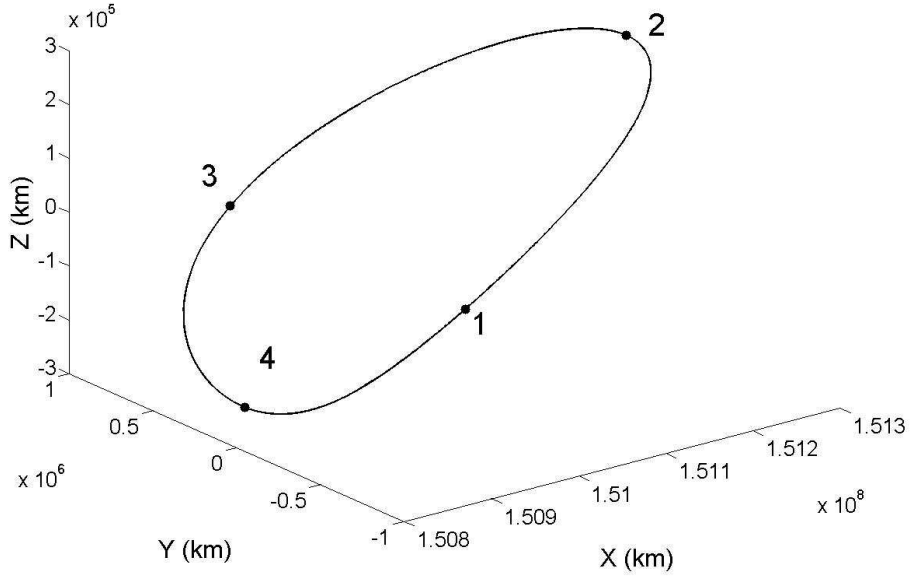

Figure 10-7 Extracted target points from a 300,000 km Z-amplitude trajectory in the Sun/Earth-Moon system used for the reference trajectory generation.

Table 10-1 Positions and times of the target points used for the reference trajectory generation

Target Point	1	2	3	4
Time (s)	5.531×10^6	9.953×10^6	1.437×10^7	1.766×10^7
X (km)	1.509×10^8	1.513×10^8	1.509×10^8	1.513×10^8
Y (km)	-7.020×10^5	0.012×10^5	7.020×10^5	0.185×10^5
Z (km)	0	2.787×10^5	0	-2.254×10^5

10.3 Station-keeping Method for the Station-keeping Module

The mission requirements for DARWIN, as well as not specifying a particular trajectory for the formation to follow, also do not stipulate how well the formation should track the reference trajectory (RT). A more simpler station-keeping approach therefore can be adopted than those introduced in sub-section 10.1.4.2 using modern control theory since they are designed to be fuel optimal, accurate reference tracking controllers. The method adopted for the Station-keeping Module (SKM) is simpler, has been used in real missions, and is called the Target Point Strategy (TPS).

10.3.1 The Target Point Strategy

The station-keeping method by Howell and Pernicka (1993) is reproduced in this sub-section for clarity. Only the necessary equations and techniques are given but the full derivation is given in the paper (Howell and Pernicka, 1993). The goal of the TPS is to keep the spacecraft close to the RT. This is achieved through the derivation of a cost function that is minimised to find a manoeuvre ΔV at a particular time, t_0, which allows the spacecraft trajectory to remain within accepted bounds of the RT. The cost function is defined as

$$J = \Delta v^T Q \Delta v + p_1^T R p_1 + v_1^T R_v v_1 + p_2^T S p_2 + v_2^T S_v v_2 + p_3^T T p_3 + v_3^T T_v v_3 \qquad (10.25)$$

where ΔV is the corrective manoeuvre at time, t_0. The vectors p_1, p_2 and p_3 represent the expected deviations of the spacecraft position from the RT at future times t_1, t_2 and t_3. The vectors v_1, v_2 and v_3 represent the expected deviations of the spacecraft velocity from the RT at the same times. These expected deviations are referenced from the projected spacecraft trajectory to the RT if no station-keeping manoeuvre is performed. Selection of the future times, or target points, are deemed arbitrary. Finally in Equation (10.25) the matrices Q, R, R_v, S, S_v, T and T_v are 3x3 weighting matrices with Q defined as symmetric positive definite and the others positive semi-definite.

The optimal control input is found through the minimisation of the cost function

$$\arg\min_{\Delta v} J \tag{10.26}$$

and this is determined through the linear equation

$$\begin{aligned}\Delta v = -\left[Q + B_{10}^T R B_{10} + B_{20}^T S B_{20} + B_{30}^T T B_{30} + D_{10}^T R_v D_{10} + D_{20}^T S_v D_{20} + D_{30}^T T_v D_{30} \right]^{-1} \\ \times \begin{bmatrix} \left(B_{10}^T R B_{10} + B_{20}^T S B_{20} + B_{30}^T T B_{30} + D_{10}^T R_v D_{10} + D_{20}^T S_v D_{20} + D_{30}^T T_v D_{30} \right) v_0 \\ + \left(B_{10}^T R A_{10} + B_{20}^T S A_{20} + B_{30}^T T A_{30} + D_{10}^T R_v C_{10} + D_{20}^T S_v C_{20} + D_{30}^T T_v C_{30} \right) p_0 \end{bmatrix}\end{aligned} \tag{10.27}$$

where the state transition matrix of the RT is sub-divided to give

$$\Phi\left(t_k, t_0\right) = \begin{bmatrix} A_{k0} & B_{k0} \\ C_{k0} & D_{k0} \end{bmatrix} \tag{10.28}$$

for target points t_1, t_2 and t_3. v_0 is the velocity difference relative to the nominal velocity at t_0 and likewise p_0 is the position difference at t_0. The calculation of the ΔV assumes that execution of the manoeuvre is instantaneous at t_0.

Howell and Pernicka (1993) also state a number of constraints that must be satisfied before the station-keeping manoeuvre be computed and executed:

- increasing magnitude of position deviation over time
- position deviation greater than a defined minimum deviation
- elapsed time from previous manoeuvre greater than a defined minimum duration
- magnitude of the computed manoeuvre is greater than a defined minimum manoeuvre magnitude

The implementation of these constraints prevents excessive and continual computation of station-keeping manoeuvres and prevents the execution of frequent small manoeuvres that may be of the same order of magnitude as the manoeuvre errors.

10.4 Station-keeping Module Approach

The goal of the Station-keeping Module (SKM) is to decide whether the upcoming reconfiguration phase is optimal to include a station-keeping manoeuvre with respect the future reconfiguration phases. These reconfiguration phases are calculated by the Science Operations Module (SOM) and extracted from the optimal tour that is output. With both the reference trajectory (RT) generation and the station-keeping methods defined this section describes how the SKM is envisioned to work.

A schematic of the data flow for the SKM is given in Figure 10-8. In-line with the manoeuvre execution constraints detailed in sub-section 10.3.1, the SKM is not executed if;

1. The Beam Combiner Spacecraft (BCS) trajectory and RT position deviation is decreasing
2. The BCS trajectory and RT position deviation is less than a defined minimum deviation and
3. The time from a previous station-keeping manoeuvre less than a defined minimum duration

The BCS is selected as the reference vehicle for the SKM calculations as it is this spacecraft that will actually be performing the station-keeping manoeuvre as planned. The other spacecraft in the formation then simply track the position of the BCS and manoeuvre accordingly to maintain their relative positions. With no SKM available the author has been unable to analyse how the definition of the minimum deviation and the minimum duration affect the results of the module. From the literature however, a minimum station-keeping manoeuvre separation of 90 days and position deviation of 5 km are reasonable values to use for the RT given (Dunham and Roberts, 2001 and Rohrbaugh and Schiff, 2002).

The main input for the SKM is the optimal tour generated by the SOM. The tour contains the observation duration for each task scheduled and so the SKM can easily determine when each reconfiguration phase within the tour will occur. The core of the SKM is the FOR loop (Figure 10-8) that sequentially selects the beginning of every reconfiguration phase for analysis. The beginning of the reconfiguration phase becomes t_0 and the algorithm progresses to select the target points t_1, t_2 and t_3. Howell and Pernicka (1993) describe the selection of the target points as ‘arbitrary’ however selecting points very far apart (in time) is likely to reduce the accuracy of the calculated ΔV. The RT (at least the initial Richardson trajectory from Figure 10-7) has a period of ~280 days and so separating the target points by 30 days encompasses almost half the desired trajectory for each station-keeping manoeuvre calculation. Unfortunately, with no working SKM to perform analysis on there is no way to ascertain how the target point temporal separation affects the magnitude of resulting manoeuvre.

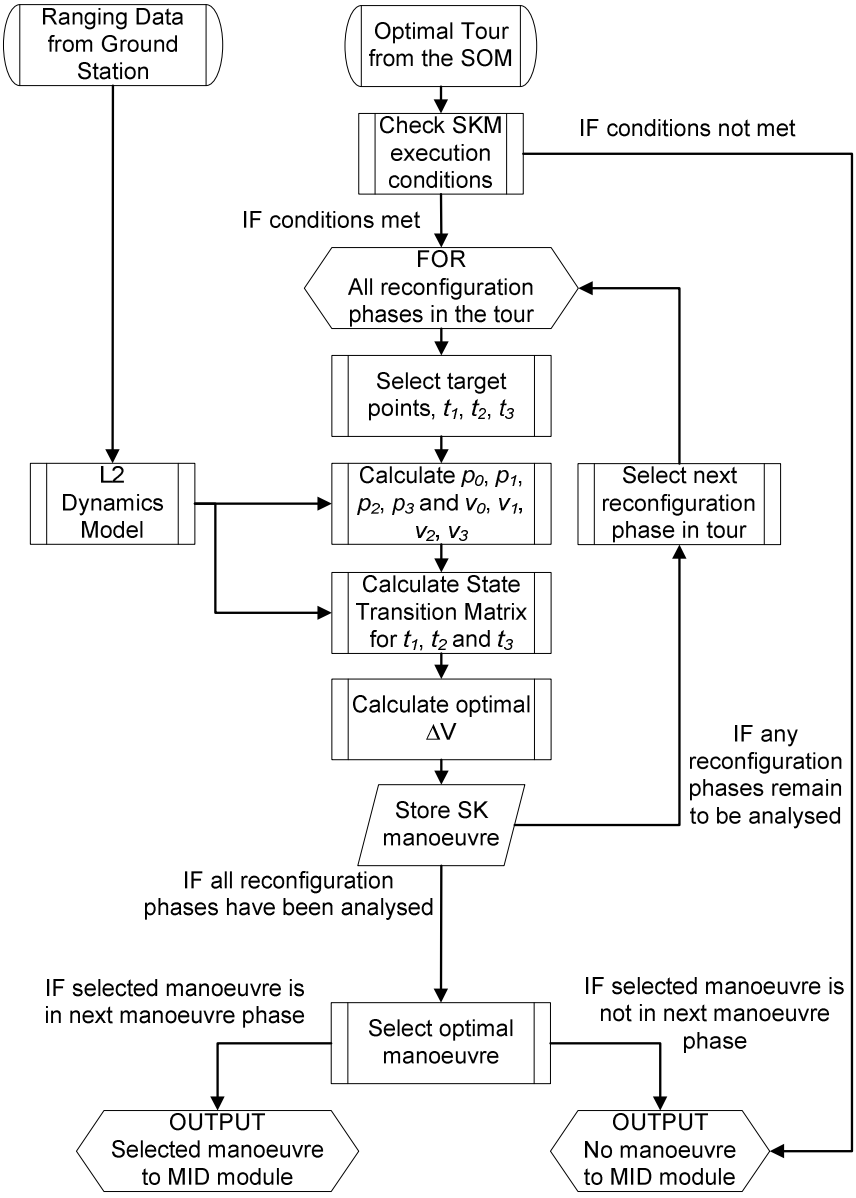

Figure 10-8 Schematic of the Station-keeping Module

Once the target points have been selected, the next step is to calculate the projected deviations from the RT the BCS will have if no manoeuvre is performed at t_0. These deviations are the vectors p_1, p_2 and p_3 and v_1, v_2 and v_3, Equation (10.25). Also required are the expected deviations (p_0 and v_0) at the manoeuvre execution time (t_0). The position and velocity deviations are calculated by integrating the current BCS trajectory using an L_2 dynamics model. The most appropriate implementation of the L_2 dynamics is the non-linear model used to generate the RT, Equation (10.7). The current BCS trajectory can be found by using a combination of the most recent ranging data from the ground station, the L_2 dynamics model and factoring in any BCS manoeuvres that have occurred since the last ranging data was received. Once the deviations at the target points have been calculated the next step is to calculate the state transition matrix of the RT at each of the target points and finally the optimal ΔV for execution at time t_0, Equation (10.27). At this stage the FOR loop ends with

the calculated manoeuvre being stored and a manoeuvre for the next reconfiguration phase generated.

Once all the reconfiguration phases within the tour have had station-keeping manoeuvres calculated for them the final stage of the SKM is to select the optimal manoeuvre for the tour. This selection follows a two-stage process. First all the manoeuvres are screened in line with the last of the manoeuvre execution constraints detailed previously in Howell and Pernicka (1993). Manoeuvres are removed from the list if the magnitude is less than a defined minimum manoeuvre magnitude. With no SKM implemented the author is unable to analyse the affect on the results this minimum magnitude will have however from the literature a magnitude of 10 cms^{-1} appears a reasonable choice (Dunham and Roberts, 2001 and Rohrbaugh and Schiff, 2002). The second stage of the SK manoeuvre selection simply involves choosing the manoeuvre from the remaining selection with the lowest ΔV associated with it. If the selected manoeuvre is due for execution at the next reconfiguration phase of the mission then the manoeuvre data is passed for Manoeuvre Information Dissemination (MID) and released to the formation when required. If the selected manoeuvre is not due for execution in the next reconfiguration phase then the selected manoeuvre is discarded and no information is passed for MID.

10.5 Future Work

The Station-keeping Module (SKM) presented in this chapter provides a method for including the station-keeping manoeuvre requirements into the manoeuvre planning strategy for the DARWIN mission. Unfortunately due to time constraints imposed on this research project the SKM remains very much a concept. To realise the full potential of the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) is it necessary for the SKM to be completed.

There are a number of steps that must be taken to realise the implementation of the SKM within the SepM-MPA:

1. The reference trajectory (RF) around L_2 must be generated. It is essential for the chosen station-keeping approach that the RF contains position, velocity and state transition matrix (STM) information. This can be found using the method developed by Howell and Pernicka (1998) and summarised in sub-section 10.2.1.
2. An L_2 dynamics model must be created. For the Target Point Strategy (TPS) station-keeping method it is necessary to integrate the trajectory forward in time to extract trajectory position, velocity and STM information at the selected target points. The suggested model is the non-linear Circular Restricted Three Body Problem (CR3BP), Equation (10.2).
3. The Target Point Strategy (TPS) for station-keeping must be coded. This is the method devised by Howell and Pernicka (1993) and summarised in sub-section 10.3.1. This algorithm should include target point selection and output an optimised station-keeping ΔV for any execution time and 3 target point combination
4. The SKM requires coding as per sub-section 10.4. This needs to include the SKM execution conditions, the TPS, manoeuvre storage and manoeuvre selection

algorithms. Once implemented the SKM should (if the execution conditions are met) plan a station-keeping manoeuvre for every reconfiguration phase in the SOM optimised tour, select the optimal station-keeping manoeuvre and pass the manoeuvre for Manoeuvre Information Dissemination (MID) if it coincides with the next reconfiguration phase.

As well as implementing the SKM within the SepM-MPA analysis needs to be performed so that the frequency of execution of the SKM is optimised. The expected time required to calculate the optimal station-keeping manoeuvre for any given tour will affect how much time is available to the SOM for its calculation of the tour. The more times the SKM is called during the mission the less time the SOM has to generate tours. This will affect the performance of the observation schedule for the mission. The frequency of execution of the SKM is determined by the SKM execution conditions (suggested values of >90 days between station-keeping manoeuvres or a minimum trajectory deviation of 5 km or a minimum station-keeping ΔV of 50 cms^{-1} are given in sub-section 10.4). Analysis of these conditions is required to optimise the frequency of SKM execution to ensure minimal station-keeping manoeuvre ΔV whilst maximising calculation time for the SOM.

10.6 Chapter Summary

This chapter dealt with the requirements for the Station-keeping Module (SKM), its implementation and limitations. The chapter began by introducing the SKM's position within the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) and highlighted the necessity for station-keeping manoeuvres. The first part of the chapter introduced the dynamic environment of the Circular Restricted Three-Body Problem (CR3BP) and introduced the concept of libration points (and trajectories in their vicinity). This was followed by a review of the literature in relation to reference trajectory generation and station-keeping techniques for libration point missions. The review emphasised the robust nature of the currently used station-keeping methods and influenced the decision to develop a new method for the SKM.

A way to generate an accurate reference trajectory (RT) about the L_2 point was introduced. The method, devised by Howell and Pernicka (1998), is a numerical method and uses the non-linear dynamics of the CR3BP to generate a Lissajous reference trajectory (RT) that is effectively continuous in position and velocity. The station-keeping method, devised by Howell and Pernicka (1993) is called the Target Point (TP) strategy and has been used in the design of a number of missions. The TPS uses the RT and a non-linear model of the dynamics of the CR3BP to minimise a cost function, resulting in a fuel optimal station-keeping manoeuvre for the manoeuvre epoch selected. The RT and TPS are used within the Station-keeping Module (SKM) to decide whether a station-keeping manoeuvre should be performed during the next reconfiguration phase. A data flow for the SKM was presented along with the conditions that bound the execution of the SKM and the execution of the station-keeping manoeuvre.

Unfortunately, due to time constraints whilst preparing this research thesis, the author was unable to complete, test and validate the SKM and so this part of the research remains

unfinished. The final section of this chapter details the future work required to realise the SKM as designed. Though the SKM remains unfinished, the author believes that the methods presented in this chapter constitute an effective first step to incorporate the station-keeping aspect of libration point missions into a unified manoeuvre planning system for formation flying missions.

11. PERFORMANCE OF THE MANOEUVRE PLANNING ARCHITECTURE

The manoeuvre planning architecture introduced in Chapter 6 represents ways of combining the manoeuvre planning optimisation modules to form a complete optimal manoeuvre planning algorithm for spacecraft formation flying missions. The optimisation modules were designed and written specifically for the Separate Modular MPA (SepM-MPA). This chapter compares the performance of the SepM-MPA to an architecture that uses the benchmark optimisation modules.

11.1 Manoeuvre Planning Architecture Re-cap and Benchmark Implementation

This sub-section provides a re-cap of the implementation of the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) and details how the Benchmark Manoeuvre Planning Architecture (B-MPA) is implemented for comparison.

11.1.1 Operation and Implementation

The SepM-MPA is the simplest of the manoeuvre planning architectures introduced in Chapter 6 and treats each of the optimisation modules as separate entities. Figure 11-1 shows a schematic of the SepM-MPA. To summarise its operation, before every manoeuvre:

- The Science Operation Module (SOM) calculates a tour that maximises the number of completed observations that can be achieved within a minimum tour duration.
- The first observation in the tour is chosen and the required formation configuration to perform that observation is defined.
- The Position Assignment Module (PAM) optimises the spacecraft post-manoeuvre positions to ensure the required formation configuration is met whilst optimising the manoeuvre duration, fuel consumption and fuel balancing as required.
- The Trajectory Design Module (TDM) uses the spacecraft post-manoeuvre positions and optimised manoeuvre duration to calculate a safe trajectory for each spacecraft that avoids collisions, thruster plume impingement and minimises fuel consumption. The optimised trajectories are then passed for manoeuvre information dissemination.
- Separately, the Station-keeping Module (SKM) uses the schedule from the SOM to plan station-keeping manoeuvres within that schedule. If the next manoeuvre phase requires a station-keeping manoeuvre, the module calculates the size and direction of the manoeuvre required and passes it for manoeuvre information dissemination.
- At the manoeuvre information dissemination (MID) stage, manoeuvre details are communicated to each individual spacecraft. Station-keeping manoeuvre information is sent first and comprises of open-loop thruster commands for the BCS. Reconfiguration manoeuvre information is then sent. For the BCS the information is open-loop thruster commands. For the TS the information is a desired trajectory relative to the BCS.

- Manoeuvres are executed simultaneously by each spacecraft. For station-keeping manoeuvres the BCS follows the open-loop thruster commands whilst the TSs track the position of the BCS and manoeuvre to maintain formation configuration. For the reconfiguration manoeuvres the BCS follows the open-loop thruster commands whilst the TS track the BCS and manoeuvre to follow the desired trajectory relative to the BCS.

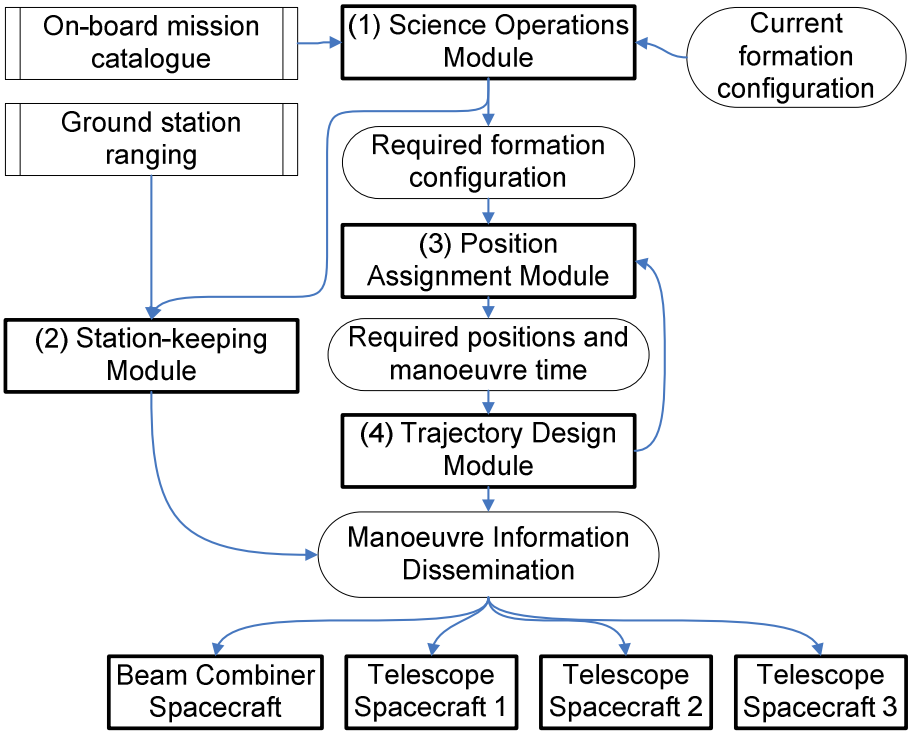

Figure 11-1 Separate Modular Manoeuvre Planning Architecture (SepM-MPA) reproduced from Figure 6-10

The SepM-MPA represents a simple way to perform formation flying manoeuvre planning whilst incorporating manoeuvre scheduling, station-keeping, collision avoidance, thruster plume avoidance, manoeuvre duration and fuel consumption optimisations and fuel balancing optimisation within a computationally non-intensive environment. Due to the incompletion of the Station-keeping Module (SKM) this module is omitted from the analysis to follow. The effect of this on the MPA is discussed in the next sub-section.

11.1.2 The Effect of the SKM Removal from the Manoeuvre Planning Architecture

The inclusion of the station-keeping module within the manoeuvre planning architecture was to incorporate stationkeeping manoeuvre planning within the architecture for planning the reconfiguration manoeuvres. This allows the effects that the reconfiguration manoeuvres have on the absolute trajectory of the spacecraft to be included within the manoeuvre planning environment. Over the course of the mission this results in a manoeuvre profile that involves long periods of multiple reconfiguration manoeuvres (planned in a zero gravity environment) punctuated by periodic station-keeping manoeuvres (planned in inertial space). The ratio of reconfiguration to stationkeeping manoeuvres will decrease throughout the

mission as the number of short observations decreases. For the comparison of the manoeuvre planning architectures in this chapter it is necessary to remove the stationkeeping module from the MPA. With the SKM incomplete as described in Chapter 10 it becomes infeasible to include it, or any aspect of it, in the MPA comparison. The effects of this removal can be viewed at two levels, the mission level and the individual manoeuvre level.

At the individual manoeuvre level the removal of the SKM from the MPA will have little effect on the performance of the MPA. When the MPA is executed the SOM is run followed by the SKM. The PAM and TDM calculate their manoeuvres within the local 'flat' space in the vicinity of the formation (as discussed in sub-section 6.3.2). The results of the PAM/TDM optimisations are not affected by the SKM. The removal of the SKM therefore will have no effect on the PAM/TDM manoeuvre plans. On occasion where the SKM does have an output, it is actioned at the manoeuvre information dissemination stage prior to the reconfiguration manoeuvre. The planned fuel consumption for the station-keeping manoeuvre will have been incorporated in the PAM/TDM planning so executing a station-keeping manoeuvre has no effect on the reconfiguration manoeuvre planning. The only effect that the removal of the SKM will have on the MPA on a manoeuvre-by-manoeuvre basis is an increase in calculation performance (since the SKM code will not be executed). This cannot be measured as there is no indication of how 'costly' running the SKM code is. At the mission level, this lack of influence of the SKM on the MPA manoeuvre plans can be extended for tours of reconfiguration manoeuvres in-between stationkeeping manoeuvres.

The influence of the SKM only comes into play when analysing tours of a long enough duration to require the execution of a stationkeeping manoeuvre (i.e. greater than 3 months). Without the SKM, no stationkeeping manoeuvre will be planned and executed and the formation will drift away from the reference trajectory due to the gravitational environment, solar radiation pressure and the execution of the reconfiguration manoeuvres. There is no doubt that the requirement for stationkeeping manoeuvres is essential to the success of the mission, however since this is a requirement regardless of the manoeuvre planning architecture adopted a fair comparison of manoeuvre planning architectures can be made over manoeuvre tours with timescales less than the frequency of the required station-keeping manoeuvres.

To model the MPA minus SKM over timescales longer than 3 months would necessitate the assumption that stationkeeping manoeuvres were being planned and executed externally and independently to the MPA. Stationkeeping manoeuvres would have to be executed in-between MPA runs and this would create differences between the expected and actual initial formation state for the manoeuvre planning after the stationkeeping manoeuvre. The MPA however is designed to cope with these differences as it already does not account for formation-keeping manoeuvres, attitude manoeuvres and other un-planned manoeuvre contingencies which cause similar discrepancies.

The incorporation of the SKM enhances the MPA by allowing the planning of stationkeeping manoeuvres to influence (and be influenced by) formation reconfiguration manoeuvres and the science observation schedule in an autonomous architecture. The removal of the SKM however will by no means affect the comparison of manoeuvre planning architectures especially over timescales less than 3 months or with the assumption that stationkeeping is planned and executed externally to the MPA.

11.1.3 Benchmark Manoeuvre Planning Architecture

In this chapter the SepM-MPA will be compared to a similar manoeuvre planning architecture that uses the benchmark planning modules as its core, the B-MPA. This is to examine whether the results from the SepM-MPA warrant the increased computational requirements to achieve them. A recap of the benchmark planning modules is provided in Table 11-1. The data flow for the B-MPA is exactly the same as for the SepM-MPA from Figure 11-1.

Table 11-1 Planning Modules of the Benchmark Manoeuvre Planning Architecture (BMPA)

Benchmark Planning Modules	Science Operations Module (BSOM) Position Assignment Module (BPAM) Trajectory Design Module (BTDM)	Generates the tour by always selecting the science task with the shortest completion time (observation + manoeuvre duration). Reconfiguration is a simple rigid-body rotation of the formation followed by baseline/configuration changes as required. Avoidance manoeuvre parameters, *X*, chosen from first viable set found in function '*TDM_findx0*'.

11.2 Comparison Set-up

The SepM-MPA incorporates the manoeuvre planning algorithms from the SOM, PAM and TDM. Analysis of the comparison of each of the modules with their respective benchmark has been performed on a manoeuvre-by-manoeuvre basis in chapters 7, 8 and 9. A similar comparison for the SepM-MPA with the B-MPA however would be pointless as the data output by the SOM and BSOM would be different thus making PAM/BPAM and TDM/BTDM comparison meaningless. To properly compare the SepM-MPA with the B-MPA requires the analysis of a tour of manoeuvres with the final fuel consumption, fuel balancing and tour duration being the comparison metrics analysed.

Table 11-2 lists the initial parameters used in the set-up of the analysis. Many of the parameters are identical to those used in the analyses in Chapters 7, 8 and 9 however there are some noticeable differences. Within the SOM the nature of the tour length has been changed. Previously in Chapter 7 the tour length was defined as a minimum tour time, T_{target}, and successive observations were added to the tour until the tour duration exceeded the value set for T_{target}. With this method the user is unable to control the number of observations that are included in any tour since the tour duration is the terminating factor. For this SepM-MPA analysis a different metric within the SOM has been used to terminate each tour calculation. Each tour within the SOM is limited to being 10 observations long. This speeds up the tour generation part of the algorithm and thus more tour iterations are generated. Though each tour within the SOM will only be 10 observations long it is desirable to examine many more manoeuvres within the context of the SepM-MPA. The SepM-MPA analysis will therefore simulate a 50 observations tour with each repetition of the SOM limited to a calculation time of 7200 sec. In addition it was decided to select TF2 as the taskflag file as in this period of

the mission there are less detection tasks available thus increasing the chance of formation reconfiguration manoeuvres (i.e. linTTN to triTTN and vice versa) to be scheduled.

Initial parameters for the PAM and TDM are also given in Table 11-2. Most of these initial parameters are the same as used in the analyses in Chapters 8 and 9 with spacecraft mass and thruster data taken from Karlsson, et al. (2004) and D'Arcio (2005). Initial spacecraft fuel has been set to reflect fuel imbalance across the fleet. The initial BCS fuel has been set so that it is not the highest across the fleet thus minimising the chances of a large single spacecraft fuel consumption anomaly as seen in sub-section 8.3.6.1. Finally fuel balancing weight, μ_{PAM} is set to 0 and 0.6 for fuel minimising and fuel balancing analysis respectively and μ_{TDM} has been set to zero for fuel minimising.

Table 11-2 Parameters for the SepM-MPA analysis

Module	Parameter	Value
SepM-MPA	Tour Length	50 observations
SOM	Taskflag File	TF2
	Startstar	200
	StartTask	Detection
	Initial anti-sun longitude (β)	β of star 200
	Standard Deviation (σ)	0.07
	Calculation Time	7200 s
	SOM Tour Length	10 observations
PAM	Spacecraft Identifiers (i)	{BCS, TS1, TS2, TS3}
	Spacecraft Mass (m_i)	{1100, 900, 900, 900} kg
	Initial Fuel Mass (mf_i)	{4.00, 4.05, 4.105, 3.9} kg
	Maximum Nominal Thrust (α_i)	{6, 6, 6, 6} mN
	Thruster I_{sp}	3300 s
	PAM Fuel Balancing Weight (μ_{PAM})	0 and 0.6
TDM	Maximum Perpendicular Thrust ($\alpha_{per,i}$)	{10, 10, 10, 10} mN
	Collision Avoidance Minimum	10 m
	Thruster Plume Half-angle	5 °
	Thruster Plume Length	10 m
	TDM Fuel Balancing Weight (μ_{TDM})	0

11.3 Manoeuvre Planning Architecture Comparison

Results of the comparison between the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) and the Benchmark Manoeuvre Planning Architecture (B-MPA) can be found in this sub-section. The tours generated were each 50 observations long starting from star 200 in the catalogue with the linTTN formation configuration.

11.3.1 Tour Comparison

Comparison between the SepM-MPA and B-MPA tours can be found in Table 11-3. Table 11-3 gives the numerical data for the tours detailing the stars visited, their order and the

Table 11-3 Generated tours from the SepM-MPA and B-MPA tour comparison

Manoeuvre Number	SepM-MPA Minimising			SepM-MPA Balancing			B-MPA		
	Star	Task	Duration (days)	Star	Task	Duration (days)	Star	Task	Duration (days)
	200	1	n/a	200	1	n/a	200	1	n/a
1	230	1	2.42	212	1	2.44	230	1	2.42
2	249	1	2.42	228	1	2.42	249	1	2.42
3	230	1	2.42	255	1	2.42	230	1	2.42
4	249	1	2.42	228	1	2.42	249	1	2.42
5	255	1	2.43	255	1	2.42	230	1	2.42
6	228	1	2.42	228	1	2.42	249	1	2.42
7	212	1	2.42	255	1	2.42	255	1	2.43
8	228	1	2.42	249	1	2.43	228	1	2.42
9	255	1	2.42	230	1	2.42	212	1	2.42
10	228	1	2.42	249	1	2.42	228	1	2.42
11	255	1	2.42	230	1	2.42	255	1	2.42
12	249	1	2.43	249	1	2.42	228	1	2.42
13	252	4	3.95	230	1	2.42	255	1	2.42
14	246	4	3.96	223	4	3.96	246	4	3.96
15	230	1	2.43	298	1	2.44	298	1	2.44
16	223	4	3.96	270	4	3.97	270	4	3.97
17	297	1	2.44	314	1	2.43	298	1	2.43
18	298	1	2.42	230	1	2.42	314	1	2.42
19	319	1	2.42	249	1	2.42	298	1	2.42
20	298	1	2.42	230	1	2.42	319	1	2.42
21	319	1	2.42	249	1	2.42	256	4	3.98
22	298	1	2.42	230	1	2.42	319	1	2.44
23	314	1	2.42	249	1	2.42	297	4	3.96
24	327	4	3.97	255	1	2.43	252	4	3.98
25	274	4	3.97	314	1	2.44	274	4	3.98
26	270	4	3.98	298	1	2.42	327	4	3.97
27	287	4	6.63	319	1	2.42	320	4	6.63
28	345	4	3.98	298	1	2.42	345	4	3.97
29	357	4	3.97	319	1	2.42	357	4	3.97
30	344	4	3.97	298	1	2.42	344	4	3.97
31	298	1	2.43	297	4	3.96	287	4	6.65
32	319	1	2.42	344	4	3.97	374	4	3.98
33	298	1	2.42	357	4	3.97	376	2	10.89
34	319	1	2.42	345	4	3.97	404	6	3.43
35	298	1	2.42	327	4	3.96	400	4	3.96
36	319	1	2.42	320	4	6.63	384	4	3.96
37	298	1	2.42	287	4	6.63	404	3	9.47
38	319	1	2.42	404	6	3.46	417	2	10.89
39	314	1	2.42	400	4	3.96	1	4	3.96
40	404	6	3.45	374	4	3.96	9	2	10.90
41	374	4	3.97	404	6	3.50	431	2	10.91
42	384	4	3.95	314	1	2.45	402	2	10.90
43	345	4	3.96	357	4	3.97	40	1	2.46
44	357	4	3.97	345	4	3.97	28	2	10.90
45	400	4	3.97	374	4	3.96	62	1	2.43
46	344	4	3.99	384	4	3.95	63	1	2.41
47	404	6	3.44	400	4	3.96	62	1	2.41
48	400	4	3.96	1	4	3.97	63	1	2.41
49	374	4	3.96	369	4	6.66	67	4	3.96
50	384	4	3.95	404	3	9.48	63	1	2.42

observation task performed for three tours; the SepM-MPA minimising tour when $\mu_{PAM} = 0$, the SepM-MPA balancing tour when $\mu_{PAM} = 0.6$ and the B-MPA tour. The tours differ greatly though the SepM-MPA minimising and B-MPA tours share the same first 4 observations. The difference in the two SepM-MPA tours highlight partly the stochastic nature of the tour calculation process and partly the different manoeuvre durations experienced for fuel minimising and fuel balancing manoeuvres. Comparing the star numbers near the end of the tours shows both SepM-MPA tours scheduling observations in roughly the same area of the sky but the B-MPA scheduling observations in an area further round in ecliptic longitude. This indicates the B-MPA tour is of greater duration than the SepM-MPA tours.

Tour durations for the three tours are given in Table 11-4. For each tour, two durations are stated. The actual tour duration is that calculated using the manoeuvre durations from the PAM/BPAM. The SOM/BSOM tour duration is calculated using the fixed formation angular rate, i.e. $\dot{\alpha} = 0.02$ °s^{-1}. The discrepancy between the two shows that the fixed angular rate used in the SOM/BSOM underestimates the PAM/BPAM manoeuvre duration by on average 27 min for the SepM-MPA minimising and B-MPA manoeuvres and on average 60 min for the SepM-MPA balancing manoeuvres. The differences between these underestimates highlight how much extra time on average is required to execute a fuel balancing manoeuvre over a fuel minimising/benchmark manoeuvre. These average underestimates can be fed back into the SOM in future versions to make the SOM tour calculation more accurate (without the explicit calculation of each manoeuvre duration using the PAM).

Table 11-4 Comparison tour performance data for the SepM-MPA and B-MPA tours

	SepM-MPA Minimising		SepM-MPA Balancing		B-MPA	
	Actual	SOM	Actual	SOM	Actual	BSOM
Tour Duration (days)	156.10	155.14	168.23	166.13	212.64	211.69
Average Task Duration (days)	3.12		3.35		4.25	
SOM/BSOM Underestimate	27.65 min		60.48 min		27.36 min	

Comparing the tour durations shows that the SepM-MPA minimising tour is 56.5 days faster than the B-MPA tour, representing an ~27 % tour duration saving by the SepM-MPA. The SepM-MPA balancing tour is 44.4 days faster, representing a 21 % tour duration saving even though the individual reconfiguration manoeuvres for the SepM-MPA balancing tour take longer to execute. These results highlight the scale of the benefits of using the SepM-MPA over the B-MPA in relation to observation scheduling performance. Even though the B-MPA tour is locally optimal for every observation, these compounded locally optimal observations do not result in globally optimal tour schedules.

11.3.2 Fuel Usage Comparison

A comparison between the SepM-MPA and B-MPA total fuel remaining is shown in Figure 11-2. The initial fuel total for the formation is 16.055 kg. The chart shows that after 50 manoeuvres the SepM-MPA minimising tour has used less fuel than the B-MPA. In total the SepM-MPA minimising tour used 44.9 g for fuel with a mean ~0.9 g per manoeuvre whilst the B-MPA used 74.1 g with a mean ~1.5 g per manoeuvre. This shows a 39.4 % decrease in fuel consumption for the SepM-MPA minimising tour relative to the B-MPA. Comparing the SepM-MPA balancing tour with the B-MPA tour shows that, as expected, the SepM-MPA balancing tour uses more fuel (~117 g with a mean 2.4 g per manoeuvre). This represents a 57.9 % increase in fuel consumption by the SepM-MPA balancing tour over the B-MPA tour. This is to be expected since the manoeuvre planning results for the PAM and TDM in Chapters 8 and 9 show that fuel balancing manoeuvres require more fuel to execute.

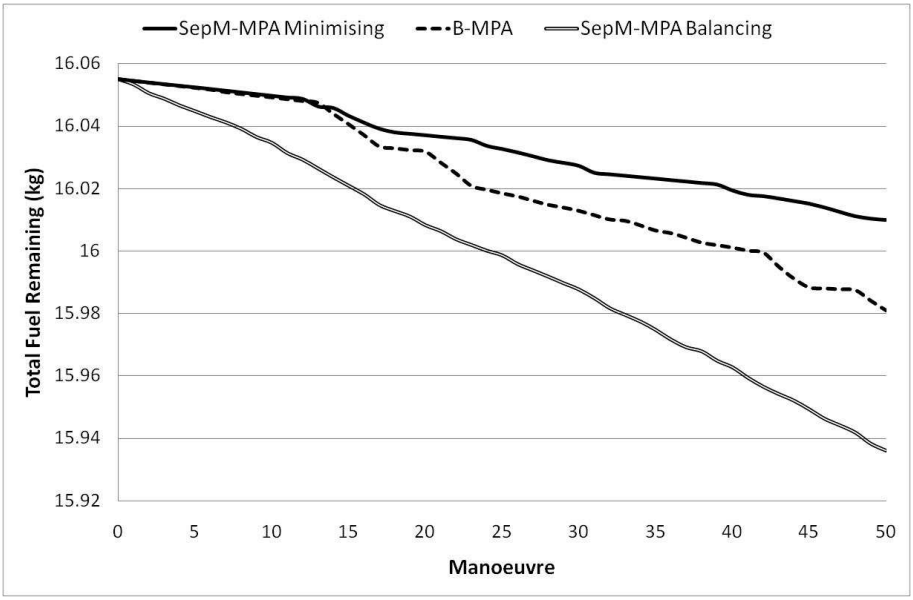

Figure 11-2 Fuel remaining comparison for the SepM-MPA and B-MPA. The chart shows than when set for fuel minimising the SepM-MPA uses least amount of fuel but when set for fuel balancing it uses the most. These results parallel the PAM tour analysis results in Chapter 8.

The scale of the decrease in fuel consumption for the SepM-MPA minimising tour is twofold, first through the PAM optimisation and then through the TDM. The shallower gradients of the data lines in Figure 11-2 represent the PAM/BPAM optimised manoeuvres that contained no manoeuvre violations and therefore did not require the execution of the TDM/BTDM. For these manoeuvres the average fuel consumed was 0.69 g for the SepM-MPA minimising manoeuvres and 0.83 g for the B-MPA. This is an ~16.8 % decrease in fuel consumption for the SepM-MPA minimising manoeuvres for PAM optimised manoeuvres. The steeper gradient lines represent manoeuvres that were optimised by the TDM/BTDM. For these manoeuvres the average fuel consumed was 2.1 g for the SepM-MPA minimising manoeuvres and 3.6 g for the B-MPA. This is an ~41.6 % decrease in fuel consumption for the SepM-MPA minimising manoeuvres for TDM optimised manoeuvres. These fuel

consumption differences between the SepM-MPA minimising manoeuvres and B-MPA are compounded by the number of times the TDM/BTDM is executed for each tour. For the SepM-MPA minimising tour the TDM is required 7 times whilst for the B-MPA tour the BTDM is required 13 times. This undoubtedly has an effect on the size of the final fuel consumption difference. However since the mean manoeuvre fuel consumption is less for the SepM-MPA minimising manoeuvres it is evident that the SepM-MPA is better than the B-MPA at fuel minimisation over a tour. A more concerted analysis of the frequency of proximity violating outputs from the PAM/BPAM is required to assess whether the PAM optimisation is less likely to produce a proximity violation than the BPAM.

This observation is mirrored when comparing the SepM-MPA balancing manoeuvres to the B-MPA. For the PAM optimised manoeuvres the SepM-MPA balancing tour average fuel consumption was 2.2 g whilst for the TDM optimised manoeuvres this was increased slightly to 2.4 g. This represents an increase of ~165 % for PAM optimised manoeuvres but a decrease of ~33 % for TDM optimised manoeuvres over the B-MPA manoeuvres. This indicates that while the PAM optimised SepM-MPA balancing manoeuvres require more fuel than the BPAM optimised B-MPA manoeuvres the TDM is much better at reducing the proximity avoidance manoeuvre fuel consumption than the BTDM. Fuel consumption for the SepM-MPA balancing tour is compounded by the number of times the TDM is executed. Of the 50 manoeuvres in the tour 44 required the TDM to plan proximity avoidance manoeuvres compared to the 13 required by the B-MPA. This high frequency results in a much larger fuel consumption of the SepM-MPA balancing tour even though its TDM manoeuvres are more fuel efficient than the BTDM manoeuvres.

11.3.3 Fuel Balancing and Performance Comparison

A comparison between the SepM-MPA and B-MPA sum of the fuel differences is shown in Figure 11-2. The initial sum of the fuel differences for the formation is 1.33 kg. All three of the tours shown show a reduction in the sum of the fuel differences in relation to the initial value. The B-MPA shows a reduction of 4.5 %, the SepM-MPA minimising tour shows a reduction of 1.8 % and the SepM-MPA balancing tour shows a reduction of 36.9 %. This clearly highlights the SepM-MPAs ability improve the fuel balancing across the fleet when required.

Using the performance metric defined in Chapter 8, the percentage increase in fuel balancing relative to the initial fuel balancing per percentage increase in fuel consumption relative to the $\mu_{PAM} = 0$, the tour performances are:

- B-MPA = $4.5\% / 65\% = 0.069$
- SepM-MPA balancing = $36.9\% / 161\% = 0.229$

This shows that even though the SepM-MPA balancing tour uses more fuel than the B-MPA, the extra amount of fuel is more efficiently used for fuel balancing than the B-MPA by ~230%

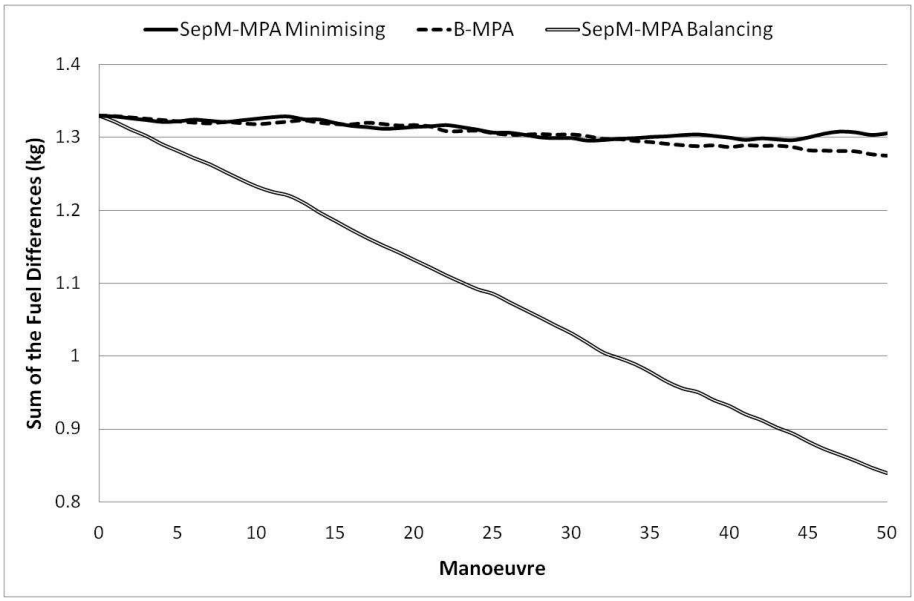

Figure 11-3 Fuel balancing comparison for the SepM-MPA and B-MPA. The chart shows that when set for fuel balancing the SepM-MPA is capable of reducing the sum of the fuel differences far better than when set for fuel minimising or than the B-MPA. These results parallel the PAM tour results in Chapter 8.

11.3.4 Calculation Time

Calculation time comparison between the different modules of the SepM-MPA and B-MPA are given in Table 11-5. The mean manoeuvre calculation time for each module is shown in the upper part Table 11-5. As expected the time taken to calculate each manoeuvre using the SepM-MPA is significantly greater than that for the B-MPA. The SepM-MPA tour calculation times are dominated by the user defined calculation time for the SOM of 7200 s (Table 11-2) whereas the BSOM is able to calculate its tour ~110 times faster.

Table 11-5 Calculation time comparison for the SepM-MPA and B-MPA

	Module	Time (sec)	Module	Time (sec)	Module	Time (sec)
Mean Manoeuvre Calculation Time	SOM	7297	SOM	7270	BSOM	66.21
	PAM	3.21	PAM	4.72	BPAM	0.04
	TDM	74.74	TDM	297	BTDM	0.30
	SepM-MPA Minimising	7375	SepM-MPA Balancing	7573	B-MPA	66.21
Total Calculation Time	SOM	364888	SOM	363534	BSOM	3310
	PAM	160	PAM	236	BPAM	2.03
	TDM	3737	TDM	14893	BTDM	15.13
	SepM-MPA Minimising	368762	SepM-MPA Balancing	378664	B-MPA	3327

Comparing the SepM-MPA minimising and balancing tours shows that the former calculated each manoeuvre on average 2.6 % faster than the latter. This is because both the PAM and TDM calculations required a longer calculation time when optimising for fuel

balancing. Coupling this with the increased frequency of TDM execution required by the SepM-MPA balancing tour shows the calculation time for the entire SepM-MPA minimising tour remains 2.6 % faster even though the entire calculation time for just the TDM is 75 % faster. This illustrates the dominance of the SOM calculation time on the total calculation time for the SepM-MPA tours.

Though the B-MPA is significantly faster than the SepM-MPA at calculating its manoeuvres the results need to be put into perspective. Firstly, the SepM-MPA is capable of generating tours up to 27 % more time efficient, with fuel savings of up to 40 % and fuel balancing gains of 34 %[12]. Secondly, the total calculation time for the SepM-MPA minimising tour is ~4.27 days. This calculation time only represents ~2.7 % of the total tour time of 156.1 days. This demonstrates that the gains achieved by manoeuvre planning using the SepM-MPA greatly outweigh the extra time required to calculate the manoeuvres.

The small percentage of total tour duration used by the SepM-MPA for this tour leaves plenty of spare capacity for the on-board processor to perform other tasks or devote more processor time to manoeuvre planning for greater optimisation results. For manoeuvre planning earlier in the mission however (i.e. TF0 or TF1) this spare capacity will be reduced. This is because the frequency of manoeuvres will be greater due to the larger availability of quicker tasks to schedule. For the SepM-MPA minimising tour introduced above the mean manoeuvre calculation time is 7375.73 sec and the observed maximum manoeuvre calculation time is 8859.02 sec. For the SepM-MPA balancing tour the mean manoeuvre calculation time is 7573.28 sec and the observed maximum manoeuvre calculation time is 7882.43 sec. These times are lower than the maximum allowed calculation time of 0.116 days (or 10022.4 sec)[13] from sub-section 7.4.4 and Table 7-2. Additionally, the SepM-MPA tours incorporate the majority of DARWIN reconfiguration manoeuvre types. These are shown in Table 11-6. Only two DARWIN manoeuvre types are not represented in the SepM-MPA minimising tour whilst only 4 are not represented in the SepM-MPA balancing tour. This data indicates that it is unlikely any manoeuvre calculation at any time in the mission will require a calculation time greater than 10^3 sec for a SOM calculation time of 7200 sec using similar hardware.

Table 11-6 Frequency of manoeuvre types in the SepM-MPA tours (L = linTTN, T = triTTN)

Baseline Change	None				Increasing				Decreasing			
Initial Configuration	linTTN		triTTN		linTTN		triTTN		linTTN		triTTN	
Final Configuration	L	T	L	T	L	T	L	T	L	T	L	T
No. Of Manoeuvres in Minimising Tour	24	0	12	0	1	3	1	2	1	2	1	3
No. Of Manoeuvres in Balancing Tour	24	0	11	0	2	4	0	2	0	2	2	3

[12] Relative to the end of tour sum of the fuel differences for the B-MPA

[13] This is the quickest observation task (10-13.2 μm spectroscopy on an M-type star (see Table 7-2)) and represents the shortest time the SepM-MPA would have available to generate a new manoeuvre plan.

11.4 Manoeuvre Planning Performance Metrics

The previous sub-sections have highlighted the differences observed between the SepM-MPA and the B-MPA. Over the 50 observation tour the SepM-MPA saved 56.5 days observation time and 29.2 grams of fuel. Extrapolating that out over the 5-year mission gives a saving of ~661 days observation time and ~341 g of fuel. As an indication of the monetary savings made adopting the SepM-MPA based on this extrapolation, comparisons can be made to existing missions.

The estimated annual operating cost of the James Webb Space Telescope (JWST) in 2015 is ~$120 M or ~$330 K per day (NASA FY2011 Budget Estimate – Astrophysics, 2010). Assuming a similar operating cost structure for DARWIN in 2030 means that using the SepM-MPA would save ~$217 M in operating costs by allowing the mission to complete almost 2 years early or allow a two year extension to the mission with no additional un-budgeted costs involved. DARWIN was due to be launched to L2 onboard two Soyuz 2 launchers with the Fregat third stage (Karlsson, et al., 2004). In 2000 this configuration had an estimated launch cost of $35 M (Soyuz/Fregat Website, 2000). The fuel savings demonstrated by the SepM-MPA for the DARWIN mission equate to only £8000 in launch costs.

It is clear from these cost comparisons that the SOM module of the SepM-MPA can realise substantial monetary savings for the planning of the operations whilst the PAM/TDM savings are negligible. Whilst the fuel savings are negligible in terms of cost they do equate to provide enough fuel for an extra ~140 reconfiguration manoeuvres (assuming manoeuvre planning using the SepM-MPA for fuel balancing and ignoring other fuel consuming operations). This extra fuel therefore could be used to help realise any mission extension that may occur due to cost savings made in spacecraft operations.

11.5 Fuel Balancing Performance Issues

It is clear from Chapters 8, 9 and sub-section 11.3 that employing aggressive fuel balancing can have significant fuel consumption issues. Chapter 8 shows that for both individual manoeuvres and over tours of manoeuvres a value of $\mu_{PAM}=0.6$ gives the optimal fuel balancing but with the penalty of an extra ~65% fuel consumption over the fuel minimised manoeuvre(s). Chapter 8 also shows however that fuel balancing performance (the amount of fuel balancing achieved as a measure of the extra fuel consumed) is optimal around $\mu_{PAM}=0.1$ for initially unbalanced fuel states and $\mu_{PAM}=0.03$ for initially balanced fuel fuel states. For the tour modelled in this chapter the initial fuel states were unbalanced but an aggressive (and hence fuel hungry) $\mu_{PAM}=0.6$ value was used for the fuel balancing over the entire tour.

For fuel balancing to be properly exploited within a tour requires an additional selection algorithm that, before each execution of the PAM, analyses the current fuel balancing state of the formation and select a μ_{PAM} value that represents the risk to the mission. If gross fuel imbalance exists (e.g. after recovery from safe mode) then perhaps more aggressive fuel balancing can be employed and if fuel imbalance is only moderate then either no, or minimal fuel balancing could be employed.

There is no doubt that fuel balancing within the manoeuvre planning environment is certainly desirable as demonstrated in Figure 8-15, Figure 8-16 and Figure 8-17. All three figures demonstrate that not employing fuel balancing, or employing too little fuel balancing, increases the fuel imbalance amongst the formation which could eventually lead to single/multiple spacecraft fuel starvation. Furthermore, non-planned manoeuvres like attitude and formation-keeping can also increase the fuel imbalance across the formation. At certain times in the mission it may just have to be necessary to sacrifice more fuel during reconfiguration manoeuvres to ensure the longevity of the mission.

11.6 TDM Execution Frequency Issues

Sub-section 11.3 highlighted the issue of the extra fuel consumed when executing TDM-planned manoeuvres over PAM-planned manoeuvres. Chapter 10 indicates that up to 50% extra fuel is required to execute TDM optimised manoeuvres over the avoidance criteria breaking PAM optimised manoeuvres. Examining the tour comparison data in sub-section 11.3 shows that the frequency in which the TDM is executed can have significant effects on the fuel consumption performance of a tour of manoeuvres and the comparison of said tours using different manoeuvre planning architectures. It is obvious therefore that the elimination of the requirement for the TDM (by ensuing the PAM always outputted positions that conformed to the avoidance criteria using straight-line trajectories) could increase the fuel consumption performance of the MPA. This is examined further in sub-section 12.3.2.

11.7 Chapter Summary

In this chapter the Separate Modular Manoeuvre Planning Architecture (SepM-MPA) was compared to a benchmark manoeuvre planning architecture (B-MPA). The chapter began with a summary of the operations of the SepM-MPA and introduced the B-MPA. The B-MPA has the same architecture as the SepM-MPA except the optimisation modules are the benchmark modules rather than the optimised modules. For this to be realised the stationkeeping module (SKM) has to be removed from the SepM-MPA and the effects of this on the performance of the SepM-MPA were discussed. The next section in the chapter discussed the set-up and initial conditions used for the MPA compassion analysis.

Three tours were compared; the B-MPA, the SepM-MPA with $\mu_{PAM} = 0$ to minimise fuel consumption and the SepM-MPA with $\mu_{PAM} = 0.6$ to optimise for fuel balancing. The results showed that over the 50 manoeuvre tour the SepM-MPA minimising tour was 27 % faster and used 40 % less fuel than the B-MPA tour. Additionally the SepM-MPA balancing tour was 21 % faster and although it used 58 % more fuel than the B-MPA it reduced the initial sum of the fuel differences by 37 % versus 5 % for the B-MPA. Analysing the calculation time for the tours showed that the B-MPA was ~110 times faster than the SepM-MPA. However the individual manoeuvre calculation time for the SepM-MPA was never recorded above 10^3 sec (the defined maximum allowable limit). The comparison performance increase was shown to generate significant monetary benefits in operations cost savings when projected over the entire mission duration. The comparison however also highlighted

deficiencies in the SepM-MPA with regards to managing the fuel balancing performance of the tour and the frequency of execution of the fuel-hungry Trajectory Design Module

The analysis and comparison of the SepM-MPA with the B-MPA has shown that in all aspects the SepM-MPA is a much more capable manoeuvre planner than the B-MPA. Additionally this capability can theoretically be realised for a manoeuvre planning system autonomously operating on-board one of the formation spacecraft by 2030. Though the Station-keeping Module (SKM) was not implemented in this analysis its influence on the manoeuvre planning results would be minimal since it would likely only be called once or twice for the durations of tour analysed. The SKMs most likely effect would be its calculation time and how it influences the total calculation time of the MPA in relation to the defined allowable maximum limit of 10^3 sec per manoeuvre.

.

12. SUMMARY, CONCLUSIONS AND FUTURE WORK

As the final chapter in this thesis this chapter will present a summary of the research project. The thesis is summarised chapter by chapter and the main findings from each presented in a concise form. This is followed by the summary of the optimisation module future work and a discussion on the future work relating to the Manoeuvre Planning Architecture (MPA). Finally the main conclusions of the research are drawn and referenced to the initial aims and objectives of the project

12.1 Thesis Summary and Findings

The summary of the thesis is split into two parts to aid comprehension. Part one summarises Chapters 1-5 covering the background to the research project, its aims, objectives and the analysis and selection of formation flying concepts that affected the design of the Manoeuvre Planning Architecture (MPA). Part two of the summary covers the remaining Chapters 6-11 summarising the design of the MPA, design and analysis of the optimisation modules and comparison of the SepM-MPA and its benchmark (B-MPA).

12.1.1 Thesis Summary Part One

Chapter 1 provided an introduction to the background and motivations for undertaking this research project and detailed the aims and objectives for the research. The motivations were driven by project's stakeholders; Cranfield University, the European Space Agency (ESA), EADS Astrium and the author himself and were as follows:

- Research into formation flying reconfiguration manoeuvre planning
- Development of a software simulator for formation flying manoeuvre planning
- Support of the DARWIN mission
- Support of extra-solar planet research

These motivations helped shape the direction of the research and the individuals representing the organisations, provided valuable feedback throughout the entire process. The discussion on the research motivations was followed by a brief introduction to ESA's DARWIN mission (the reference mission for the research) before a description of the problem statement was presented. The problem was characterised in three parts, spacecraft formation flying manoeuvres, formation flying at Lagrange points and formation flying autonomy. These three research topics became the drivers for the aims and objectives for the project. The chapter ended with a discussion on the research contributions of the project. These included the synergy of the many different types of formation flying manoeuvre planning found in the literature to one unified MPA, the inclusion of a DARWIN-based observation scheduling algorithm and a proposed method to incorporate station-keeping manoeuvre planning as part of the unified manoeuvre planning process.

In Chapter 2 the motivation concepts of extra-solar planet research and ESA's DARWIN mission were introduced and discussed in-depth. The first part of the chapter dealt with the search for extra-solar planets and began with a discussion on the motivational aspects for this research. The motivation is two-fold with one side of the research community (namely the astrophysicists and geophysicists) interested in understanding the processes of planetary formation and evolution and the other side (namely the astrobiologists and astrochemists) interested in life detection studies. These motivations have led to a slew of innovative planet detection techniques that have thus far discovered over 300 exoplanets in the last 15 years. A review of these detection techniques was followed by an introduction to the present and future space missions that are designed to further this research field. The second part of Chapter 2 provided an in-depth review of the DARWIN mission predominantly taken from the 2004 ESA Mission Assessment Study (Karlsson, et al., 2004). The review discussed the nature of the mission to L_2 and its relation to spacecraft formation flying. Also given was detailed information relating to the mission constraints, science observations requirements, guidance, navigation and control modes and the DARWIN spacecraft themselves that affect the design and development of any formation flying manoeuvre planning architecture.

The third chapter of this thesis discussed spacecraft formation flying and began with the definition of two distinct distributed spacecraft system families:

- Spacecraft flying in formation – where the formation is a result of orbit/trajectory design and each spacecraft is manoeuvred independently from the ground
- Spacecraft formation flying –where the formation is a result of autonomous active control and spacecraft are manoeuvred relative to another tracked spacecraft in the formation

This distinction was necessary to refine the field of research of this thesis and numerous examples were given in the text of both types of distributed spacecraft system. Chapter 3 continued with an introduction to planning, co-ordination and control concepts unique to spacecraft formation flying missions. These included methods for organising the manoeuvre planning capability of the formation, methods to co-ordinate formation flying manoeuvres and autonomous control methods required to execute the manoeuvres. The review of these concepts and methods was important for the design of any formation flying manoeuvre planning architecture.

Chapter 4 gave a brief introduction into the concept of mathematical optimisation and its necessity for spacecraft formation flying manoeuvre planning. The generic optimisation problem was defined using a cost function and forward-referenced with its usage later in the thesis. The solution space of the problem was also described to provide a background for the optimisation limitations found in Chapter 9. This was followed by a review of the various methods used to solve optimisation problems, their advantages and disadvantages and their usage later in the thesis.

In Chapter 5 the various concepts introduced in the preceding chapters were assessed and selections made that affected the design of the MPA. The first assessment related to the GNC modes introduced in Chapter 2 and the following conclusions were drawn:

- The MPA was to be designed for the reconfiguration mode of DARWIN only.
- The execution of station-keeping manoeuvres is restricted to the reconfiguration mode only and therefore requires planning with the formation reconfiguration manoeuvres.

The second assessment related to the formation flying concepts introduced in Chapter 5 and the following conclusions were drawn:

- The MPA was to be designed assuming the top-down autonomous manoeuvre planning organisation structure with the Beam Combiner Spacecraft (BCS) as the S_{FP} and the Telescope Spacecraft (TS) as S_{NP}s.
- The leader-follower control co-ordination was selected with the BCS as the leader spacecraft and the TSs as the followers.
- The control strategy chosen was of a dual nature that was separate between the BCS and the TS:
    - BCS follows a planned open-loop thrust schedule for both reconfiguration and station-keeping manoeuvres.
    - TS follow planned trajectories relative to the BCS for reconfiguration manoeuvres and formation-keep with the BCS for station-keeping manoeuvres using a closed-loop controller.

This chapter also discussed the concept of manoeuvre error mitigation for real-time proximity avoidance monitoring and mitigation. A method for manoeuvre error mitigation was introduced involving the RF metrology system for position determination and a set of avoidance plans depending on the nature of the manoeuvre malfunction. The final assessment of Chapter 5 related to the choice of optimisation technique and methods employed to solve the formation flying manoeuvre planning problem. These choices were driven by stakeholder motivations and the selections made previously in the chapter. The main points were:

- The MPA was to be designed as a series of multi-variable optimisation modules each optimising different aspects of the formation flying manoeuvre planning problem.
- Implementation of the MPA was to be in the MATLAB® software environment allowing the use of embedded optimisation algorithms where suitable and requiring the coding of bespoke algorithms where necessary.

Chapter 5 represented the end of the first part of the thesis.

12.1.2 Thesis Summary Part Two

This sub-section summarises the second part of the thesis, Chapters 6-0, and highlights the findings from each chapter.

12.1.2.1 Manoeuvre Planning Architecture

The design of the MPA was detailed in Chapter 6. The manoeuvre planning systems model for DARWIN was conceptualised and presented along with a trade-off analysis that lead to the emergence of four optimisation modules designed to solve and optimise all aspects of the formation flying manoeuvre planning problem. The optimisation modules were incorporated into the Separate Modular Manoeuvre Planning Architecture (SepM-MPA), the designed MPA for this research project. The SepM-MPA treats the modules individually and sequentially and outputs the required thrust profile or relative trajectories for the BCS and TSs respectively. The limitations of the SepM-MPA were detailed and explained as part of the trade-off between calculation time efficiency and optimisation accuracy. A response to the SepM-MPA limitations was also given but will be discussed in greater detail in sub-section 12.3.2. The chapter ended with an explanation of the Manoeuvre Information Dissemination (MID), the execution of the planned manoeuvre and a discussion of the operational hardware required to execute the architecture.

12.1.2.2 Science Operations Module

In Chapter 7 the Science Operations Module (SOM) was designed, implemented, analysed and compared to a benchmark algorithm (the BSOM). The constraints placed on the science operations planning section of the DARWIN mission were discussed in-depth followed by a review of the previous contributions found in the literature relating to operations scheduling. The problem definition and the SOM/BSOMs approach in finding a solution were detailed and the background to the analysis was given. The analysis followed the performance of the BSOM, the tuning of the variables for the SOM, the performance of the SOM and the comparison between the SOM and BSOM. The main findings were:

- With a maximum calculation time of 10^3 sec the SOM was always able to find better performing tours than the BSOM given the appropriate standard deviation for the mission time. This performance increase yielded results of up to 3 hrs/day observation time saved using the SOM over BSOM observation schedules.
- The standard deviation had to be in the range $0.01 < \sigma < 0.1$ and decrease throughout the mission duration to ensure maximum performance from the SOM was maintained.

The chapter concluded with recommendations for future work but these will be summarised in more detail in sub-section 12.3.1.

12.1.2.3 Position Assignment Module

The Position Assignment Module (PAM) design was covered in Chapter 8 of the thesis. The chapter began with a discussion of the manoeuvre planning constraints posed for formation flying missions before a review of the previous contributions found in the literature relating to manoeuvre planning position assignment was presented. The next part of the chapter detailed the world model within which the algorithm was designed before in-depth

design of the PAM was given. The performance of the PAM was analysed and its output compared to the benchmark version (the BPAM). The main findings were:

- For single manoeuvres the PAM was always able to achieve better spacecraft positions than the BPAM.
    - For fuel minimising manoeuvres the PAM saved on average 6.5 % fuel over the BPAM over DARWIN-like manoeuvres analysed. This however required an increased manoeuvre duration of on average 10 % over the BPAM.
    - For fuel balancing manoeuvres the PAM on average was able to increase fuel balancing by 0.35 % over the initial fuel imbalance for the example given. This however required an increase fuel consumption of on average 70 % with variable manoeuvre duration changes.
    - Optimal fuel minimisation was achieved when $\mu_{PAM} = 0$ (since the problem is reduced to a single objective single variable problem) and optimal fuel balancing achieved when $0.5 \leq \mu_{PAM} \leq 0.7$.

The output of the PAM was also analysed for its use over consecutive manoeuvres in a tour of manoeuvres. A fuel balancing performance metric was defined where the performance measured the amount of fuel balancing achieved for the amount of extra fuel being consumed. The main findings were:

- For a tour of manoeuvres the single manoeuvre findings were repeated i.e. optimal fuel minimisation is achieved when $\mu_{PAM} = 0$ and optimal fuel balancing achieved when $0.5 \leq \mu_{PAM} \leq 0.7$.
- For initially fuel balanced formations $\mu_{PAM} = 0.03$ gave the highest performance
- For initially fuel unbalanced formations $\mu_{PAM} = 0.1 - 0.2$ gave the highest performance.

The μ_{PAM} values were highlighted as they were deemed important for the design of a 'μ_{PAM} - selection' algorithm that would be required for the SepM-MPA to autonomously decide the most appropriate value of μ_{PAM} for the current fuel balancing state. Also affecting the design of a 'μ_{PAM} - selection' algorithm were the reported manoeuvre planning anomalies found during the analysis. Whilst these were shown to be extremely rare, it was deemed important to at least plan against their recurrence. Chapter 8 ended with a review of some future work concepts that will be expanded further in sub-section 12.3.1.

12.1.2.4 Trajectory Design Module

Chapter 9 described the design, implementation and analysis of the Trajectory Design Module (TDM). A brief description of the manoeuvre planning constraints tackled by the TDM was followed by a review of the previous contributions found in the literature relating to spacecraft avoidance manoeuvre planning. The world model within which the TDM was

designed was described along with an in-depth explanation of the TDM design and implementation. Analysis of the TDM revealed that the initially chosen optimisation algorithm, the PatternSearch (PS) algorithm from MATLABs Genetic Algorithm and Direct Search Toolbox (GADS), was unable to consistently optimise the cost function. A solution was found using the Genetic Algorithm (GA) from GADS which gave a 35 % decrease in cost but incurred a 470 % calculation time penalty over the PS algorithm. For the analysis of DARWIN-like manoeuvres the main findings showed:

- 10 % of $\mu_{PAM} = 0$ PAM optimised manoeuvres required execution of the TDM.
- 44 % of $\mu_{PAM} = 0.6$ PAM optimised manoeuvres required execution of the TDM
- Up to 50 % more fuel was required for $\mu_{TDM} = 0$ TDM over the PAM manoeuvres.
- Up to 300 % more fuel was required for $\mu_{TDM} = 0.6$ TDM over the PAM manoeuvres.
- For $\mu_{TDM} = 0.6$ the increase in fuel balancing was only ~1.5 % and due to the increased fuel costs made $\mu_{TDM} \neq 0$ an unattractive manoeuvre planning option.

Chapter 9 ended with a review of some future work concepts that will be expanded further in sub-section 12.3.1.

12.1.2.5 Station-keeping Module

The Station-keeping Module (SKM) was introduced in Chapter 10. The chapter began with a description of the libration points and the related mathematical preliminaries. It was then shown that periodic motion about the collinear libration points can be computed from the equations describing the simplified dynamic environment and used as reference trajectories for spacecraft. A review of different reference trajectory generation techniques was followed by a review of the past and proposed station-keeping techniques. The next two sub-sections of the chapter detailed the selected numerical reference trajectory generation technique and the selected station-keeping method, the Target Point Strategy (TPS). This was followed by a description of how the SKM was envisaged to operate. Unfortunately, due to a lack of time available, the author was unable to satisfactorily code and validate the SKM. Therefore no results or analysis could be reported. The final part of the chapter detailed the future work that would be necessary to complete the SKM.

12.1.2.6 Manoeuvre Planning Architecture Comparison

The final chapter in part two of this thesis described the comparison between the SepM-MPA and a benchmark manoeuvre planning architecture (dubbed the B-PAM). Following a re-cap of the SepM-MPA operation and an introduction to the optimisation modules used in the B-MPA the parameters were defined for the tours to be generated in the comparison. The B-MPA was compared against the SepM-MPA for both fuel minimising and fuel balancing.

The tour comparison yielded the following results over a 50 manoeuvre tour:

- The SepM-MPA minimising tour finished 56.5 days before the B-MPA tour giving a time saving of ~27 %.
- The SepM-MPA balancing tour finished 44 days before the B-MPA tour giving a time saving of ~21 %.
- Both the SOM and BSOM underestimated the manoeuvre duration calculated by the PAM/BPAM by ~27 min for the SepM-MPA minimising and B-MPA tours and ~60 min for the SepM-MPA balancing tour.

The fuel usage comparison results over a 50 manoeuvre tour showed:

- The SepM-MPA minimising tour used 40 % less fuel than the B-MPA.
- The SepM-MPA balancing tour used 58 % more fuel than the B-MPA but was able to reduce the sum of the fuel differences by 37 % compared to the B-MPA reduction of 5 %.
- The fuel balancing efficiency of the SepM-MPA balancing tour was ~ 230 % better than the B-MPA tour.

Comparing the calculation times for the tours showed:

- The B-MPA calculated the tour ~ 110 times faster than the SepM-MPA.
- Calculation durations for the SepM-MPA only represented a fraction of the actual tour duration.
- Average and maximum manoeuvre calculation duration for the SepM-MPA never exceeded the maximum allowed calculation time of 10^3 sec.

The comparison of the SepM-MPA with the B-MPA showed that in all aspects (apart from the calculation time) the SepM-MPA was a superior manoeuvre planning algorithm.

12.2 Conclusions

The stated aim of the research project from Chapter 1 was:

To design optimal manoeuvre planning algorithms for use with separated spacecraft interferometry missions at L_2 *(but specifically in support of the DARWIN mission) to enable the safe execution of formation flying reconfiguration manoeuvres. Planning these manoeuvres should allow the maximum science return to be realised for the mission through a combination of schedule optimisation, manoeuvre optimisation and optimal fuel management across all spacecraft in the formation. The planning algorithms should also be of sufficient simplicity to enable their inclusion as part of an on-board autonomous guidance, navigation and control sub-system.*

The Separate Modular Manoeuvre Planning Architecture (SepM-MPA) developed in this thesis represents one way of satisfying this aim.

The Architecture designed in this thesis represents a unique way to tackle the formation flying manoeuvre planning problem by viewing it at a mission systems level. The inclusions of operations scheduling and spacecraft station-keeping in this discussion are unique to this thesis. The analysis of the Architecture shows that operations scheduling has a far greater affect on mission cost than the manoeuvre fuel/time management solutions offered in other analyses. To realise the full benefits however of better operations scheduling (i.e. extended mission duration) requires the careful management of fuel consumption and fuel balancing offered by the Architecture.

The requirements analysis for the Science Operation Module demonstrates the complexity of operations scheduling for a formation flying based exoplanet science mission. Many interlinked temporal constraints create a complex optimisation environment requiring careful navigation to find the optimal observation schedules. The Science Operations Module represent a unique approach for optimising schedules within this complex environment and although is CPU-limited is shown to perform exceptionally within the calculation time constraints present within the mission.

Whilst the implementation of the Station-keeping Module was not fully realised in this thesis, the inclusion of station-keeping within a formation flying manoeuvre planning strategy has never been formally presented in other works. The design of the SKM, both as an individual optimisation module, and its inclusion in the SepM-MPA, represents a robust, unique and plausible approach to incorporate station-keeping manoeuvre optimisation into a formation flying manoeuvre planning system.

The SepM-MPA and its associated optimisation modules were all designed with computational efficiency as a core driver. Some of the modules are quicker than others but the Architecture (minus the SKM) has been shown to generate manoeuvre plans that within the defined minimum calculation duration meet the planning requirements. When iterated to plan a 50 observation tour in the middle stages of the DARWIN mission the SepM-MPA demonstrated an enhanced scheduling efficiency with time savings of 27 %, an enhanced fuel consumption efficiency with fuel savings of 40 % and an enhanced fuel balancing efficiency with fuel balancing 7 times better than a benchmark manoeuvre planning algorithm.

Whilst the development and analysis of the SepM-MPA has raised a number of further questions, the requirement for more analysis and a desire for improved optimisation modules, the author believes that the Architecture represents a simple, effective and cost effective way to plan and optimise formation flying reconfiguration manoeuvres.

12.3 Future Work

In this sub-section the future work opportunities that have arisen from this project will be discussed. In the individual optimisation module chapters (7-1) the final sub-sections discussed future work concepts to make the individual modules perform better. These will be summarised first followed by a discussion of the future work for the Manoeuvre Planning Architecture (MPA) as a whole.

12.3.1 Optimisation Module Future Work

This sub-section will summarise the future work concepts found in the individual optimisation module chapters relating to improving the module. Any improvements previously discussed relating to the MPA will be covered in the next sub-section.

12.3.1.1 Science Operations Module

Work required to improve the SOM was given in sub-section 7.9 and concentrated on three main themes; incorporating the remaining scheduling constraints, performing a comprehensive data analysis and optimising the coding and execution of the algorithm.

Some of the scheduling constraints introduced in Chapter 7 were not implemented into the version of the SOM presented in this thesis. The inclusion of these non-implemented constraints (and the more accurate modelling of the implemented constraints) will enable the SOM to generate more accurate tours. Obviously the BSOM would also have to include these additional constraints and so any comparison between the two would likely yield the same results. Additionally, Chapter 11, revealed that the SOM underestimated the formation manoeuvre time by ~27 min and ~60 min for fuel minimising and fuel balancing manoeuvres respectively. This underestimation should be accounted for in future versions of the SOM to enhance the accuracy of the tours generated.

Many of the results presented in Chapter 7 appear dependent on 'fixed' parameters that were not varied for the analysis. The tour starting star, length of the target tour duration and the planet/star probability will all affect the number of iterations required and standard deviation required to optimise the tour performance. Whilst restricting the number of iterations to as many as can be achieved within 7200 sec is a fair standard to use comprehensive analysis on the other 'fixed' parameters is required to see if any further patterns (like decreasing σ over the mission duration) emerge and quantitative results can be found.

The final suggested improvement for the SOM was related to coding optimisation. The performance of the SOM was related to the number of tour iterations that can be generated within the calculation time given. A SOM algorithm version that can complete one tour calculation quicker using the same hardware should perform better than the SOM presented in this thesis. An example of utilising memory resources to achieve more iterations was given however there are likely many other ways that the SOM coding could be optimised.

12.3.1.2 Position Assignment Module

Work required to improve the PAM was given in sub-section 8.4 with the non-architecture suggestion relating to the use of ΔV optimisation. It was found in Chapter 8 that the large single spacecraft fuel consumption anomaly was due partly to fuel and mass imbalance between the spacecraft in the formation and partly due to the lack on constraints within the optimisation problem definition. Removing the fuel mass terms from the PAM cost function and replacing them with the corresponding ΔV terms prevented the PAM from reproducing the anomaly using the same initial conditions. Using ΔV optimisation instead of fuel optimisation for the PAM removes the negative influence that spacecraft mass can have

on the PAM optimisation. Whilst this has only been demonstrated once a full repetition of the PAM analysis using the ΔV optimisation is required to assess and compare its performance with the fuel optimisation version of the PAM presented in this thesis.

12.3.1.3 Trajectory Design Module

Work required to improve the TDM was given in sub-section 9.4 and centred on improving the TDM solution space and alternative trajectory generation methods. The solution space for the TDM cost function was shown to be difficult for optimisation algorithms to navigate and resulted in the adoption of a more complex genetic algorithm (GA) to find suitable solutions. By removing the infinity costs for proximity violations within the TDM cost function a simpler cost function can be used. This should allow simpler optimisation algorithms to navigate the space and speed up the GA. The proximity violations can then be incorporated as constraints within the optimisation problem definition. The entire TDM analysis should be repeated with this new cost function to assess its performance.

The second future work concept for the TDM was the adoption of alternative trajectory generation methods. A number of examples were given including the alteration of the nominal thrust, the use of alternative perpendicular thrust profiles and the removal of the constraint that forces the avoidance thrust to be perpendicular to the nominal thrust. These are interesting concepts to examine to see whether their performance justifies their additional complexity.

12.3.1.4 Station-keeping Module

Work required to complete the Station-keeping Module was given in sub-section 10.5. A number of steps are presented detailing how the reference trajectory generated is to be incorporated along with the Target Point station-keeping strategy and an L_2 dynamic model to make the SKM. After complete integration and validation of the SKM a number of suggested analyses are presented that will help optimise the SKM execution conditions that l affect the frequency of SKM execution.

12.3.2 Manoeuvre Planning Architecture Future Work

The optimisation module further work summary presented in sub-section 12.3.1 omits those further work suggestions from Chapters 7-1 that would require a change to the operation of the SepM-MPA. From the analysis in Chapter 11 the greatest effect in the fuel optimality of the tour is the frequency of calls to the TDM. From the analysis in Chapter 6 the design decision to separate position assignment and proximity violation control into two separate optimisation modules was based upon the desire to minimise the average calculation time of the optimisation. The unintended consequence of this is, by allowing the PAM to generate positions that break the proximity rules, a more fuel costly solution is adopted in the TDM. The way to avoid this is to integrate the TDM into the PAM and this can be realised in two ways.

12.3.2.1 Simplified Modular Manoeuvre Planning Architecture

One approach to embed the TDM with the PAM involves modifying the cost function of the PAM to assign a high cost to any position combination that involves a collision or thruster plume proximity violation. This entirely negates the requirement for the full TDM (as only the proximity-check is required) and so is called the simplified modular manoeuvre planning architecture (SimM-MPA) as in Figure 12-1. The main benefit of the SimM-MPA is that the modified PAM will always output positions that conform to the proximity rules and with no TDM in the loop could produce better fuel optimised manoeuvres. In addition, the average calculation time may drop since the TDM never gets called. However this must be tempered by the increased complexity in the PAM cost function that requires a proximity check for every position combination it examines. The main disadvantage of the SimM-MPA for fuel management is its rejection of potential optimal positions (with proximity violations) in favour of positions without proximity violations. In some cases a TDM modified manoeuvre from PAM optimised positions as calculated by the SepM-MPA may still be more optimal than a manoeuvre optimised within the SimM-MPA.

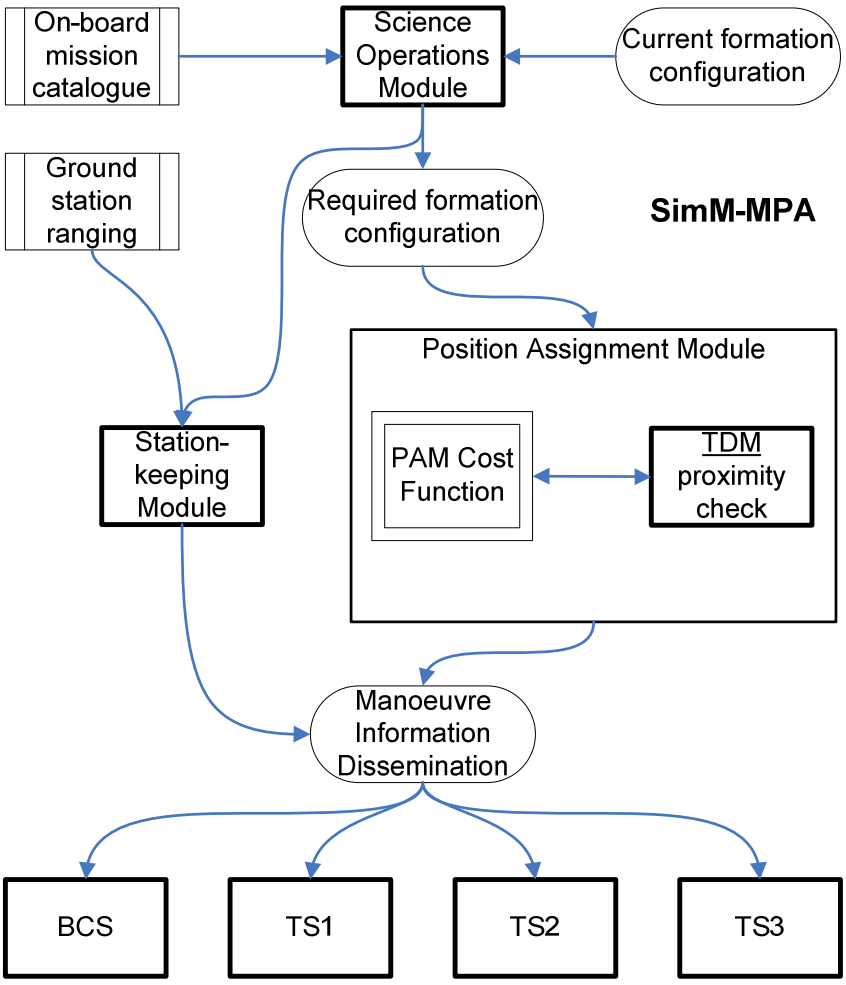

Figure 12-1 Simplified Modular-Manoeuvre Planning Architecture

Some basic analysis can be performed to examine the potential of the SimM-MPA. The TDM execution time from Chapter 11 revealed that on average it takes the TDM ~75 sec for fuel minimising manoeuvres and ~300 sec for fuel balancing manoeuvres to find safe, fuel optimised trajectories using the reference hardware. Since this only represents 1 % and 3 % of the mean total calculation time for SepM-MPA minimising and SepM-MPA balancing manoeuvres respectively it is clear that the TDM does not have a large influence on the SepM-MPA calculation time. For this reason the SimM-MPA would seem redundant. For the fuel management issue however, relating the reduced optimality of TDM manoeuvres on PAM optimised manoeuvres, the SimM-MPA would be required. Comparing the PAM results for the SimM-MPA with the TDM results from the SepM-MPA using an identical tour would show whether the fuel management of the SimM-MPA was hindered by the restricted PAM.

12.3.2.2 Integrated Position Assignment Module Manoeuvre Planning Architecture

Another way to embed the TDM into the PAM involves the integration of the TDM into the PAM cost function to produce the integrated PAM manoeuvre planning architecture (IPAM-MPA) as in Figure 12-2. For the IPAM-MPA whenever the PAM cost function analyses a position combination a proximity check is carried out. However, instead of rejecting combinations that cause proximity violations (as in the SimM-MPA), the cost function calls the TDM to optimise safe trajectories for that position combination. This solves the issue posed by the SimM-MPA rejecting optimal solutions but adds a further level in computational complexity over both the SimM-MPA and the SepM-MPA.

A rough calculation can be made to access the increased calculation time required for the IPAM-MPA to generate one manoeuvre plan. It is assumed that the PAM optimisation algorithm requires 100 iterations to find the best solution for any initial conditions. For a fuel balancing manoeuvre (from Chapter 9) 44 % of those iterations would require the execution of the TDM. From Chapter 11 the mean calculation time for the TDM on fuel balanced spacecraft positions is ~300 sec. So the TDM component of the calculation takes ~13200 sec (or 3.67 hrs). Even before the SOM and PAM calculation time components are added it is clear that the IPAM-MPA would not be able to complete the manoeuvre plan within the 10^3 sec limit imposed in Chapter 7.

This calculation time however is not prohibitively long for manoeuvre planning in the later stages of the mission when the observation phases of the mission, and hence the available calculation times, are much longer. The improvements in planning performance of the IPAM-MPA over the SepM-MPA later in the mission may justify a switch to the IPAM-MPA at a specified mission epoch. For this reason the development of the IPAM-MPA is recommended as potential future work for this research project.

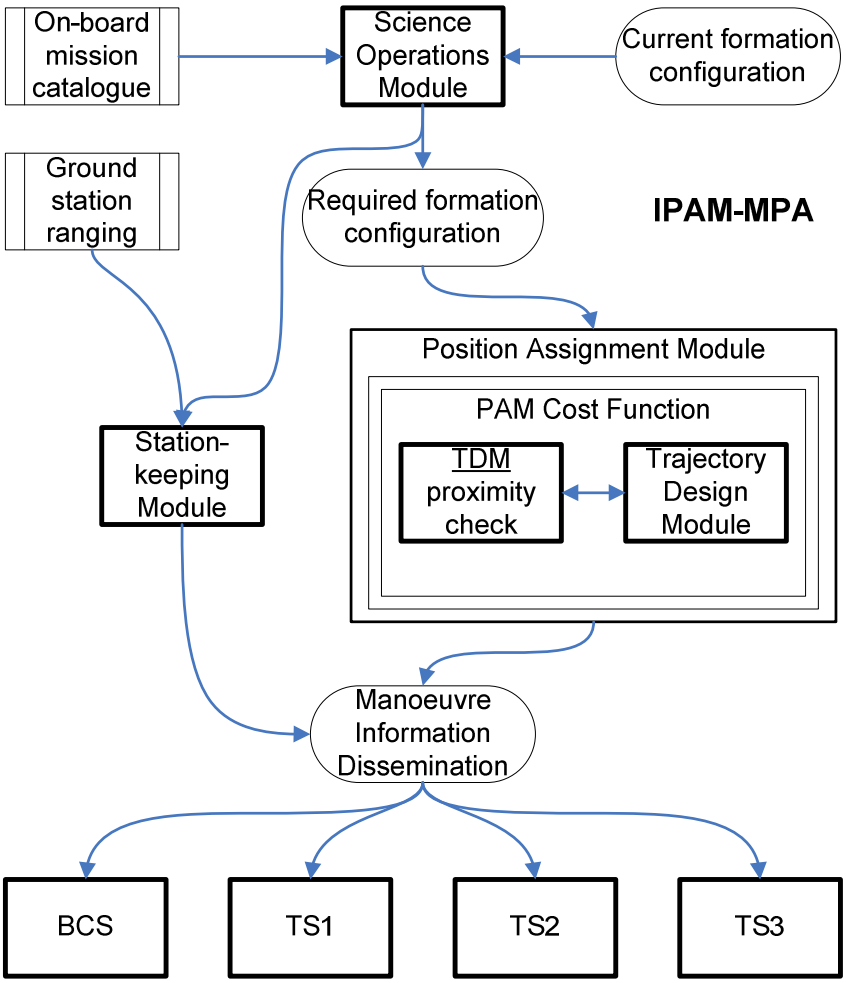

Figure 12-2 Integrated Position Assignment Module-Manoeuvre Planning Architecture

12.3.2.3 Global Optimisation Manoeuvre Planning Architecture

The SepM-MPA has been designed so that it is capable of running autonomously on-board one of the DARWIN spacecraft circa 2030. The simplifications made in developing the SepM-MPA in Chapter 6 have realised this at the potential cost however in the time/fuel optimality of the mission. Even removing one of the simplest trades, as in the IPAM-MPA, pushes the calculation cost of the architecture above the limitations on the defined on-board system. The requirements for on-board planning autonomy have been driven by reduced operating costs for the mission. A ground-based planning architecture, employing much more powerful computer hardware, could run a full global manoeuvre planning architecture, potentially out-performing the optimisations of the SepM-MPA, and thus giving further cost savings for the mission. If the extra costs saved covered the costs of running the ground-based hardware then the deployment of the autonomous SepM-MPA would not be necessary. A full cost analysis of the SepM-MPA therefore would required a more in-depth study into ground-based global manoeuvre planning architecture optimisation and as such is a final future work recommendation from this research project.

13. REFERENCES

Abramson, M.A., "MATLAB Genetic Algorithm and Direct Search Toolbox", The Mathworks Inc, Natwick, Massachusetts, 2006.

Açikmese, A.B., Mettler, E., Breckenridge, W.G., Macenka, S.A. and Tubbs, E.F. "L_2 Earth Atmosphere Observatory: Formation Guidance, Metrology and Control Synthesis", *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Providence, Rhode Island, 16-19 August, 2004. AIAA 2004-5212.

Aeroflex Gaisler (2008), "LEON 3 Multiprocessing CPU Core", http://www.gaisler.com/doc/leon3_product_sheet.pdf, Last accessed: 20 March 2010.

Aeroflex Gaisler (2010), "LEON 4 32-bit processor core", http://www.gaisler.com/doc/LEON4_32-bit_processor_core.pdf, Last accessed: 20 March 2010.

Ankersen, F. "Coordinate Systems and Rotation Definitions for Darwin and Demonstration Mission", *ESA Internal Document*, December 2003.

BAE Systems, "RAD750™", http://www.aero.org/conferences/mrqw/2002-papers/A_Burcin.pdf, Last accessed: 20 March 2010.

Bailey, C.A., McLain, T.W. and Beard, R.W., "Fuel-Saving Strategies for Dual Spacecraft Interferometry Missions", *The Journal of the Astronautical Sciences*, Vol. 49, No. 3, pp. 469-488, July-September 2001.

Bansal, N., et al., "Approximation Algorithms for Deadline TSP and Vehicle Routing with Time Windows", *STOC 2004*, Chicago, Illinois, 13-15 June, 2004.

Beard, R.W. and Hadaegh, F.Y. "Fuel Optimized Rotation for Satellite Formations in Free Space", *Proceedings of the American Control Conference*, San Diego, California, June 1999.

Beard, R.W., Lawton, J. and Hadaegh, F.Y. "A Coordination Architecture for Spacecraft Formation Control", *IEEE Transactions on Control Systems Technology*, Vol. 9, No. 6, 2001.

Beard, R.W., McLain, T.W. and Hadaegh, F.Y., "Fuel Equalized Retargeting for Separated Spacecraft Interferometry", *American Control Conference*, Philadelphia, Pennsylvania, June 1998.

Beard, R.W., McLain, T.W. and Hadaegh, F.Y. "Fuel Optimization for Constrained Rotation of Spacecraft Formations", *Journal of Guidance, Control and Dynamics*, Vol. 23, No. 2, pp. 339-346, 2000.

Bennet, D.J. and McInnes, C.R. “Pattern Transition in Spacecraft Formation Flying via the Artificial Potential Field Method and Bifurcation Theory”, *3rd International Symposium on Formation Flying, Missions and Technologies*, ESA-ESTEC, Noordwijk, The Netherlands, 23-25 April, 2008.

Blackwood, G.H., et al., “StarLight Mission: A Formation Flying Stellar Interferometer”, In: *Proceedings of the SPIE*, Volume 4852 - “Interferometry in Space”, pp. 463-480, 2003.

Borde, J., Teston, F., Santandrea, S. and Boulande, S. “Feasibility of the PROBA-3 Formation Flying Demonstration Mission as a Pair of Microsats in GTO”, *55th International Astronautical Congress*, Vancouver, Canada, October 4-8, 2004. IAC-04-U.1.04.

Borucki, W.J. and Summers, A.L., “The Photometric Method of Detecting Other Planetary Systems”, *Icarus*, Vol. 58, Issue 1, pp. 121-134, 1984.

Brown, O. and Eremenko, P. “The Value Proposition for Fractionated Spacecraft Architectures”, *Space 2006*, San Jose, California, 19-21 September, 2006.

Brucker, P., Hilbig, T. and Hurink, J., “A Branch and Bound Algorithm for a Single-Machine Scheduling Problem with Positive and Negative Time Lags”, *Discrete Applied Mathematics*, Vol. 94, pp. 77-99, 1999.

Burgon, R. “The Search and Detection of Extra-Solar Planets Using Gravitational Microlensing”, BSc. Thesis, School of Physics and Astronomy, St Andrews University, Fife, UK, 1997.

Caitlin, K.A. and McLaughlin, C.A. “Earth-Moon Triangular Libration Point Spacecraft Formations”, *Journal of Spacecraft and Rockets*, Vol. 44, No. 3, pp. 660-670, 2007.

Campbell, M.E. “Planning Algorithm for Multiple Satellite Clusters”, *Journal of Guidance, Control and Dynamics*, Vol. 26, No. 5, pp 770-780, 2003.

Campbell, M. and Schetter, T. “Comparison of Multiple Agent-Based Organisations for Satellite Constellations”, *Journal of Spacecraft and Rockets*, Vol. 39, No. 2, pp.274-283, 2002.

Campbell, M.E., Zanon, D. and Kulkarni, J. “Cluster Planning and Control for Spacecraft Formations”, *14th AAS/AIAA Spaceflight Mechanics Conference*, Maui, Hawaii, 8-12 February, 2004. AAS 04-254.

Carlson, B.A., Pernicka, H.J. and Balakrishnan, S.N. “Spacecraft Formation Flight about Libration Points”, *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Providence, Rhode Island, 16-19 August 2004. AIAA 2004-4737.

Chabot, T. and Udrea, B. “XEUS Mission Guidance Navigation and Control”, *AIAA Guidance, Navigation and Control Conference and Exhibit*, Keystone, Colorado, 21-24 August, 2006. AIAA-2006-6587.

Charbonneau, D., Brown, T., Noyes, R.W. and Gilliland, R., “Detection of an Extrasolar Planet Atmosphere”, *Astrophysical Journal*, Vol. 568, pp. 377-384, 2002.

Chein, S. et al, “Automated Planning and Scheduling for Goal-Based Autonomous Spacecraft”, *Intelligent Systems and Their Applications*, IEEE, Vol. 13, Issue 5, pp. 50-55, September-October 1998.

Coello Coello, C.A., Van Veldhuizen, D.A. and Lamont, G.B., “Evolutionary Algorithms for Solving Multi-objective Problems”, Kluwer Academic/Plenum Publishers, New York, 2002.

Coleman, T.F. and Zang, Y. “MATLAB Optimization Toolbox”, The Mathworks Inc, Natwick, Massachusetts, 2005.

Colomb, F.R. and Varotto, C.F. “SAC-C and the AM Constellation: Three Years of Achievements”, *International Conference on Recent Advances in Space Technologies (RAST)*, Istanbul, Turkey, 20-22 November, 2003.

Conkey, D.E., Dell, G.T., Good, S.M. and Bristow, J. “EO-1 Formation Flying Using AutoCon™”, *IEEE Aerospace Conference*, Vol. 7, pp. 55-61, 2000.

COROT website “COROT Scientific Results”, http://smsc.cnes.fr/COROT/GP_res_scie.htm. Last accessed: 3 June 2009.

Cover, T.M. and Thomas, J.A., “Elements of Information Theory”, Wiley, New York, 1991, Chapter 2, pp. 13.

CPLEX, 2008. Internet resource - http://www.ilog.com/products/cplex/. Last accessed: 26 June 2008.

Cranfield University Space Research Centre website, “Space Research Centre – Research”, http://www.cranfield.ac.uk/soe/departments/aerospaceengineering/space/page14639.jsp. Last accessed: 22 May 2009.

Cranfield University Statement of Work, “Sponsored PhD on Satellite Formation Flying”, CP/PhD/02.2005, Issue 3, Cranfield University, Wharley End, Bedfordshire, UK, October 2006.

Crauwels, H.A.J., Potts, C.N. and Van Wassenhove, L.N., "Local Search Heuristics for the Single Machine Total Weighted Tardiness Scheduling Problem", *INFORMS Journal of Computing*, Vol. 10, No. 3, Summer 1998.

D'Arcio, L. "*Darwin Payload Definition Document*", ESA Technical Note (SCI-A/2005/301/DARWIN/DMS/LdA), Issue 1, Revision 2, August, 2005.

da Silva Curiel, A., Boland, L., Cooksley, A., Bekhti, M., Stephens, P., Sun, W. and Sweeting, M. "First results from the disaster monitoring constellation (DMC)", *Acta Astronautica*, Vol. 56, pp. 261-271, 2005.

Davidson, M., Maleville, L., Boulade, S., McQuade, F. and Ankersen, F. "GNC System for the Deployment and the Fine Control of the Darwin Free-Flying Interferometer", *7th International Conference on Dynamics and Control of Systems and Structures in Space*, Greenwich, UK, 16-20 July, 2006.

de Boer, P-T., Kroese, D.P., Mannor, S. and Rubinstein, R.Y., "A Tutorial on the Cross-Entropy Method", *Annals of Operations Research*, Vol. 134, pp 19-67, 2005.

Deming, D., Seager, S., Richardson, J. and Harrington, J., "Infrared radiation from an extrasolar planet", *Nature*, Vol. 434, pp. 740 – 743, 2005.

den Hartog, R., "Darwin All Sky Target Stars Catalogue, Issue 1, November 2005", private communication, May, 2006.

de Selding, P.B., "Five-Satellite RapidEye Constellation Launched Successfully", *Space News*, Vol. 19, No. 34, pp. 6, 1 Sept, 2008

DLR website, "The habitable zone for stars of different sizes", http://www.dlr.de/en/DesktopDefault.aspx/tabid-5170/8702_read-15322/gallery-1/gallery_read-Image.1.7761/. Last accessed: 2 June 2009.

Dorigo, M. and Di Caro, G., "The Ant Colony Optimization Meta-Heuristic", In: *New Ideas in Optimization*, pp. 11-32. McGraw Hill, London, UK, 1999.

Dorigo, M., Di Caro, G. and Gambardella, L.M., "Ant Algorithms for Discrete Optimization", *Artificial Life*, Vol. 5, No. 2, pp. 137-172, Spring 1999.

Dunham, D.W. and Roberts, C.E. "Stationkeeping Techniques for Libration-Point Satellites", *The Journal of the Astronautical Sciences*, Vol. 49, No. 1, pp. 127-144, January-March, 2001.

ESA ATV Information Kit, "The Automated Transfer Vehicle – ATV Overview", http://esamultimedia.esa.int/docs/ATV/infokit/english/01_ATVOverview.pdf, 2008. Last accessed: 23 October 2008.

ESA Cluster website, “Cluster”. http://sci.esa.int/science-e/www/area/index.cfm?fareaid=8. Last accessed: 28 October 2008.

ESA Cosmic Vision website, “Cosmic Vision 2015-2025”. http://sci.esa.int/science-e/www/area/index.cfm?fareaid=100. Last accessed: 4 June 2009.

ESA Cross-Scale website, “Cross-Scale”. http://sci.esa.int/science-e/www/area/index.cfm?fareaid=109. Last accessed: 30 October 2008.

ESA GAIA website, “GAIA Extra-solar Planets”, http://sci.esa.int/science-e/www/object/index.cfm?fobjectid=40577. Last accessed: 4 June 2009.

ESA Herschel website, “Herschel”, http://sci.esa.int/science-e/www/area/index.cfm?fareaid=16. Last accessed: 6 February 2009.

ESA JWST website, “JWST”, http://sci.esa.int/science-e/www/area/index.cfm?fareaid=29. Last accessed: 6 February 2009.

ESA News website 2006, “ESA launches new initiative to foster research”, http://www.esa.int/esaCP/SEM8889L6VE_index_0.html. Last accessed: 20 May 2009.

ESA News website 2007, “COROT discovers its first exoplanet and catches scientists by surprise”, http://www.esa.int/esaCP/SEMCKNU681F_index_0.html. Last accessed: 3 June 2009.

ESA Planck website, “Planck”, http://sci.esa.int/science-e/www/area/index.cfm?fareaid=17. Last accessed: 6 February 2009.

ESA PROBA-3 website, “Proba-3 Science Payload”, http://www.esa.int/techresources/ESTEC-Article-fullArticle_par-28_1220880350211.html, 2008. Last accessed: 23 October 2008.

ESA Space Science website (2009a), “How to Find an Extrasolar Planet”, http://www.esa.int/esaSC/SEMYZF9YFDD_index_1.html. Last accessed: 3 June 2009.

ESA Space Science website (2009b), “COROT discovers smallest exoplanet yet, with a surface to walk on”, http://www.esa.int/esaSC/SEM7G6XPXPF_index_0.html. Last accessed: 6 July 2009.

ESO (European Southern Observatory) website, “The Radial Velocity Method”, http://www.eso.org/gallery/v/ESOPIA/illustrations/phot-22e-07.tif.html. Last accessed: 3 June 2009.

Farquhar, R.W., et al. “Trajectories and Orbital Manoeuvres for the First Libration-Point Satellite”, *Journal of Guidance and Control*, Vol. 3, No. 6, pp. 549-554, November-December, 1980.

Ferri, P. and Sorensen, E., “Automated Mission Operations for Rosetta”, IN: *Proceedings of the 5th International Symposium on Space Mission Operations and Ground Data Systems: SpaceOps 98*, Tokyo, Japan, June 1-5, 1998.

Folta, D., Hartman, K., Howell, K. And Marchand, B. “Formation Control of the MAXIM L_2 Libration Point Mission”, *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Providence, Rhode Island, 16-19 August, 2004.

Folta, D and Hawkins, A. “Results of NASA’s First Autonomous Formation Flying Experiment: Earth Observing-1 (EO-1)”, *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Monterey, Canada, 5-8 August, 2002.

Franca, P.M., Mendes, A. and Moscato, P., “A Memetic Algorithm for the Total Tardiness Single Machine Scheduling Problem”, European *Journal of Operations Research*, Vol. 132, pp. 224-242, 2001.

Fridlund, C.V.M. “DARWIN – The Infrared Space Interferometer”, IN: *Proceedings of the Conference ‘Darwin and Astronomy – The Infrared Space Interferometer’*, Stockholm, Sweden, 17-19 November, 1999. ESA SP-451.

Galileo Website, “European Commission Transport”, http://ec.europa.eu/transport/galileo/index_en.htm, 2009. Last accessed: 13 May 2009.

Gaisler, J. “LEON 1 Processor – First Evaluation Results”, http://klabs.org/DEI/Processor/sparc/Papers/gaisler.pdf, Last accessed: 20 March 2010.

Globalstar Website, “Globalstar – Our Technology”, http://www.globalstareurope.com/en/content.php?cid=601, 2009. Last accessed: 13 May 2009.

GLONASS Website, “Russian Space Agency – Information-Analytical Centre”, http://www.glonass-ianc.rsa.ru/pls/htmldb/f?p=202:1:16827457493506720876, 2009. Last accessed: 13 May 2009.

Gomez, G. et al. “Station-Keeping Strategies for Translunar Libration Point Orbits”, *Advances in the Astronautical Sciences*, Vol. 99, pp. 949-964, 1998.

GPS Website, “Global Positioning System – Serving the World”, www.gps.gov, 2009. Last accessed: 13 May 2009.

Grady, J. “IXO Project Overview”, IXO *Facility Science Team Meeting*, Goddard Spaceflight Centre, Maryland, MD, 20-22nd August, 2008. https://conxproj.gsfc.nasa.gov/resources/meetings/DOCS/fstaug08/04_Grady_Project%20Update_FST_8-20-08.pdf. Last accessed: 24 October 2008.

Grun, E., Zook, H.A., Fechtig, H. and Giese, R.H. “Collisional Balance of the Meteoritic Complex”, *Icarus*, Vol. 62, pp 244-272, 1985.

Gurfil, P., Idan, M. and Kasdin, N.J. “Neurocontrol of Spacecraft Formation Flying in the Elliptic Restricted Three-Body Problem”, *AIAA Guidance, Navigation and Control Conference and Exhibit*, Monterey, California, 5-8 August, 2002. AIAA 2002-4962.

Gurfil, P., Idan, M. and Kasdin, N.J. “Adaptive Neural Control of Deep-Space Formation Flying”, *Journal of Guidance Control and Dynamics*, Vol. 25, No. 3, 2003.

Gurfil, P. and Kasdin, N.J. “Stability and Control of Spacecraft Formation Flying in Trajectories of the Restricted Three-Body Problem” *Acta Astronautica*, Vol. 54, pp. 433-453, 2004.

Hamilton, N.H., Folta, D. and Carpenter, R. “Formation Flying Satellite Control around the L_2 Sun-Earth Libration Point”, *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Monterey, California, 5-8 August, 2002.

Hammer, J., Piper, G., Thorp, O. and Watkins, J. “Investigating Virtual Structure Based Control Strategies for Spacecraft Formation Flying Manoeuvres”, *AIAA Guidance, Navigation and Control Conference and Exhibit*, Providence, Rhode Island, 16-19 August, 2004.

Hernádvölgyi, I.T., “Solving the Sequential Ordering Problem with Automatically Generated Lower Bounds”, IN: *Proceedings of Operations Research 2003*, Heidelberg, Germany, pp. 355–362, 2003.

Howell, K.C. and Marchand, B.G. “Design and Control of Formations near the Libration Points of the Sun-Earth/Moon Ephemeris System”, *GSFC Flight Mechanics Symposium*, Greenbelt, MD, October, 2003.

Howell, K.C. and Pernika, H.J. “Numerical Determination of Lissajous Trajectories in the Restricted Three-Body Problem”, *Celestial Mechanics*, Vol. 41, pp. 107-124, 1988.

Howell, K.C. and Pernicka, H.J. “Stationkeeping Method for Libration Point Trajectories”, *Journal of Guidance and Control*, Vol. 16, No. 1, pp. 151-159, 1993.

Hughes, S.P., "Formation Design and Sensitivity Analysis for the Magnetospheric Multiscale Mission (MMS)", *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Honolulu, Hawaii, 18 - 21 August 2008.

Hurkens, C. A. and Woeginger, G. J., "On the Nearest Neighbor Rule for the Traveling Salesman Problem", *Operations Research Letters*, Vol. 32, No. 1, pp. 1-4, 2004.

Iridium Website, "Iridium – How it works", http://www.iridium.com/about/howitworks.php, 2009. Last accessed: 13 May 2009.

Izzo, D.R. "Formation Flying Linear Modelling", *5th Dynamics and Control of Systems and Structures in Space*, Cambridge, UK, 14-18 July, 2002.

JACARA (The Joint Australian Centre for Astrophysical Research in Antarctica) website, "Science Project for Pilot - Planetary Microlensing in the Inner Galaxy", http://www.phys.unsw.edu.au/jacara/pilotscience.php. Last accessed: 3 June 2009.

JPL Case Studies website. "Earth Atmosphere Observatory at L_2". http://start1.jpl.nasa.gov/caseStudies/eao-L2.cfm. Last accessed: 26 October 2008.

JPL TPF-I website. "Exoplanet Interferometry Technology: Technology Plan", 21 July, 2008. http://planetquest.jpl.nasa.gov/TPF-I/tpf_currentStatus.cfm. Last accessed: 26 October 2008.

Junge, O et al. "Identification of Halo Orbits for Energy Efficient Formation Flying", *IN: Proceedings of the International Symposium Formation Flying*, Toulouse, 2002.

Kaltenegger, L., Karlsson, A., Fridlund, M. and Absil, O. "Overview of the Darwin Mission" *Towards Other Earths: DARWIN/TPF and the Search for Extrasolar Terrestrial Planets*, Heidelberg, Germany, 22-25 April, 2003. ESA SP-539.

Karlsson, A., Kaltenegger, L., den Hartog, R., d'Arcio, L., Kilter, M., Ch. Erd and Ankersen, F. "Darwin, TTN+ Mission Design Assessment", *ESA Internal Document (SCI-A/2004/187/Darwin/DMS)*, Issue 1, Revision 2, November 2004.

Kennedy, J. and Eberhart, R. C., "Particle Swarm Optimization", In: *Proceedings of the IEEE International Conference on Neural Networks*, Piscataway, New Jersey, pp. 1942-1948, 1995.

Kepler website, "Capabilities of Various Planet Detection Methods", http://kepler.nasa.gov/sci/capabilities.html. Last accessed: 4 June 2009.

Kidder, S.Q., Kankiewicz, J.A. and Vonder Haar, T.H. "The A-Train: How Formation Flying is Transforming Remote Sensing", *Joint 2007 EUMETSAT Meteorological Satellite*

Conference and the 15th Satellite Meteorology & Oceanography Conference of the American Meteorological Society, Amsterdam, The Netherlands, 24-28 September, 2007.

Kim, Y., Mesbahi, M. and Hadaegh, F.Y., "*Duel-Spacecraft Formation Flying in Deep-Space: Collision-Free Reconfigurations*", Journal of Guidance, Control and Dynamics, Vol. 26, No. 2, March-April, 2003.

King, L.B., Parker, G.G., Deshmukh, S., and Chong, J-H., "A Study of Inter-spacecraft Coulomb Forces and Implications for Formation Flying, *Journal of Propulsion and Power*, Vol. 19, No. 3, pp. 497-505, 2003.

Kulkarni, J.E., Campbell, M.E. and Dullerud, G.E. "Stabilization of Spacecraft Flight in Halo Orbits: An H∞ Approach", *IEEE Transactions on Control Systems Technology*, Vol. 14, No. 3, May 2006.

Kunieda, H., Parmar, A. and White, N. "Announcing: The International X-ray Observatory (IXO)". http://ixo.gsfc.nasa.gov/news/2008/ixo_announcement.html. Last accessed: 24 October 2008.

Lagadec, K., Lebas, J. and Ankersen, F. "Precision Formation Flying for the Darwin Interferometer", *5th International ESA Conference on Spacecraft Guidance, Navigation and Control Systems*, Frascati, Italy, 22-25 October, 2002. ESA-SP-516.

Lawson, P.R. et al. "Terrestrial Planet Finder Interferometer (TPF-I) Whitepaper for the AAAC Exoplanet Task Force", 2 April, 2007. http://planetquest.jpl.nasa.gov/TPF-I/TPFIwhitepaper.pdf. Last accessed: 26 October 2008.

Lawton, R.L. and Beard, R.W. "Synchronized Multiple Spacecraft Rotations", *Automatica*, Vol. 38, pp. 1359-1364, 2002.

Liu, G.P., Yang, J.B. and Whidbourne, J.F. "Multiobjective Optimisation and Control", Research Studies Press Ltd, Baldock, Herts, UK, Chapter 1, pp. 3-4, 2003.

Mailhe, L.M. and Guzman, J.J. "Initialization and Resizing of Formation Flying using Global and Local Optimization Methods", *IEEE Aerospace Conference Proceedings*, Big Sky, Montana, 6-13 March, 2004.

Mao, S. and Paczynski, B., "Gravitational Microlensing by Double Stars and Planetary Systems", *Astrophysical Journal*, Vol. 374, Part - 2 Letters, pp. L37-L40, 1991.

Marchand, B.G. and Howell, K.C. "Control Strategies for Formation Flight in the Vicinity of the Libration Points" *Journal of Guidance, Control and Dynamics*, Vol. 28, No. 6, 2005.

Mayor, M and Queloz, D., "A Jupiter-mass companion to a solar type star", Nature, Vol. 378, pp. 355-359, 1995.

McQuade, F., Ward, R. and McInnes, C.R. “The Autonomous Configuration of Satellite Formations Using Generic Potential Functions”, *3rd International Workshop on Satellite Constellations and Formation Flying*, Pisa, Italy, February, 2003.

Merz, P. and Freisleben, B., “Genetic Local Search for the TSP: New Results”, *IEEE International Conference on Evolutionary Computation*, Indianapolis, Indiana, 13-16 April 1997.

Mettle, E., Acikmese, A.B., Breckenridge, W.G., Macenka, S.A. and Tubbs, E.F. “Earth Atmosphere Observatory Formation at L_2”, *Space 2004 Conference and Exhibit*, San Diego, California, 28-30 September, 2004.

MOST Science website, “MOST - Canada’s First Space Telescope”, http://www.astro.ubc.ca/MOST/science.html. Last accessed: 3 June 2009.

Mueller, J.B. “A Multiple-Team Organization for Decentralized Guidance and Control of Formation Flying Spacecraft”, *AIAA 1st Intelligent Systems Technical Conference*, Chicago, Illinois, Sep. 20-22, 2004. AIAA-2004-6249.

Muscettola, N., Pell, B., Hansson, O. and Mohan, S., “Automating Mission Scheduling for Space-Based Observatories”, *Robotic Telescopes: Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy*, G.W. Henry and J.A. Eaton (eds). Astronomical Society of the Pacific, Provo, UT.

NASA FY2011 Budget Estimate – Astrophysics, http://www.nasa.gov/pdf/428151main_Astrophysics.pdf, Last Accessed: 5 April 2010.

NASA, “Vision for Space Exploration”, NP-2004-01-334-HQ, NASA Headquarters, Wahsington, DC, February, 2004.

NASA Report, “Overview of the DART Mishap Investigation Results – For Public Release”, http://mynasa.nasa.gov/pdf/148072main_DART_mishap_overview.pdf, 2006. Last accessed: 23 October 2008.

Nelder, J.A. and Mead, R., “*A Simplex Method for Function Minimisation*”, Computer Journal, Vol. 7, No. 4, pp 308-313, 1965.

Papadimitriou, C.H. and Yannakakis, M. “Shortest Paths without a map”, Theoretical Computer Science, Vol. 84, Issue 1, pp. 127-150, 1991.

Penin, L.F., Araujo, J. and Avila, N. “Design and Evaluation of Optimal Reconfiguration Manoeuvres for Separated Spacecraft Interferometry”, *Acta Astronautica*, Vol. 57, pp. 330-340, 2005.

Pernicka, H.J., Carlson, B.A. and Balakrishnan, S.N. “Spacecraft Formation Flight About Libration Points Using Impulsive Manoeuvring”, *Journal of Guidance, Control and Dynamics*, Vol.29, No. 5, pp. 1122-1130, 2006.

Persson, S., Jacobsson, B. and Gill, E. “PRISMA – Demonstration Mission for Advanced Rendezvous and Formation Flying Technologies and Sensors”, *56th International Astronautical Conference*, Fukuoka, Japan, 17-21 October, 2005. IAC-05-B5.6.B.07.

Petre, R., White, N.E., Tananbaum, H., Hornschemeier, A., Bookbinder, J., Garcia, M., Grady, Jean. and Kilbourne, C. “The Status of the Constellation-X Mission”, IN: *Proceedings of the SPIE*, Vol. 6686, pp. 66860B-66860B-9, 2007.

Pettazzi, L., Izzo, D. and Theil, S., “Swarm Navigation and Reconfiguration using Electrostatic Forces”, *7th International Conference on Dynamics and Control of Systems and Structures in Space*, Greenwich, UK, pp. 257-267, 16-20 July, 2006.

Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P., “Numerical Recipes in C: The Art of Scientific Computing”, 2nd Edition, Cambridge University Press, 1992.
Rahmani, A., Jalali, M.A. and Pourtakdoust, S.H. “Optimal Approach to Halo Orbit Control”, *AIAA Guidance, Navigation and Control Conference and Exhibit*, Austin, Texas, 11-14 August, 2003.

Richards, A., Schouwenaars, T., How, J.P. and Feron, E. “Spacecraft Trajectory Planning with Avoidance Constraints Using Mixed-Integer Linear Programming”, *Journal of Guidance, Control and Dynamics*, Vol. 25, No. 4, pp 755-764, July-August, 2002.

Richards, A., How, J., Schouwenaars, T. and Feron, E., “Plume Avoidance Manoeuvre Planning Using Mixed Integer Linear Programming”, *AIAA Guidance, Navigation and Control Conference and Exhibit*, Montreal, Canada, 6-9 August, 2001.

Richardson, D.L. “Analytical Construction of Periodic Orbits about the Collinear Points”, *Celestial Mechanics*, Vol. 22, pp. 241-253, 1980a.

Richardson, D.L. “Halo-Orbit Formulation for the ISEE-3 Mission”, *Journal of Guidance and Control*, Vol. 3, No. 6, pp. 543-548, November-December, 1980b.

Richardson, D.L. and Cary, N.D. “A Uniformly Valid Solution for Motion About the Interior Librations Point of the Perturbed Elliptic-Restricted Problem”, *AAS/AIAA Astrodynamics Specialist Conference*, Nassau, Bahamas, July 28-30, 1975. AAS 75-021.

Roberts, J. A. “Satellite formation flying for an interferometry mission”, *PhD Thesis*, Cranfield University Press, Bedfordshire, UK, 2005.

Roberts, J.A. "Development of a Relative Motion Model for Satellite Formation Flying Around L_2", *Cranfield CERES Staff Publications*, Cranfield University, Bedfordshire, UK, 2004. https://dspace.lib.cranfield.ac.uk/handle/1826/788. Last accessed: 22 May 2009.

Roberts, J.A and Roberts, P.C.E. "The Development of High Fidelity Linearized J2 Models for Satellite Formation Flying Control", *Advances in the Astronautical Sciences*, Vol. 119, Part 1, pp. 913-934, 2004.

Roberts, P.C.E., Bowling, T.S. and Hobbs, S.E. "MUSTANG: A Technology Demonstrator for Formation Flying and Distributed Systems Technologies in Space", *5th Dynamics and Control of Systems and Structures in Space*, Cambridge, UK, 14-18 July, 2002.

Rohrbaugh, D. and Schiff, C. "Stationkeeping Approach for the Microwave Anisotropy Probe (MAP)", *AIAA/ASS Astrodynamics Specialist Conference and Exhibit*, Monterey, California, 5-8 August, 2002.

Rui, X., Ping-yuan, C. and Xiao-fei, X., "Realization of Multi-Agent Planning System for Autonomous Spacecraft", *Advances in Engineering Software*, Vol. 36, pp. 266-272, 2005.

Rui, X., Ping-yuan, C., Xiao-fei, X. and Hu-tao, C., "Design for Autonomous Mission Planning System", *Aircraft Engineering and Aerospace Technology*, Vol. 75, No. 4, pp. 365-371, 2003.

Ruilier, C., Sghedoni, M. and Krawczyk, R. "DARWIN System Assessment Study Summary Report", *Alcatel Internal Document* (DW_SAS-ASP-TN-457), Issue 1, February, 2007.

Scharf, D.P., Hadaegh, F.Y. and Ploen, S.R. "A Survey of Spacecraft Formation Flying Guidance and Control (Part I): Guidance", IN: *Proceedings of the American Control Conference*, Denver, Colorado, +-June 4-6, pp. 1733-1739, 2003.

Schneider, J. "The Extrasolar Planets Encyclopaedia", http://exoplanet.eu/catalog.php. Last accessed: 3 June 2009.

Science@NASA website, "In Search of ET's Breath", http://science.nasa.gov/headlines/y2002/10jan_exo-atmospheres.htm. Last accessed: 3 June 2009.

Seereeram, L., Ravichandran, M., Smith, R. and Beard, R. "Multi-spacecraft Trajectory Optimization, Using Genetic Algorithm Techniques", IN: *Proceedings of the IEEE Aerospace Conference*, Vol.7, pp. 99-108, 2000.

Segerman, A.M. and Zedd, M.F., "Preliminary Planar Formation-Flight Dynamics near Sun-Earth L_2 Point", *AAS/AIAA Space Flight Mechanics Meeting*, Ponce, PR, 2003. AAS-03-113.

Sengupta, P. and Vadali, S.R. “Modelling and Control of Interferometric Formation in the Vicinity of the Collinear Libration Points”, *Space Flight Mechanics*, Vol. 1, 2005, AAS 05-161.

Sigurdsson, S., Richer, H.B., Hansen, B.M., Stairs, I.H. and Thorsett, S.E., “A Young White Dwarf Companion to Pulsar B1620-26: Evidence for Early Planet Formation”, Science, Vol. 301, No. 5630, pp. 192.196, 2003.

Simo, C. et al. “On the Optimal Station Keeping Control of Halo Orbits”, *Acta Astronautica*, Vol. 15, No. 6/7, pp. 391-397, 1987.

Singh, G. and Hadaegh, F.Y., “Collision Avoidance Guidance for Formation-Flying Applications”, *AIAA Guidance, Navigation and Control Conference and Exhibit*, Montreal, Canada, 6-9 August, 2001.

Smith, B., Rajan, K. and Muscettola, N., "Knowledge Acquisition for the Onboard Planner of an Autonomous Spacecraft", *European Knowledge Acquisition Workshop*, Spain 1997.

Smith, D.R., Ambrosi, R.M., Holland, A.D., Hutchinson, I.B. and Wells, A. “The Prompt Particle Background and Micrometeriod Environment at L_2 and its Implications for Eddington”, *2nd Eddington Workshop*, Palermo, 9-11 April, 2003. ESA SP-538.

Smith, L., “First pictures taken of planet outside the solar system: Fomalhaut b”, *The Times Online*, http://www.timesonline.co.uk/tol/news/uk/science/article5149705.ece. Last accessed: 3 June 2009.

Smith, R.S. and Hadaegh, F.Y. “Control of Deep-Space Formation-Flying Spacecraft: Relative Sensing and Switched Information”, *Journal of Guidance, Control and Dynamics*, Vol. 28, No. 1, 2005.

Soyuz/Fregat Website, 2000. http://www.orbireport.com/Launchers/Soyuz_U-Fregat/. Lastr accessed: 5 April 2010.

Stella, L., Santoro, G.E. and Tosatti, E., “Optimization by Quantum Annealing: Lessons from Simple Cases”, *Physical Review B*, Vol. 72, No. 1, 014303, 2005.

Stron, R. and Price, K., “Differential Evolution – A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces”, *Journal of Global Optimization*, Vol. 11, No. 4, pp. 341-359, December, 1997.

Sultan, C., Seereeram, S. and Mehra, R.K., “Minimization and Equalization of Energy for Formation Flying Reconfiguration”, *Proceedings of the IEEE International Conference on Robotics and Automaton*, New Orleans, LA, April, 2004a.

Sultan, C., Seereeram, S. and Mehra, R.K., "Matrix Inequalities and Energy Optimal Reconfiguration for Formation Flying Spacecraft", *AIAA Guidance, Navigation and Control Conference and Exhibit*, Providence, Rhode Island, 16-19th August, 2004b.

Sultan, C., Seereeram, S. and Mehra, R.K., "Energy Optimal Reconfiguration for Large Scale Formation Flying", *Proceedings of the American Control Conference*, Boston, MA, June 30 – July 2, 2004c.

Sultan, C., Seereeram, S. and Mehra, R.K., "Energy Optimal Multi-Spacecraft Relative Reconfiguration of Deep Space Formation Flying", *43rd IEEE Conference on Decision and Control*, Atlantis, Paradise Island, Bahamas, 14-17 December, 2004d.

The Mathworks Inc., "MATLAB® – The Language of Technical Computing", Natwick, Massachusetts, Version R2006a, 2006.

Thurman, R. and Wolfolk, P.A. "The Geometry of Halo Orbits in the Circular Restricted Three-body Problem", *Technical Report GCG95*, Geometry Centre, University of Minnesota, 1996.

Tillerson, M., Breger, L. and How, J.P. "Distributed Coordination and Control of Formation Flying Spacecraft", *IEEE American Control Conference*, Denver, Colorado, June 4-6, 2003.

Tillerson, M., Inalhan G. and How, J.P. "Co-ordination and Control of Distributed Spacecraft Systems Using Convex Optimization Techniques", *International Journal of Robust and Nonlinear Control*, Vol. 12, pp 207-242, 2002.

TPF-C website, "Terrestrial Planet Finder Coronagraph", http://planetquest.jpl.nasa.gov/TPF-C/tpf-C_index.cfm. Last accessed: 4 June 2009.

TPF-I website, "Terrestrial Planet Finder Interferometer", http://planetquest.jpl.nasa.gov/TPF-I/tpf-I_index.cfm. Last accessed: 4 June 2009.

Vaddi, S.S., Alfriend, K.T., Vadali, S.R. and Sengupta, P. "Formation Establishment and Reconfiguration Using Impulsive Control", *Journal of Guidance, Control and Dynamics*, Vol. 28, No.2, pp. 262-268, 2005.

VanDyke, M.C. and Hall, C.D. "Decentralized Coordinated Attitude Control within a Formation of Spacecraft", *Journal of Guidance, Control and Dynamics*, Vol. 29, No.5, 2006.

Vane, D. "The CloudSat Mission and the A-Train: A Revolutionary Approach to Observing Earth's Atmosphere", IN: *Proceedings of the IEEE Aerospace Conference*, pp. 1-5, 2008.

Vignal, P. and Pernicka, H. "Low-Thrust Spacecraft Formation Keeping", *Journal of Spacecraft and Rockets*, Vol. 43, No. 2, pp. 466-475, 2006.

Wallner, O. “DARWIN System Assessment Study Summary Report”, *Astrium Internal Document* (DARWIN-SR), Issue 1, Revision 0, December, 2006.

Wang, P.K.C. and Hadaegh, F.Y. “Coordination and Control of Multiple Microspacecraft Moving in Formation”, The Journal of the Astronautical Sciences, Vol. 44, No. 3, pp. 315-355, 1996.

Wei, R and Beard, R.W. “Decentralized Scheme for Spacecraft Formation Flying via the Virtual Structure Approach”, *Journal of Guidance, Control and Dynamics*, Vol. 27, No.1, 2004.

Wenzel, W. and Hamacher, K., “Stochastic Tunnelling Approach for Global Minimization of Complex Potential Energy Landscapes”, *Physical Review Letters*, Vol. 82, No. 15, pp. 3003-3007, April 1999.

Wie, B. “Space Vehicle Dynamics and Control”, AIAA Education Series, AIAA Inc, 1801 Alexander Bell Drive, Reston, VA 20191, 1998.

Wikipedia (2009a), http://en.wikipedia.org/wiki/Gauss-Newton_algorithm, last accessed 24 March 2009.

Wikipedia (2009b), http://en.wikipedia.org/wiki/Levenberg-Marquardt_algorithm, last accessed 24 March 2009.

Williams, K. et al. “Genesis Halo Orbit Station-Keeping Design”, *15th International Symposium on Spaceflight Dynamics*, Biarritz, France, 26-30 June, 2000.

WMAP Website, “Lagrange Points of the Sun-Earth system”, http://map.gsfc.nasa.gov/media/990529/index.html. Last accessed: 4 June 2009.

Wong, H. and Kapila, V. “Adaptive Nonlinear Control of Spacecraft Near Sun-Earth L_2 Lagrange Point”, *IN: Proceedings of the American Control Conference*, Denver, Colorado, 4-6 June, 2003.

Xin, M., Balakrishnan, S.N. and Pernicka, H.J. “Deep-Space Spacecraft Formation Flying Using Θ-D Control”, *AIAA Guidance, Navigation and Control Conference and Exhibit, Providence*, Rhode Island, 16-19 August, 2004. AIAA 2004-4784.

Xin, M., Balakrishnan, S.N. and Pernicka, H.J. “Position and Attitude Control of Deep-Space Spacecraft Formation Flying Via Virtual Structure and Θ-D Technique”, *AIAA Guidance, Navigation and Control Conference and Exhibit*, San Francisco, California, 15-18 August, 2005. AIAA 2005-6090.

Yang, G., Yang, Q., Kapila, V., Palmer, D. and Vaidyanathan, R. “Fuel Optimal Manoeuvres for Multiple Spacecraft Formation Reconfiguration Using Multi-Agent Optimization”, *International Journal of Robust and Nonlinear Control*, Vol. 12, pp. 243-283, 2002.

Zeilik, M and Smith, E.v.P., “Introduction to Astronomy and Astrophysics, Second Edition”, Saunders Collage Publishing, Philadelphia, PA, Chapter 7, pp. 135-140, 1987.

Zink, M., Krieger, G., Fiedler, H. and Moreira, A. “The TanDEM-X Mission: Overview and Status”, *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 45, Issue 11, Part 1, pp. 3317-3341, 2007.