

Autonomous and Autonomic Systems: With Applications to NASA Intelligent Spacecraft Operations and Exploration Systems

Walt Truszkowski Harold L. Hallock Christopher Rouff Jay Karlin James Rash Mike Hinchey Roy Sterritt



The NASA Monographs in Systems and Software Engineering series addresses cutting-edge and groundbreaking research in the fields of systems and software engineering. This includes in-depth descriptions of technologies currently being applied, as well as research areas of likely applicability to future NASA missions. Emphasis is placed on relevance to NASA missions and projects.

Walt Truszkowski, Harold L. Hallock, Christopher Rouff, Jay Karlin, James Rash, Mike Hinchey, and Roy Sterritt

Autonomous and Autonomic Systems: With Applications to NASA Intelligent Spacecraft Operations and Exploration Systems

With 56 Figures



Walt Truszkowski NASA Goddard Space Flight Center Mail Code 587 Greenbelt MD 20771, USA Walt.Truszowski@nasa.gov

Harold L. Hallock NASA Goddard Space Flight Center (GSFC) Greenland MD 20771, USA harold.l.hallock@nasa.gov

Christopher Rouff Lockheed Martin, Advanced Technology Laboratories Arlington, VA, USA christopher.rouff@lmco.com

Jay Karlin Viable Systems, Inc. 4710 Bethesda Ave. #516 Maryland MD 20814, USA jkarlin1@comcast.net James Rash NASA Goddard Space Flight Center Mail Code 585 Greenbelt MD 20771, USA James.L.Rash@nasa.gov

Mike Hinchey Lero-the Irish Software Engineering Research Centre University of Limerick Limerick Ireland

Roy Sterritt Computer Science Research Institute University of Ulster Newtownabbey, County Antrim Northern Ireland r.sterritt@ulster.ac.uk

ISSN 1860-0131 ISBN 978-1-84628-232-4 e-ISBN 978-1-84628-233-1 DOI 10.1007/978-1-84628-233-1 Springer Dordrecht Heidelberg London New York

British Library Cataloguing in Publication Data A catalogue record for this book is available from the British Library

Library of Congress Control Number: 2009930628

© Springer-Verlag London Limited 2009

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may only be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licenses issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers.

The use of registered names, trademarks, etc., in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

The publisher makes no representation, express or implied, with regard to the accuracy of the information contained in this book and cannot accept any legal responsibility or liability for any errors or omissions that may be made.

Cover design: SPi Publisher Services

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

In the early 1990s, NASA Goddard Space Flight Center started researching and developing autonomous and autonomic ground and spacecraft control systems for future NASA missions. This research started by experimenting with and developing expert systems to automate ground station software and reduce the number of people needed to control a spacecraft. This was followed by research into agent-based technology to develop autonomous ground control and spacecraft. Research into this area has now evolved into using the concepts of autonomic systems to make future space missions self-managing and giving them a high degree of survivability in the harsh environments in which they operate.

This book describes much of the results of this research. In addition, it aims to discuss the needed software to make future NASA space missions more completely autonomous and autonomic. The core of the software for these new missions has been written for other applications or is being applied gradually in current missions, or is in current development. It is intended that this book should document how NASA missions are becoming more autonomous and autonomic and should point to the way of making future missions highly autonomous and autonomic. What is not covered is the supporting hardware of these missions or the intricate software that implements orbit and attitude determination, on-board resource allocation, or planning and scheduling (though we refer to these technologies and give references for the interested reader).

The book is divided into three parts. The first part gives an introduction to autonomous and autonomic systems and covers background material on spacecraft and ground systems, as well as early uses of autonomy in space and ground systems. The second part covers the technologies needed for developing autonomous systems, the use of software agents in developing autonomous flight systems, technology for cooperative space missions, and technology for adding autonomicity to future missions. The third and last part discusses applications of the technology introduced in the previous chapters to spacecraft constellations and swarms, and also future NASA missions that will need the discussed technologies. The appendices cover some detailed information on spacecraft attitude and orbit determination and some operational scenarios of agents communicating between the ground and flight software. In addition, a list of acronyms and a glossary are given in the back before the list of references and index.

In Part One of the book, Chap. 1 gives an overview of autonomy and autonomic systems and why they are needed in future space missions. It also gives an introduction to autonomous and autonomic systems and how we define them in this book. Chapter 2 gives an overview of ground and flight software and the functions that each supports. Chapter 3 discusses the reasons for flight autonomy and its evolution over the past 30 plus years. Chapter 4 mirrors Chap. 3 for ground systems.

In Part Two, Chap. 5 covers the core technologies needed to develop autonomous and autonomic space missions, such as planners, collaborative languages, reasoners, learning technologies, perception technologies and verification and validation methods for these technologies. Chapter 6 covers designing autonomous spacecraft from an agent-oriented perspective. It covers the idea of a flight software backbone and the spacecraft functions that this backbone will need to support, subsumption concepts for including spacecraft functionality in an agent context, and the concept of designing a spacecraft as an interacting agent. Chapter 7 covers the technologies needed for cooperative spacecraft. It starts by discussing the need for cooperative spacecraft, a model of cooperative autonomy, mission management issues for cooperation, and core technologies for cooperative spacecraft. Chapter 8 covers autonomic systems and what makes a system autonomic, why autonomicity is needed for future autonomous systems, and what functions would be needed to make future missions autonomic.

Part Three starts with Chap. 9, which discusses spacecraft constellations, cooperation between or among the spacecraft in the constellation, difficulties in controlling multiple cooperative spacecraft, and a multiagent paradigm for constellations. Chapter 10 gives an overview of swarm technology, some example missions that are being proposed that use this technology, and issues in developing the software for swarm-based missions. Chapter 11 discusses some future missions that NASA is planning or developing conceptually. This chapter discusses how the technology discussed in the previous chapters would be applied to these missions, as well as additional technology that will need to be developed for these missions to be deployed.

The Appendix offers additional material for readers who want more information concerning attitude and orbit determination and control, or concerning operational scenarios of agents communicating between the ground and flight software. This is followed by a list of acronyms used in the book and a glossary of terms. All references are included in the back of the book.

There are three types of people who will benefit from reading this book. First are those who have an interest in spacecraft and desire an overview of ground and spaceflight systems and the direction of related technologies. The second group comprises those who have a background in developing current flight or ground systems and desire an overview of the role that autonomy and autonomic systems may play in future missions. The third group comprises those who are familiar with autonomous and/or autonomic technologies and are interested in applying them to ground and space systems.

Different readers in each of the above groups may already have some of the background covered in this book and may choose to skip some of the chapters. Those in the first group will want to read the entire book. Those in the second group could skip Chap. 2 as well as Chaps. 3 and 4, though the latter two may be found interesting from an historical view. The third group of people could skip or skim Chap. 5, and though they may already be familiar with the technologies discussed in Chaps. 6–8, they may find the chapters of interest to see how AI technologies are applied in the space flight domain.

We hope that this book will not only give the reader background on some of the technologies needed to develop future autonomous and autonomic space missions, but also indicate technology gaps in the needed technology and stimulate new ideas and research into technologies that will enable future missions possible.

MD, USA MD, USA VA, USA MD, USA MD, USA Limerick, Ireland Belfast, Northern Ireland Walt Truszkowski Lou Hallock Christopher Rouff Jay Karlin James Rash Mike Hinchey Roy Sterritt

Acknowledgements

There are a number of people who made this book possible. We would like to thank them for their contributions. We have made liberal use of their work and contributions to Agent-based research at NASA Goddard.

We would like to thank the following people from NASA Goddard Space Flight Center who have contributed information or some writings for material that covered flight software: Joe Hennessy, Elaine Shell, Dave McComas, Barbara Scott, Bruce Love, Gary Welter, Mike Tilley, Dan Berry, and Glenn Cammarata.

We are also grateful to Dr. George Hagerman (Virginia Tech) for information on Virginia Tech's Self Sustaining Planetary Exploration Concept, Dr. Walter Cedeño (Penn State University, Great Valley) for information on Swarm Intelligence, Steve Tompkins (NASA Goddard) for information on future NASA missions, Drs. Jonathan Sprinkle and Mike Eklund (University of California at Berkeley) for information on research into UUVs being conducted at that institution, Dr. Subrata Das and his colleagues at Charles River Analytics of Boston for their contribution of information concerning AI technologies, and to Dr. Jide Odubiyi for his support of Goddard's agent research, including his major contribution to the development of the AFLOAT system.

We would like to thank the NASA Goddard Agents Group for their work on the agent concept testbed (ACT) demonstration system and on spacecraft constellations and the NASA Space Operations Management Office (SOMO) Program for providing the funding. We would also like to thank Drs. Mike Riley, Steve Curtis, Pam Clark, and Cynthia Cheung of the ANTS project for their insights into multiagent systems.

The work on autonomic systems was partially supported at the University of Ulster by the Computer Science Research Institute (CSRI) and the Centre for Software Process Technologies (CSPT), which is funded by Invest NI through the Centres of Excellence Programme, under the EU Peace II initiative.

X Acknowledgements

The work on swarms has been supported by the NASA Office of Systems and Mission Assurance (OSMA) through its Software Assurance Research Program (SARP) project, Formal Approaches to Swarm Technologies (FAST), administered by the NASA IV&V Facility and by NASA Goddard Space Flight Center, Software Engineering Laboratory (Code 581).

Contents

Part I Background

1	Inti	Introduction					
	1.1	Direct	tion of New Space Missions	5			
		1.1.1	New Millennium Program's Space Technology 5	5			
		1.1.2	Solar Terrestrial Relations Observatory	6			
		1.1.3	Magnetospheric Multiscale	7			
		1.1.4	Tracking and Data Relay Satellites	8			
		1.1.5	Other Missions	8			
	1.2	Auton	nation vs. Autonomy vs. Autonomic Systems	9			
		1.2.1	Autonomy vs. Automation	9			
		1.2.2	Autonomicity vs. Autonomy	10			
	1.3	Using	Autonomy to Reduce the Cost of Missions	13			
		1.3.1	Multispacecraft Missions	14			
		1.3.2	Communications Delays	15			
		1.3.3	Interaction of Spacecraft	16			
		1.3.4	Adjustable and Mixed Autonomy	17			
	1.4	Agent	Technologies	17			
		1.4.1	Software Agents	19			
		1.4.2	Robotics	21			
		1.4.3	Immobots or Immobile Robots	23			
	1.5	Summ	nary	23			
2	Ove	erview	of Flight and Ground Software	25			
	2.1	Groun	nd System Software	25			
		2.1.1	Planning and Scheduling	27			
		2.1.2	Command Loading	28			
		2.1.3	Science Schedule Execution	28			
		2.1.4	Science Support Activity Execution	28			
		2.1.5	Onboard Engineering Support Activities	28			

		2.1.6	Downlinked Data Capture	29
		2.1.7	Performance Monitoring	29
		2.1.8	Fault Diagnosis	29
		2.1.9	Fault Correction	30
		2.1.10	Downlinked Data Archiving	30
		2.1.11	Engineering Data Analysis/Calibration	30
		2.1.12	Science Data Processing/Calibration	31
	2.2	Flight	Software	31
		2.2.1	Attitude Determination and Control, Sensor	
			Calibration, Orbit Determination, Propulsion	33
		2.2.2	Executive and Task Management, Time Management,	
			Command Processing, Engineering and Science Data	
			Storage and Handling, Communications	34
		2.2.3	Electrical Power Management, Thermal Management,	
			SI Commanding, SI Data Processing	34
		2.2.4	Data Monitoring, Fault Detection and Correction	34
		2.2.5	Safemode	35
	2.3	Flight	vs. Ground Implementation	35
9	El:	lat A	tonomy Evolution	97
0	1 Hg		a for Elight Autonomy	01 90
	0.1	211	Satisfying Mission Objectives	- 30 - 30
		3.1.1 3.1.2	Satisfying Spacecraft Infrastructure Needs	- 39 - 47
		3.1.2 3.1.3	Satisfying Operations Staff Needs	50
	32	Brief I	History of Existing Flight Autonomy Capabilities	54
	0.2	3 2 1	1970s and Prior Spacecraft	55
		322	1980s Spacecraft	57
		323	1990s Spacecraft	59
		324	Current Spacecraft	61
		325	Flight Autonomy Capabilities of the Future	63
	3.3	Currei	t Levels of Flight Automation/Autonomy	66
			5	
4	Gro	und A	utonomy Evolution	69
	4.1	Agent-	-Based Flight Operations Associate	69
		4.1.1	A Basic Agent Model in AFLOAT	70
		4.1.2	Implementation Architecture for AFLOAT Prototype	73
		4.1.3	The Human Computer Interface in AFLOAT	75
	1.0	4.1.4	Inter-Agent Communications in AFLOAT	76
	4.2	Lights	Out Ground Operations System	78
		4.2.1	The LOGOS Architecture	18
	4.9	4.2.2	An Example Scenario	80
	4.5	Agent	Oromient of the ACT A cost Auglitude	81 01
		4.3.1	Overview of the AU1 Agent Architecture	81
		4.3.2	Architecture Components	83
		4.3.3	Datanow Between Components	87

4.3.4	ACT Operational Scenario	88
4.3.5	Verification and Correctness	90

Part II Technology

5	Cor	re Tech	nnologies for Developing Autonomous	
	and	Auto	nomic Systems	95
	5.1	Plan 7	Technologies	95
		5.1.1	Planner Overview	95
		5.1.2	Symbolic Planners	98
		5.1.3	Reactive Planners	99
		5.1.4	Model-Based Planners	100
		5.1.5	Case-Based Planners	101
		5.1.6	Schedulers	103
	5.2	Collab	orative Languages	103
	5.3	Reason	ning with Partial Information	103
		5.3.1	Fuzzy Logic	104
		5.3.2	Bayesian Reasoning	105
	5.4	Learni	ing Technologies	106
		5.4.1	Artificial Neural Networks	106
		5.4.2	Genetic Algorithms and Programming	107
	5.5	Act Te	echnologies	108
	5.6	Percep	otion Technologies	108
		5.6.1	Sensing	108
		5.6.2	Image and Signal Processing	109
		5.6.3	Data Fusion	109
	5.7	Testin	g Technologies	110
		5.7.1	Software Simulation Environments	110
		5.7.2	Simulation Libraries	112
		5.7.3	Simulation Servers	113
		5.7.4	Networked Simulation Environments	113
6	Age	ent-Bas	sed Spacecraft Autonomy Design Concepts	115
	6.1	High I	Level Design Features	115
		6.1.1	Safemode	116
		6.1.2	Inertial Fixed Pointing	116
		6.1.3	Ground Commanded Slewing	117
		6.1.4	Ground Commanded Thruster Firing	117
		6.1.5	Electrical Power Management	118
		6.1.6	Thermal Management	118
		6.1.7	Health and Safety Communications	118
		6.1.8	Basic Fault Detection and Correction	118
		6.1.9	Diagnostic Science Instrument Commanding	119
		6.1.10	Engineering Data Storage	119

6.2	Remot	e Agent Functionality11	9
	6.2.1	Fine Attitude Determination	0
	6.2.2	Attitude Sensor/Actuator and Science Instrument	
		Calibration	1
	6.2.3	Attitude Control	1
	6.2.4	Orbit Maneuvering	2
	6.2.5	Data Monitoring and Trending12	2
	6.2.6	"Smart" Fault Detection, Diagnosis, Isolation,	
		and Correction	3
	6.2.7	Look-Ahead Modeling12	3
	6.2.8	Target Planning and Scheduling	3
	6.2.9	Science Instrument Commanding and Configuration 12	4
	6.2.10	Science Instrument Data Storage	
		and Communications12	4
	6.2.11	Science Instrument Data Processing12	4
6.3	Space	craft Enabling Technologies	5
	6.3.1	Modern CCD Star Trackers	5
	6.3.2	Onboard Orbit Determination	5
	6.3.3	Advanced Flight Processor	6
	6.3.4	Cheap Onboard Mass Storage Devices12	6
	6.3.5	Advanced Operating System	6
	6.3.6	Decoupling of Scheduling from Communications 12	7
	6.3.7	Onboard Data Trending and Analysis12	7
	6.3.8	Efficient Algorithms for Look-Ahead Modeling12	7
6.4	AI En	abling Methodologies12	7
	6.4.1	Operations Enabled by Remote Agent Design12	8
	6.4.2	Dynamic Schedule Adjustment Driven by Calibration	
		Status	9
	6.4.3	Target of Opportunity Scheduling Driven by Realtime	
		Science Observations	9
	6.4.4	Goal-Driven Target Scheduling	0
	6.4.5	Opportunistic Science and Calibration Scheduling 13	1
	6.4.6	Scheduling Goals Adjustment Driven by Anomaly	
		Response	1
	6.4.7	Adaptable Scheduling Goals and Procedures	2
	6.4.8	Science Instrument Direction of Spacecraft	
		Operation	2
	6.4.9	Beacon Mode Communication	3
	6.4.10	Resource Management	4
6.5	Advan	tages of Remote Agent Design13	4
	6.5.1	Efficiency Improvement	5
	6.5.2	Reduced FSW Development Costs	7
6.6	Missio	n Types for Remote Agents	8
	6.6.1	LEO Celestial Pointers	9
	6.6.2	GEO Celestial Pointers	1
	6.6.3	GEO Earth Pointers	1

		6.6.4	Survey Missions	142
		6.6.5	Lagrange Point Celestial Pointers	142
		6.6.6	Deep Space Missions	144
		6.6.7	Spacecraft Constellations	144
		6.6.8	Spacecraft as Agents	145
	-			
7	Coc	operat	ive Autonomy	147
	7.1	Need	for Cooperative Autonomy in Space Missions	148
		7.1.1	Quantities of Science Data	148
		7.1.2	Complexity of Scientific Instruments	148
	7.0	7.1.3	Increased Number of Spacecraft	148
	7.2	Gener	al Model of Cooperative Autonomy	149
		7.2.1	Autonomous Agents	149
		(.2.2)	Agent Cooperation	101
	7 2	1.2.3 Space	cooperative Actions	100
	(.5	5pace	Science Dispring	100
		729	Mission Dlanning	150
		733	Socuence Planning	157
		7.3.3 7 3 1	Command Sequencer	158
		7.3.4	Science Data Processing	158
	74	Space	craft Mission Viewed as Cooperative Autonomy	158
		7.4.1	Expanded Spacecraft Mission Model	158
		7.4.2	Analysis of Spacecraft Mission Model	161
		7.4.3	Improvements to Spacecraft Mission Execution	162
	7.5	An Ez	xample of Cooperative Autonomy: Virtual Platform .	164
		7.5.1	Virtual Platforms Under Current Environment	165
		7.5.2	Virtual Platforms with Advanced Automation	166
	7.6	Exam	ples of Cooperative Autonomy	167
		7.6.1	The Mobile Robot Laboratory at Georgia Tech	169
		7.6.2	Cooperative Distributed Problem Solving Research	
			Group at the University of Maine	169
		7.6.3	Knowledge Sharing Effort	170
		7.6.4	DIS and HLA	170
		7.6.5	IBM Aglets	171
0	A		· C	179
8		Over	ic Systems	1/3
	0.1	8 1 1	What are Autonomic Systems?	173
		819	Autonomic Properties	175
		813	Necessary Constructs	175
		814	Evolution vs. Revolution	178
		815	Further Reading	170
	8.2	State	of the Art Research	
	0.2	8.2.1	Machine Design	
		J		

	8.2.2	Prediction and Optimization	 	 . 180
	8.2.3	Knowledge Capture and Representation	 	 . 181
	8.2.4	Monitoring and Root-Cause Analysis	 	 . 181
	8.2.5	Legacy Systems and Autonomic Environments	 	 . 182
	8.2.6	Space Systems	 	 . 183
	8.2.7	Agents for Autonomic Systems	 	 . 183
	8.2.8	Policy-Based Management	 	 . 183
	8.2.9	Related Initiatives	 	 . 184
	8.2.10	Related Paradigms	 	 . 184
8.3	Resear	ch and Technology Transfer Issues	 	 . 185

Part III Applications

9	Aut	onomy	y in Spacecraft Constellations	189
	9.1	Introd	uction	189
	9.2	Const	ellations Overview	190
	9.3	Advar	tages of Constellations	193
		9.3.1	Cost Savings	193
		9.3.2	Coordinated Science	194
	9.4	Apply	ing Autonomy and Autonomicity to Constellations	194
		9.4.1	Ground-Based Constellation Autonomy	195
		9.4.2	Space-Based Autonomy for Constellations	195
		9.4.3	Autonomicity in Constellations	196
	9.5	Intelli	gent Agents in Space Constellations	198
		9.5.1	Levels of Intelligence in Spacecraft Agents	199
		9.5.2	Multiagent-Based Organizations for Satellites	200
	9.6 Grand View			
		9.6.1	Agent Development	204
		9.6.2	Ground-Based Autonomy	204
		9.6.3	Space-Based Autonomy	205
10	Swa	rms ir	a Space Missions	207
	10.1	Introd	uction to Swarms	208
	10.2	Swarn	n Technologies at NASA	209
		10.2.1	SMART	210
		10.2.2	NASA Prospecting Asteroid Mission	212
		10.2.3	Other Space Swarm-Based Concepts	214
	10.3	Other	Applications of Swarms	215
	10.4	Auton	omicity in Swarm Missions	216
	10.5	Softwa	are Development of Swarms	217
		10.5.1	Programming Techniques and Tools	
		10.5.2	Verification	
	10.6	Future	e Swarm Concepts	

11	Concluding Remarks.22311.1 Factors Driving the Use of Autonomy and Autonomicity
	Missions
A	Attitude and Orbit Determination and Control231
В	Operational Scenarios and Agent Interactions 235B.1Onboard Remote Agent Interaction Scenario235B.2Space-to-Ground Dialog Scenario239B.3Ground-to-Space Dialog Scenario240B.4Spacecraft Constellation Interactions Scenario242B.5Agent-Based Satellite Constellation Control Scenario246B.6Scenario Issues247
С	Acronyms
D	Glossary
Ref	ferences
Ind	lex

Background

Introduction

To explore new worlds, undertake science, and observe new phenomena, NASA must endeavor to develop increasingly sophisticated missions. Sensitive new instruments are constantly being developed, with ever increasing capability of collecting large quantities of data. The new science performed often requires multiple coordinating spacecraft to make simultaneous observations of phenomena. The new missions often require ground systems that are correspondingly more sophisticated. Nevertheless, the pressures to keep mission costs and logistics manageable increase as well.

The new paradigms in spacecraft design that support the new science bring new kinds of mission operations concepts [165]. The ever-present competition for national resources and the consequent greater focus on the cost of operations have led NASA to utilize adaptive operations and move toward almost total onboard autonomy in certain mission classes [176, 195]. In NASA's new space exploration initiative, there is emphasis on both human and robotic exploration. Even when humans are involved in the exploration, human tending of space assets must be evaluated carefully during mission definition and design in terms of benefit, cost, risk, and feasibility.

Risk is a major factor supporting the use of unmanned craft: the loss of human life in two notable Shuttle disasters has delayed human exploration [160], and has led to a greater focus on the use of automation and robotic technologies where possible. For the foreseeable future, it is infeasible to use humans for certain kinds of exploration, e.g., exploring the asteroid belt, for which the concept autonomous nano technology swarm (ANTS) mission was posed – discussed in Chap. 10 – where uncrewed miniature spacecraft explore the asteroid belt. A manned mission for this kind of exploration would be prohibitively expensive and would pose unacceptable risks to human explorers due to the dangers of radiation among numerous other factors.

Additionally, there are many possible missions where humans simply cannot be utilized for a variety of reasons such as the long mission timeline reflecting the large distances involved. The Cassini mission taking 7 years to reach Titan, the most important of Saturn's moons, is an example. Another

© Springer-Verlag London Limited 2009

example is Dawn, a mission to aid in determining the origins of our universe, which includes the use of an altimeter to map the surface of Ceres and Vesta, two of the oldest celestial bodies in our solar system.

More and more, these unmanned missions are being developed as autonomous systems, out of necessity. For example, almost entirely autonomous decision-making will be necessary because the time lag for radio communications between a craft and human controllers on earth will be so large that circumstances at the craft will no longer be relevant when commands arrive back from the earth. For instance, for a rover on Mars during the months of mission operations when the earth and Mars are on opposite sides of the sun in their respective orbits, there would be a round-trip delay of upwards of 40 min (earth) in obtaining responses and guidance from mission control.

Historically, NASA missions have been designed around a single spacecraft carrying one instrument, or possibly a small number of related instruments. The spacecraft would send its raw data back to earth to a dedicated ground control system which had a dedicated staff that controlled the spacecraft and would trouble-shoot any problems. Many of the new missions are either very complex, or have long communications delays, or require very fast responses to mission circumstances. Manual control by humans becomes problematic or impractical for such missions. Consequently, the onboard software to operate the missions is increasingly complex, placing increasing demands on software development technologies.

In anticipation of these missions, NASA has been doing research and development into new software technologies for mission operations. Two of these technologies enable autonomous and autonomic operations. Technology for autonomy allows missions to operate on their own with little or no directions from humans based on internal goals. Autonomicity builds on autonomy technology by giving the mission what is called *self-awareness*. These technologies will enable missions to go to new planets or distant space environments without constant real-time direction from human controllers.

Recent robotic missions to Mars have required constant inputs and commands from mission control to move the rovers only inches at a time, as a way to ensure that human controllers would not be too late in learning about changed conditions or circumstances or too late in returning appropriate directions to the rovers. With 20 min required for one-way communications between the earth and Mars, it takes a minimum of 40 min for mission control to receive the most recent video or sensor inputs from a robot and send the next commands back. Great care is required when moving a robot on Mars or other distant location, because if it flipped over or got stuck, the mission could be ended or become severely limited. There often elapsed several hours between the receipt of new images from Mars and the transmission of movement commands back to the robot. The result of the delays was a great, but unavoidable, limitation on exploration by the robots.

If these missions could instead operate autonomously, much more exploration could be accomplished, since the rovers, reacting independently and immediately to conditions and circumstances, would not have the long wait times for communications, or for human recalculation and discussion of the next moves. Until recently, there was not enough computational power onboard spacecraft or robots to run the software needed to implement the needed autonomy. Microprocessors with requisite radiation-hardening for use in space missions now have greatly increased computing power, which means these technologies can be added to space missions giving them much greater capabilities than were previously possible.

This book discusses the basics of spaceflight and ground system software and the enabling technologies to achieve autonomous and autonomic missions. This chapter gives examples of current and near-term missions to illustrate some of the challenges that are being faced in doing new science and exploration. It also describes why autonomy in missions is needed not only from an operations standpoint, but also from a cost standpoint. The chapter ends with an overview of what autonomy and autonomicity are and how they differ from each other, as well as from simple automation.

1.1 Direction of New Space Missions

Many of the planned future NASA missions will use multiple spacecraft to accomplish their science and exploration goals. While enabling new science and exploration, multispacecraft and advanced robotic missions, along with the more powerful instruments they carry, create new challenges in the areas of communications, coordination, and ground operations. More powerful instruments will produce more data to be downlinked to mission control centers. Multispacecraft missions with coordinated observations will mean greater complexity in mission control. Controlling operations costs of such missions will present significant challenges, likely entailing streamlining of operations with fewer personnel required to control a spacecraft.

The following missions that have been recently launched or that will be launched in the near future illustrate the types of missions NASA is planning.

1.1.1 New Millennium Program's Space Technology 5

The New Millennium Program's (NMP) Space Technology 5 (ST5) [171] is a technology mission that was launched in March of 2006 with a 90 day mission life. The goal of the mission was to demonstrate approaches to reduce the weight, size, and, therefore, cost of missions while increasing their capabilities. The science it accomplished was mapping the intensity and direction of magnetic fields within the inner magnetosphere of the earth. To accomplish this, it used a cluster of three 25-kg class satellites (Fig. 1.1). Each micro satellite had a magnetometer onboard to measure the earth's magnetosphere simultaneously from different positions around the earth. This allowed scientists to determine the effects on the earth's magnetic field due to the solar wind and other phenomena. 6



Fig. 1.1. New Millennium Program (NMP) Space Technology 5 (ST5) spacecraft mission (image credit: NASA)

ST5 was designed so that the satellites would be commanded individually and would communicate directly to the ground. There was a 1-week period of "lights out" operation where the microsats flew "autonomously" with preprogrammed commands in a test to determine whether commanding could be done onboard instead of from ground stations. In the future, this approach would allow the spacecraft to react to conditions more quickly and save on ground control costs. Preprogrammed commands are considered only as a step toward autonomy and as a form of automation, since the commands are not determined by internal goals and the current conditions of the spacecraft software.

1.1.2 Solar Terrestrial Relations Observatory

The Solar Terrestrial Relations Observatory (STEREO) mission, launched in March 2006, is studying the Sun's coronal mass ejections (CMEs), which are powerful eruptions in which as much as 10 billion tons of the Sun's atmosphere can be blown into interplanetary space [37]. CMEs can cause severe magnetic storms on earth. STEREO tracks CME-driven disturbances from the Sun to earth's orbit, producing data for determining the 3D structure and dynamics of coronal and interplanetary plasmas and magnetic fields, among other things. STEREO comprises two spacecraft with identical instruments (along with ground-based instruments) that will provide a stereo reconstruction of solar eruptions (Fig. 1.2).



Fig. 1.2. STEREO mission spacecraft (image credit: NASA)

The STEREO instrument and spacecraft will operate autonomously for most of the mission due to the necessity of the two spacecraft maintaining exact distances from each other to obtain the needed stereographic data. The spacecraft use a beacon mode where events of interest are identified. Broadcasts are then made to the ground for notification and any needed support. The beacon mode automatically sends an alert to earth when data values exceed a threshold level. Each of the STEREO spacecraft is able to determine its position, orientation, and orbit, and react and act autonomously to maintain proper position. Autonomous operation means significant savings over manual operation.

1.1.3 Magnetospheric Multiscale

The Magnetospheric Multiscale (MMS) mission, scheduled to launch in 2014, will use a four-spacecraft cluster that will study magnetic reconnection, charged particle acceleration, and turbulence in regions of the earth's magnetosphere (Fig. 1.3). The four spacecraft will be positioned in a hexahedral or quad tetrahedral configuration with separations of 10 km up to a few 10 s of thousands of kilometers [51]. This arrangement will allow three-dimensional (3D) structures to be described in both the magnetosphere and solar wind.

Distances between the Cluster spacecraft will be adjusted during the mission to study different regions and plasma structures. Simultaneous measurements from the different spacecraft will be combined to produce a 3D picture of plasma structures. The spacecraft will use interspacecraft ranging and communication and autonomous operations to maintain the correct configuration and the proper distances between spacecraft.



Fig. 1.3. Magnetospheric Multiscale (MMS) spacecraft (image credit: NASA)

1.1.4 Tracking and Data Relay Satellites

The Tracking and Data Relay Satellites (TDRS) relay communications between satellites and the ground (Fig. 1.4, [186]). The original TDRS required ground commands for movement of the large single-access antennas. TDRS-H, I and J autonomously control the antenna motion and adjust for spacecraft attitude according to a profile transmitted from the ground. This autonomous control greatly reduces the costs of operations and the need for continuous staffing of ground stations.

1.1.5 Other Missions

Other proposed or planned near-term missions that will have onboard autonomy include the following:

- Two Wide-angle Imaging Neutral-atom Spectrometers (TWINS) mission, which will stereoscopically image the magnetosphere.
- Geospace Electrodynamic Connections (GEC) mission, which is a cluster of three satellites that will study the ionosphere-thermosphere (2013+).
- Laser Interferometer Space Antenna (LISA) mission, which will consist of three spacecraft to study gravitational waves (2020).
- Magnetotail Constellation (MC) mission, which will consist of 30+ nanostatellites that will study the earth's magnetic and plasma flow fields (2011+).

Relying on spacecraft that coordinate and cooperate with each other, NASA will be able to perform new science that would be difficult or impossible to do with a single spacecraft. And recognition of the challenges of



Fig. 1.4. Tracking and Data Relay Satellites (TDRS) spacecraft (image credit: NASA)

real-time control of such complex missions would lead naturally to designing them to operate autonomously, with goals set at a higher level by human operators and scientists.

1.2 Automation vs. Autonomy vs. Autonomic Systems

In this section, the differences between automation, autonomy, and autonomicity are discussed. This establishes working definitions for this book, and aids in understanding the current state of flight automation/autonomy/autonomicity.

1.2.1 Autonomy vs. Automation

Since "autonomy" and "automation" seem to have a wide range of definitions, it is appropriate to establish how those terms will be used in the context of this book. Both terms refer to processes that may be executed independently from start to finish without any human intervention. Automated processes simply replace routine manual processes with software/hardware ones, which follow a step-by-step sequence that may still include human participation. Autonomous processes, on the other hand, have the more ambitious goal of emulating human processes rather than simply replacing them.

An example of an automated ground data trending program would be one that regularly extracts from the data archive a set list of telemetry parameters (defined by the flight operations team (FOT)), performs a standard statistical analysis of the data, outputs in report form the results of the analysis, and generates appropriate alerts regarding any identified anomalies. So, in contrast to an autonomous process, in this case, the ground system performs no independent decision making based on realtime events, and a human participant is required to respond to the outcome of the activity.

An automated onboard process example could be an attitude determination function not requiring a priori attitude initialization. The steps in this process might be as follows. On acquiring stars within a star tracker, an algorithm compares the measured star locations and intensities to reference positions within a catalog, and identifies the stars. Combining the reference data and measurements, an algorithm computes the orientation of the star tracker to the star field. Finally, using the known alignment of the star tracker relative to the spacecraft, the spacecraft attitude is calculated. The process is an automated one because no guidance from FOT personnel is required to select data, perform star identification, or determine attitudes. However, the process does not define when the process should begin (it computes attitudes whenever star data are available), simply outputs the result for some other application to use, and in the event of an anomaly that causes the attitude determination function to fail, takes no remedial action (it just outputs an error message). In fact this calculation is so easily automatable that the process described, up to the point of incorporating the star tracker alignment relative to the spacecraft, now can be performed within the star tracker itself, including compensations for velocity aberration.

On the other hand, the more elaborate process of autonomy is displayed by a ground software program that independently identifies when communications with a spacecraft is possible, establishes the link, decides what files to uplink, uplinks those files, accepts downlinked data from the spacecraft, validates the downlinked data, requests retransmission as necessary, instructs the freeing-up of onboard storage as appropriate, and finally archives all validated data. This would be an example of a fully autonomous ground process for uplink/downlink.

Similarly, a flight software (FSW) program that (a) monitors all key spacecraft health and safety (H&S) data, (b) identifies when departures from acceptable H&S performance has occurred and (c) independently takes any action necessary to maintain vehicle H&S, including (as necessary) commanding the spacecraft to enter a safemode that can be maintained indefinitely without ground intervention, would be a fully autonomous flight fault detection and safemode process.

1.2.2 Autonomicity vs. Autonomy

In terms of computer-based systems-design paradigms, autonomy implies "self-governance" and "self-direction," whereas autonomic implies "selfmanagement." Autonomy is self-governance, requiring the delegation of responsibility to the system to meet its prescribed goals. Autonomicity is system self-management, requiring automation of responsibility including some decision making for the successful operation of the system). Thus, autonomicity will often be in addition to self-governance to meet the system's own required goals. It may be argued at the systems level that the success of autonomy requires successful autonomicity. Ultimately, ensuring success in terms of the tasks requires that the system be reliable [158].

For instance, the goals of a system may be to find a particular phenomenon using its onboard science instrument. The system may have autonomy (the self-governance/self-direction) to decide between certain parameters. The goals to ensure the system is fault tolerant and continues to operate under fault conditions, for instance, would not fall directly under this specific delegated task of the system (its autonomy), yet ultimately the task may fail if the system cannot cope with uncertain dynamic changes in the environment. From this perspective, the autonomic and self-management initiatives may be considered specialized forms of autonomy, that is, the autonomy (self-governance, self-direction) is specifically to manage the system (to heal, protect, configure, optimize, and so on).

Taking the dictionary definitions (Table 1.1) of autonomous and autonomic, autonomous essentially means "self-governing" (or "self-regulating," "self-directed") as defined [109, 162]. "Autonomic" is derived from the noun "autonomy," and one definition of autonomous is autonomic, yet the main difference in terms of the dictionary definitions would relate to speed; autonomic being classed as "involuntarily," "reflex," and "spontaneous."

"Autonomic" became mainstream within computing in 2001 when IBM launched their perspective on the state of information technology [63]. Autonomic computing has four key self-managing properties [69]:

- Self-configuring
- Self-healing
- Self-optimizing
- Self-protecting

With these four properties are four enabling properties:

Self-aware: of internal capabilities and state of the managed component Self-situated: environment and context awareness Self-monitor and self-adjust: through sensors, effectors, and control loops

In the few years since autonomic computing became mainstream, the "self-x" list has grown as research expands, bringing about the general term "selfware" or "self-*," yet the four initial self-managing properties along with the four enabling properties cover the general goal of self management [156].

The tiers for Intelligent Machine Design [101, 136, 139, 140] consist of a top level (reflection), a middle level (routine), and a bottom level (reaction). Reaction is the lowest level where no learning occurs, but is the immediate response to state information coming from sensory systems. Routine is the middle level, where largely routine evaluation and planning behaviors take place.

Table 1.1. Dictionary definitions of autonomic, autonomicity and autonomous [109]

au·to·nom·ic (àwt nómmik) *adj.*

- 1. Physiology.
 - (a) Of, relating to, or controlled by the autonomic nervous system.
- (b) Occurring involuntarily; automatic: an autonomic reflex.
- 2. Resulting from internal stimuli; spontaneous.

au·ton·o·mic·i·ty (àwt nóm i síttee)

n.

1. The state of being autonomic.

au·ton·o·mous (aw tónn m s)

adj.

- 1. Not controlled by others or by outside forces; independent: an autonomous judiciary; an autonomous division of a corporate conglomerate.
- 2. Independent in mind or judgment; self-directed.

3.

- (a) Independent of the laws of another state or government; self-governing.
- (b) Of or relating to a self-governing entity: an autonomous legislature.
- (c) Self-governing with respect to local or internal affairs: an autonomous region of a country.

4. Autonomic.

[From Greek autonomos: auto-, *auto-* + nomos, *law*]

It receives input from sensors as well as from the reaction level and the reflection level. Reflection, a meta-process where the mind deliberates about itself, at the top level receives no sensory input and has no motor output; it receives input from below. The levels relate to speed. For instance, reaction should be immediate, whereas reflection is consideration over time. Other variations of this three-tier architecture have been derived in other research domains (see Table 1.2) [158]. Each approach is briefly discussed below.

In the future communications-paradigms research domain, a new construct, a knowledge plane, has been identified as necessary to act as a pervasive system element within the network to build and maintain high-level models of the network. These indicate what the network is supposed to do to provide communication services and advice to other elements in the network [28]. It will also work in coordination with the management plane and data planes. This addition of a knowledge plane to the existing data and management/control planes would form a three-tier architecture with data, management/control, and knowledge planes [28].

In the late 1990s, Defense Advanced Research Projects Agency (DARPA)/ ISO's autonomic information assurance (AIA) program studied defense

		DARPA/ISO's		Self-directing and
Intelligent	Future	autonomic	NASA's	self-managing
machine	comms.	information	science	system
design	paradigm	assurance	mission	potential
Reflection	Knowledge	Mission	Science	Autonomous
	plane	plane		
Routine	Management	Cyber	Mission	Selfware
	control	plane		
	plane			
Reaction	Data	Hardware	Command	Autonomic
	plane	plane	sequence	

Table 1.2. Various similar three-tier approaches in different domains

mechanisms for information systems against malicious adversaries. The program developed an architecture consisting of three planes: mission, cyber, and hardware. One finding from the research was that fast responses are necessary to counter advanced cyber-adversaries [87], similar to a reflex action discussed earlier.

As will be defined later in this book, NASA's science mission management, from a high level perspective, may be classified into:

- Science planning: Setting the science priorities and goals for the mission.
- Mission planning: Involving the conversion of science objectives to instrument operations and craft maneuvering, housekeeping and scheduling, comm-

unications link management, etc.

Sequence planning: Production of detailed command sequence plans.

These versions of a high-level, three-tier view of self-governing and selfmanaging systems may be generalized into autonomous-selfware-autonomic tiers. Of course, this is intended neither to be prescriptive nor to be in conflict with other views of autonomic systems. The intention in examining and viewing systems in this light is to assist in developing effective systems.

1.3 Using Autonomy to Reduce the Cost of Missions

Spacecraft operations costs have increasingly concerned NASA and others and have motivated a serious look into reducing manual control by automating as many spacecraft functions as possible. Under current designs and methods for mission operations, spacecraft send their data (engineering and science) to earth for processing and receive commands from analysts at the control center. As the complexity and number of spacecraft increase, it takes a proportionately larger number of personnel to control the spacecraft. Table 1.3 shows some current and future missions with the number of people needed to operate them [122]. People-to-spacecraft ratios are shown (a) for past and present missions based on current technology, and (b) for expected future multispacecraft missions with the current technology and operations approaches.

			Number of people	Current	Goal
		Number of	to operate with	people:	people:
Year	Mission	spacecraft	current technology	S/C	S/C
2000	WMAP	1	4	4:1	—
2000	Iridium	66	200	3:1	—
2000	GlobalStar	48	100	2:1	—
2007	NMP ST5	3	12	_	1:1
2012	MC	30-40	120-160	_	1:10

 Table 1.3. Ratio goals for people controlling spacecraft

 $W\!MAP$ Wilkinson Microwave Anisotropy Probe; $N\!MP$ New Millennium Program; MC magnetotail constellation

The figure also shows the operator-to-spacecraft goal for the future missions. Missions capable of performing the desired science will achieve the operator-to-spacecraft ratio goals only if designed to operate without intensive control and direct commanding by human operators. Clearly, a combination of automation, autonomy, and autonomicity will be needed.¹

In many cases, multispacecraft missions would be impossible to operate without near-total autonomy. There are several ways autonomy can assist multispacecraft missions. The following section describes some of the ways by which autonomy could be used on missions to reduce the cost of operations and perform new science.

1.3.1 Multispacecraft Missions

Flying multispacecraft missions can have several advantages, including:

- Reducing the risk that the entire mission could fail if one system or instrument fails
- Making multiple observations of an object or event at the same time from multiple locations (giving multiple perspectives or making the equivalent of a large antenna from many small ones)
- Reducing the complexity of a spacecraft by reducing the number of instruments and supporting subsystems
- Replacing or adding an instrument by adding a new spacecraft into an already existing constellation or swarm

The Wilkinson Microwave Anisotropy Probe (WMAP) mission, launched in 2001, was forecast to use an average of four people to operate the mission (Table 1.3). This mission consists of a single spacecraft and utilizes a small number of people for operations. The Iridium satellite network has 66

¹To simplify, in the remainder of this book, since autonomicity builds on autonomy, we will simply refer to the combination of autonomous and autonomic systems simply as autonomy, except where explicitly noted.

satellites, and it is estimated that initially, its primary operations consisted of 200 people, approximately three people per spacecraft – a 25% improvement over the WMAP mission. The GlobalStar satellite system consists of 48 satellites and requires approximately 100 people to operate it, approximately two people per satellite. The Iridium and GlobalStar satellites are mostly homogeneous, which makes for easier operations than heterogeneous satellites. Understandably, many of the future NASA missions are proposing multiple homogeneous spacecraft.

The last two examples in Table 1.3 show future NASA multisatellite missions along with the operation requirement using current technology, and the goal. Keeping to the current technology of approximately four people per spacecraft, the cost of operations will become prohibitive as the number of spacecraft increases. Missions are currently being planned and proposed that will include tens and hundreds of spacecraft. The most effective way of avoiding excessive cost of operations is by reducing the operators-to-spacecraft ratio. To do this, operators need to work at a higher level of abstraction and be able to monitor and control multiple spacecraft simultaneously.

In addition to saving costs on operations, autonomy can play a vital role in reducing the size of the communications components. This, in turn, reduces weight and the cost of the mission not only in components, but also in the amount of lift needed to put the spacecraft into orbit. Historically, the mission principal investigator (PI) would want all of the data to be transmitted back to earth for archival purposes and for rechecking calculations. Earlier, instruments did not generate as much data as they do now and it was not an issue in transmitting it because the onboard resources were available. Newer instruments now produce more data and many mission are now flying more remote orbits. Both of these require higher-gain antennas and more power, which increase costs, including greater cost for launch. An alternative is to do onboard processing of science data or transmit less data by stripping out nonscience data, both of which result in less data to download. This is an example of one tradeoff, i.e., in basic terms, design the mission to either (a) download all of the data, or (b) download only part of it (thereby reducing the cost of one part of the mission) and doing more science by adding more instruments.

1.3.2 Communications Delays

Autonomous onboard software is needed when communications can take more than a few minutes between the spacecraft and the ground. When communications is lengthened, mission risks increase because monitoring the spacecraft in real-time (or near real-time) by human operators on earth becomes less feasible. The mission is then flown less on a real-time, current basis, and the operator needs to stay ahead of the mission by visualizing what is happening and confirming it with returned data.

Though *science of opportunity* is not necessarily restricted by communication delays, it often has required a human in the loop. This can present a problem since the phenomenon may no longer be observable by the time initial data are transmitted back to earth, a human analyzes it, and commands are sent back to the spacecraft. Autonomy can be useful in this area because it enables immediate action to be taken to observe the phenomenon. Challenges in this area include devising the rules for determining whether the new science should interrupt other on-going science. Many factors may be involved, including the importance of the current observation, the time to react to the new science (it could be over by the time the instrument reacts), the spacecraft state, and H&S issues. The rules would have to be embedded in an onboard expert system (or other "intelligent" software) that could be updated as new rules are learned. An example of a mission that does autonomous onboard target-of-opportunity response is the Swift mission. It has a survey instrument that finds possible new gamma ray bursters, determines whether an object has high priority, and has autonomous commanding that can slew the spacecraft so the narrow-field-of-view instrument can observe it.

For spacecraft H&S, a large communications latency could mean that a spacecraft could be in jeopardy unknown to human operators on the earth and could be lost before any corrective commands could be received from earth. As in the case of science of opportunity, many aspects need to be taken into account, to be embedded in onboard software, and to be updateable as the mission progresses.

1.3.3 Interaction of Spacecraft

Spacecraft that interact and coordinate with each other, whether formation flying or performing the science itself, may also have to communicate with a human operator on earth. If the spacecraft does not have onboard autonomy, the human operator performs analysis and sends commands back to the spacecraft. For the case of formation flying, having the formation coordination done autonomously or semiautonomously through inter-spacecraft communications can save the lag time for downloading the appropriate data and waiting for a human operator to analyze it and return control commands. In many cases, inter-spacecraft coordination can also save spacecraft resources such as power and communication. However, the more the spacecraft interact and coordinate among themselves, the larger the onboard memory needed for the state space for keeping track of the interactions.

Whether formation flying needs autonomy depends on how accurately spacecraft separations and orientations need to be maintained, and what perturbing influences affect the formation. If the requirements are loose and the perturbations (relative to the requirements) are small, the formation can probably be ground-managed, with control only applied sporadically. If the requirements are stringent and the perturbations (relative to the requirements) are large, then control will need to be exercised with minimum delay and at high rates, necessitating autonomous formation control.

1.3.4 Adjustable and Mixed Autonomy

Complete autonomy may not be desirable or possible for some missions. In these missions, adjustable and mixed autonomy may need to be used [132]. In adjustable autonomy, the level of autonomy of the spacecraft can be varied depending on the circumstances or the desires of mission control. The autonomy can be adjusted to be either complete, partial, or no autonomy. In these cases, the adjustment may be done automatically by the spacecraft depending on the situation (i.e., the spacecraft may ask for help from mission control), or may be requested by mission control either to help the spacecraft accomplish a goal or to perform an action manually. Challenges in adjustable autonomy include knowing when it needs to be adjusted, as well as how much and how to make the transition between levels of autonomy.

In mixed autonomy, autonomous agents and people work together to accomplish a goal or perform a task. Often the agents perform the low level details of the task (analogous to the formatting of a paper in a word processor), while the human performs the higher-level functions (e.g., analogous to writing the words of the paper). Challenges in this area are how to get humans working with the agents, how to divide the work up between the humans and agents, and how to impart to the humans a sense of cooperation and coordination, especially if the levels of autonomy are changing over time.

1.4 Agent Technologies

Agent technologies have found themselves in many domains with very different purposes and competencies. This section discusses some of the issues involved in the design and implementation of agents, and then it will focus on three important classes: software agents, robots, and immobots (immobile robots).

Figure 1.5 lists some of the attributes used to describe agents. From the top-level viewpoint, the two most important attributes of an agent are its purpose and the domain in which it operates. All other attributes can be inferred from these two. It is from these attributes that the agent will be designed and technologies selected.

Figure 1.5 also shows three important classes of agents. Software agents are intelligent systems that pursue goals for their human owners. An example would be an information locator that receives some objectives from its owner, interacts with electronic information sources, locates the desired information, organizes and prioritizes it, and finally presents it to the owner. Software agents exist in a virtual computer world and their sensors and actuators are distributed among the computer systems with which they interact. They may



Fig. 1.5. Agent attributes

be able to migrate from one system to another, and some software agents can interact with other agents to cooperate on achieving common objectives.

Robots are mobile systems that pursue goals in the physical world. Robots are outwardly focused, with the primary goals of measuring and interacting with the external world using their sensors and actuators. Although the field of robotics has been active for many years, cooperative robotics is a relatively new area of investigation. A space-based robotics example would be the Mars Sojourner whose expertise was collecting scientific information on Mars' surface. This successful agent achieved its goals with constrained autonomy and limited ability to cooperate. After Sojourner's top-level goals were downloaded from earth, it would perform multiple steps to achieve the desired objective. During task execution, it would continually monitor the situation and protect itself from unexpected events.

Immobots are immobile robotic systems that manage a distributed network of physical sensors and actuators. They are inwardly focused, with goals to monitor and maintain the general health of the overall system. There is so far only limited experience with immobots cooperating with other agents. A modern factory floor would be an example of an immobot. Integrated throughout the shop floor is a network of sensors that the immobot monitors along with actuators that are used when some action is needed. If a dangerous situation arises, the human operator must be notified and the situation explained.

These examples are all quite different, but in each of them, some form of computer-based agency is necessary for them to carry out their mission.

The types of sensors and actuators change from agent to agent. Software agents usually sense and manipulate the computational environment by transporting information. This information is a mixture of data to use or interpret and commands to execute. Robots' and immobots' primary sensors measure and sense the physical world and their actuators modify and interact with the physical world. In addition to their primary sensors, many robots have the ability to communicate with other agents to receive orders and report problems and results.

Many agents have mobility. Mobility in robots is easy to understand since many robots travel to achieve their objectives. Some software agents are mobile and are able to move themselves (their code and data) from one computing platform to another. Software agents may move to gain access to resources they cannot access efficiently from their original computing platform.

Agents work in either the physical or the virtual world. When in the physical world, the sensors and actuators take up space, cost money, undergo wear and tear, and consume resources while performing their mission. If the mission goal requires the physical world to be sensed and manipulated, then these costs must be paid. Software agents live in a virtual world made up of one or more computers connected by networking. In this world, moving information fulfills the roles of sensing and acting and the only direct resource consumed is computation.

1.4.1 Software Agents

Software agents are being used in many domains and encompass a wide range of technologies. At least three broad categories of software agents are being developed and applied:

Informational agents: Informational agents interact with their owner to determine the types and quantities of information the owner desires. These agents then utilize electronic sources to locate the appropriate information, which the agent then organizes and formats for presentation to the owner.

As an example, currently FSW provides subscription services so that onboard applications can subscribe to spacecraft ephemeris information and receive that information when it is calculated. An agent-based enhancement of this arrangement might entail a scenario like the following. Suppose that predicted spacecraft ephemeris is generated onboard once per second (which is typical of current spacecraft) and represents, second by second, the best available estimation of the spacecraft position and velocity. The informational agent ensures that the information is provided once per second to applications that need it.

But occasionally the spacecraft needs to plan and schedule an orbit stationkeeping maneuver, which would be planned so as to minimize disruption to high priority onboard activities, but could (depending on orbit geometric constraints) disrupt scheduled routine onboard activities, including science. To devise the plan and incorporate the maneuver into the schedule, it needs the best predicted spacecraft ephemeris data available for an interval covering a few days in the future from the current time. When the spacecraft's onboard maneuver agent determines that this planning must be done, the maneuver agent would contact the informational agent and describe its needs, including both the time interval to be covered and the required modeling accuracy. The informational agent then passes the request to the onboard ephemeris agent, which determines whether or not it can create a special ephemeris product that can meet the maneuver agent's needs. If the ephemeris agent can do so, the informational agent supplies it to the maneuver agent when the data becomes available.

If the accuracy requirement cannot be met using the latest extended precision seed vectors that were most recently received from the ground, the informational agent requests that the communications agent send a message to the ground station that an ephemeris update will have to be uplinked before the maneuver planning agent can do its job.

Normally, the ground would then schedule a tracking event to update the ephemeris knowledge and uplink an update to the spacecraft. However, if time is short, the ground may elect to use the improved ephemeris information to plan and schedule the next stationkeeping maneuver and uplink it (including the entire schedule with all the spacecraft attitude adjustment and subsequent rocket- or thruster-firing parameters) to the spacecraft. The informational agent would then pass on the stationkeeping maneuver (and the whole schedule itself) to the onboard scheduling agent, informing the maneuver agent that it no longer needs to worry about the next stationkeeping maneuver.

Personal assistant agents: Personal assistant agents act like a personal secretary or assistant. They know the owner's goals, schedule, and personal demands and help the owner manage his or her activities. More advanced forms of assistants can interact with other agents or humans to offload activities.

One could imagine adapting and employing a personal-assistant type of agent in the spacecraft operations domain by providing such an assistant to a planner/scheduling agent in a scenario like the following. Suppose that normally the planner/scheduler agent receives requests to schedule various activities from other applications (attitude control, orbit maneuvering, power, communications, science instruments, etc.) and interleaves all these requests to produce a conflict-free schedule. However, occasionally a request will come either from an odd source (like a realtime ground command) or a request might arrive after the schedule has been generated and would affect the viability of activities already scheduled.

A personal assistant could be useful to the scheduling agent as a means to intercept the "out-of-the-blue" requests and figure out what to do with them. In other words, the personal assistant could decide that it is a request important enough to bother the scheduler regarding changing the existing schedule (a realtime ground command would always be that important), or could decide the request is not important enough to make
a plan modification (in which case it might send it back requesting that the submitter determine whether the request could be re-submitted at a later time). The personal assistant might also determine that other agents need to be consulted and could set up a "meeting" between the various relevant agents to try to resolve the question.

Buying/negotiating agents: Buying/negotiating agents are asked to acquire some product or service. The agent interacts with potential suppliers, negotiates the best overall deal, and sometimes completes the transaction. One could imagine this type of agent adapted for use in the spacecraft operations domain, acting as an electrical power agent tasked with managing onboard power resources, in the following scenario.

Suppose that the power agent monitors overall power resources and has allocated power so as to satisfy all customers' needs (for example, attitude control subsystem (ACS), propulsion, communication, thermal, science instruments 1, 2, and 3, etc.). Suddenly the power agent realizes that a large portion of one of the solar arrays is not producing electrical power, leading to what will soon be an insufficient amount of power to satisfy all the customers' needs.

The power agent knows what the spacecraft critical functions are and immediately allocates whatever power those functions need. There remains enough power to continue to perform some science, but not all the science currently scheduled. At this point, the negotiating agent (possibly the power agent itself) "talks" with each of the three science instruments and the science scheduler to decide how best to allocate the remaining power. The scheduler points out which science (done by which instruments) are the highest priority from a mission standpoint.

The science instruments (which have already been guaranteed the power they will need to enter and maintain safemode) report what their special needs are when they transition from safemode to do their science (for example, warm-up times, re-calibrations, etc.). The scheduler then produces a draft modified schedule, which the negotiating agent checks for power validity, i.e., do they have enough power to "buy" the proposed schedule. If they do, the negotiating agent contacts ground control (through the communications agent) to obtain ground control's blessing to change the schedule, while ground control figures out how to restore (if possible) nominal power capabilities. If ground control cannot be contacted, the negotiating agent will approve executing the draft schedule until it hears from the ground or some new problem develops (or, less likely, the original problem disappears on its own).

1.4.2 Robotics

Intelligent robots are self contained mechanical systems under the guidance of computerized control systems. Intelligent robots have a long history that goes back to the beginning of computer control. While they share many features of software agents, the complex constraints placed on them by their physical environment have forced robot designers to augment the technologies used in software agents or to develop completely different architectures.

The sensors used by robots measure physical quantities such as environmental image properties, direction and speed of motion, and effector tactile feedback. They have many attributes that mean complexity in robot designs. The complex nature of the physical world makes exact sensor measurements difficult. Two readings taken moments apart or readings taken by two "identical" and "healthy" devices often differ. Physical sensors can fail in many ways. Sometime the failure will result in a device that will give correct results intermittently. Some sensors require complex processing and in many situations, the information is difficult to interpret. A vision system using even an off-the-shelf charge-coupled device (CCD) array can easily supply millions of bytes of image data every second. Image processing techniques must be used to examine the data and determine the features of the image relevant to the robot control system.

Actuators are used to make changes in the physical world. Examples are opening a valve, moving a wheel, or firing an engine. Like sensors, actuators can fail in complex ways due possibly to design deficiencies, wear and tear, or damage caused by the environment. The designers of actuators must also deal with complex interactions with the rest of the robot and the environment. For example, if a robot arm and its cargo are heavy, then actions that move the arm will also apply torque to the robot body that can significantly affect the sensor readings in other parts of the robot and can even alter the position of the robot on the supporting surface.

Since both sensors and actuators have complex failure modes, robust systems should keep long-term information on the status of internal systems and develop alternative plans accounting for known failures. The unexpected will happen, so robust systems should actively detect failures, attempt to determine their nature, and plan alternative strategies to achieve mission objectives. Some systems reason on potential failures during planning in an attempt to minimize the effects that possible sensor or actuator failures could have on the ultimate outcome.

Robots exist in a world of constant motion. This requires the robot to continually sense its environment and be prepared to change its plan based on unexpected events or circumstances. For example, a robotic arm attempting to pick up an object in a river bed must be prepared to adapt to changes in the object's position caused by dynamically changing river currents. Many robotic systems use reactive control systems to perform such low-level tasks. They continually sense and analyze the environment and their effect on it and, within constraints, dynamically change their strategy until the objective is achieved. The realities of reactive control systems often make them interact poorly with the slower and symbolic, high-level control systems.

Because of their mobile nature, many robots commit a large percentage of their resources to navigation. Sensors must support detection, measurement, and analysis of all relevant aspects of the environment to enable the robot to see potential paths and recognize potential hazards. Actuators run motor and drive trains to move the robot. Finally, complex software analyzes the current situation in light of the goals and determine the best path to take to reach the destination.

1.4.3 Immobots or Immobile Robots

Immobot is a recent term described by Williams and Nayak [193]. An immobot is a large distributed network of sensors and actuators under the guidance of a computerized control system. While immobots share core robotic technologies, they differ in structure and perspective. Immobots have a robotic control system surrounded by a large number of fixed sensors and actuators connected by a network. These sensors and actuators are physically embedded in the environment they are attempting to measure and control, and the sensors can be located at great distances from the control system.

The primary objectives of immobots are different from those of robots. Robots spend considerable resources (hardware, consumables, and software) on externally focused activities such as navigation, sensing the world, and environmental manipulation. The immobot's sensors and actuators are fixed into their environment and their focus is internal. Their resources are dedicated to managing the environment they control to get the system into appropriate configurations to achieve objectives and to monitor and mitigate problems that occur. Immobots often monitor their systems for years, and yet when certain events occur, they need to react in real-time.

1.5 Summary

Incorporating further degrees of autonomy and autonomicity into future space missions will be crucial not only in reducing costs, but also in making some missions possible at all. The following chapters give an overview of how space and ground software has operated in the past, the enabling technology to make autonomous and autonomic missions possible, some applications of autonomy in past systems, and future missions where this technology will be critical.

Overview of Flight and Ground Software

To provide a context for later chapters, this chapter presents brief summaries of the responsibilities and functionalities of traditional ground and flight systems as viewed from the framework of a total system process, followed by highlights of the key drivers when making flight-ground trades. Details in the areas of attitude and orbit determination and control, mission design, and system engineering [47,84,92,189,191], which are essential for successful space missions, are beyond the scope of this book, but are well developed and interesting in their own right.

2.1 Ground System Software

Traditionally, the ground system has been responsible almost entirely for spacecraft planning and scheduling (P&S), establishment of communications for uplink and downlink, as well as science data capture, archiving, distribution, and (in some cases) processing. The ground has also shouldered the bulk of the calibration burden (both science and engineering) and much of the job of health and safety (H&S) verification. And when major onboard anomalies or failures arise, flight operation team (FOT) personnel are charged with determining the underlying cause and developing a long-term solution.

So the traditional ground system has occupied the ironic position of having most of the responsibility for managing the spacecraft and its activities, and yet (with the exception of the planning and scheduling function) relying on the spacecraft to provide nearly all information required for carrying out those responsibilities. Today, in a more conventional workplace setting, this kind of work organization might be analyzed (from a reengineering perspective) to be an artificially fragmented process with unnecessary management layering leading to degraded efficiency and wasteful costs.

In the context of reengineering, the standard solution to this sort of problem is to re-integrate the fragmented process or processes by empowering the local areas where the information originates, and eliminating unneeded management layers whose only function is to shuttle that information between boxes on a classic pyramid-shaped table of organization while providing nonvalue added redundant checking and counter-checking.

Leaving the conventional fully earth-embedded workplace behind and returning to the modern spacecraft control center, the reengineering philosophy is still a valid one, except now the analysis of the system's fundamental processes must extend into the spacecraft's (or spacecrafts') orbit (or orbits) and must include trades between ground system functionality and flight system functionality, as will be discussed later.

A prerequisite to perform these flight-ground trades is to identify all the components of the spacecraft operations process, initially without considering whether that component is performed onboard or on the ground. The following is a breakdown of operations into a set of activity categories, at least in the context of typical robotic (i.e., uncrewed) space missions. The order of the categories is, roughly, increasing in time relative to the end-to-end operations process, from defining spacecraft inputs to utilizing spacecraft outputs, though some activities (such as fault detection and correction, FDC) are continuous and in parallel with the main line.

- 1. P&S
- 2. Command loading (including routine command-table uplink)
- 3. Science schedule execution
- 4. Science support activity execution
- 5. Onboard engineering support activities (housekeeping, infrastructure, ground interface, utility support functions, onboard calibration, etc.)
- 6. Downlinked data capture
- 7. Data and performance monitoring
- 8. Fault diagnosis
- 9. Fault correction
- 10. Downlinked data archiving
- 11. Engineering data analysis/calibration
- 12. Science data processing/calibration

Many of the operations activity groups listed above currently are partially automated at this time (for example, P&S, and data and performance monitoring), and may well become fully autonomous (within either the ground or flight systems) in the next 10 years. Some of these functions are already largely performed autonomously onboard. A working definition of the difference between autonomy and automation was supplied in Chap. 1. A description of the current state of the art of onboard autonomy/automation will be supplied in Chap. 3. Next, we will briefly discuss each of the steps in the overall spacecraft operations process.

2.1.1 Planning and Scheduling

Especially for low earth orbit (LEO) missions, the ground system P&S function traditionally has been responsible for generating a detailed desired optimized timeline of spacecraft activities, sometimes (as in the case of Hubble Space Telescope (HST)) based on rather complex predictive modeling of the spacecraft's environment and anticipated behavior. The ground system then recasts (and augments) the timeline information in an appropriate manner for processing on the spacecraft. The timeline is then executed onboard via a (largely) time-driven processor. Often along with the nominal, expected timeline, the ground interleaves it with a large array of alternate branches, to be executed in place of the nominal timeline when certain special conditions or anomalies are encountered. The resulting product is a highly coupled, timedependent mass of data, which, in the past, occupied a substantial fraction of available onboard storage.

Ironically, the ground system's creation of the timeline data block is itself a process almost as highly time-dependent as the execution of the actual timeline onboard. HST provides a particularly complex example. Long-term scheduling (look-ahead intervals of several months to a year) was used to block-out accepted proposal targets within allowed geometric boundary conditions. The geometry factors are typically dominated by Sun-angle considerations, with additional contributions from issues such as moon avoidance, maximizing target orbital visibility, obtaining significant orbital dark time or low sky brightness, and meeting linkages between observations specified by the astronomer.

On the intermediate term (a few weeks to a couple of months), LEO spacecraft targets were ordered and scheduled relative to orbital events such as South Atlantic Anomaly (SAA) entrance/exit and earth occultations, and the duration of their associated observations was estimated based on required exposure time computations. Concurrently, support functions requiring scheduling, like communications, were interleaved with the science-target scheduling.

Lastly, on the short term (a day to a week), final detailed scheduling (both of science targets and support functions) to a precision of seconds was performed using the most accurate available model data (for example, the most recent predicted spacecraft ephemeris), with the possibility for including new targets (often referred to as targets of opportunity (TOOs)) not previously considered.

At times, the software needed to support the intermediate and short-term scheduling process performed by the ground system has been massive, complex, and potentially very brittle. Further, multiple iterations of the process, frequently involving a lot of manual intervention (at considerable expense), were often required to produce an error-free schedule. Although considerable progress has been made in streamlining this process and reducing its associated costs, the mathematical modeling remains fairly sophisticated and some amount of operational inefficiency is inevitable due to the necessity of relying on approximations during look-ahead modeling.

2.1.2 Command Loading

By contrast to the P&S function, command loading is quite straightforward. It consists of translating the directives output from planning and scheduling (plus any realtime directives, table loads, etc., generated at and output from the control center) into language/formats understandable by the flight computer and compatible with the communications medium. As communications protocols and the input interfaces to flight computers become more standard-ized, this ground system function will become steadily more automated via commercial off-the-shelf (COTS) tools.

2.1.3 Science Schedule Execution

Science schedule execution refers to all onboard activities that directly relate to performing the science mission. They include target acquisition, science instrument (SI) configuration, and SI operation on target (for example, exposure time management).

2.1.4 Science Support Activity Execution

Science support activities are those that are specifically performed to ensure the success of the science observation, but are not science observations themselves, nor are they routine housekeeping activities pertaining to maintenance of a viable observing platform. They are highly mission/SI specific activities and may include functionality such as optical telescope assembly (OTA) calibration and management and SI direction of spacecraft operation (such as pointing adjustment directives). These activities may be performed in immediate association with ongoing science, or may be performed as background tasks disjoint from a current observation. Although executed onboard, much (if not all) of the supporting calculations may be done on the ground and the results uplinked to the flight computer in the form of tables or commands.

2.1.5 Onboard Engineering Support Activities

Onboard engineering support activities are routine housekeeping activities pertaining to maintenance of a viable observing platform. The exact form of their execution will vary from spacecraft to spacecraft, but general categories are common within a mission type (e.g., geosynchronous earth orbit (GEO) earth-pointer, LEO celestial-pointer, etc.). Engineering support activities include angular momentum dumping, data storage and management, antenna pointing, attitude and orbit determination and/or prediction, attitude control, and orbit stationkeeping. These activities may be performed in immediate association with ongoing science, or may be performed as background tasks disjoint from a current observation. Again, although executed onboard, some of the supporting calculations may be done on the ground and the results uplinked to the flight computer in the form of tables or commands.

2.1.6 Downlinked Data Capture

Capture of downlinked telemetry data is rather straightforward and highly standardized. This ground system function will become steadily more automated via COTS tools.

2.1.7 Performance Monitoring

Monitoring of spacecraft performance and H&S by checking the values of telemetry points and derived parameters is a function that is currently shared between flight and ground systems. While critical H&S monitoring is an onboard responsibility (especially where triggers to safemode entrance are concerned), the ground, in the past, has performed more long-term nonrealtime quality checking, such as hardware component trending and accuracy analysis, as well as analysis of more general performance issues (e.g., overall observing efficiency).

2.1.8 Fault Diagnosis

Often the term "FDC" has been used in connection with spacecraft H&S autonomy. Such terminology tends to conceal an important logical step in the process, which in the past has been exclusively the preserve of human systems engineers. This step is the diagnosis of the fundamental cause of problems based on measured "symptoms."

Traditionally, prior to launch, the systems and subsystem engineers would identify a whole host of key parameters that needed to be monitored onboard, specify tolerances defining in-range vs. out-of-range performance, and identify FSW responses to be taken in realtime and/or FOT responses to be taken in near-realtime. But what has actually occurred is that the engineers have started with a set of failure scenarios in mind, identified the key parameters (and their tolerances/thresholds) that would measure likely symptoms of those failures, figured out how to exclude possible red-herrings (i.e., different physical situations that might masquerade as the failure scenario under consideration), and (in parallel) developed corrective responses to deal with those failures. So the process of transitioning from the symptoms specified by the parameters to the correction action (often a static table entry) that constitutes the diagnosis phase conceptually actually occurs (prelaunch) in the reverse order and the intellectual content is sketchily stored rigidly onboard.

In the postlaunch phase, the systems engineers/FOT may encounter an unanticipated problem and must perform a diagnosis function using the telemetry downlinked to the ground. In such cases, operations personnel must rely on their experience (possibly with other spacecraft) and subject matter expertise to solve the problem. When quick solutions are achieved, the process often used is that of pattern recognition (or, more formally, case-based reasoning, as will be discussed later), i.e., the recognition of a repetition of clues observed previously when solving a problem in the past. The efficient implementation of a capability of this sort lies in the domain of artificial intelligence. Failing at the pattern recognition level, a more lengthy general analytical phase (the human equivalent of state modeling) typically ensues that is manually intensive and at times very expensive.

So a spacecraft called upon to diagnose and isolate its own anomalies is being asked not just to emulate the capabilities of human beings. In fact, it is being asked to emulate the capabilities of the most senior and knowledgeable individuals associated with the operations staff. Therefore, as a FSW implementation of this function must by its nature be extremely costly, a very careful trade must be conducted prior to migrating this function to the spacecraft.

2.1.9 Fault Correction

Currently at GSFC, generating a plan to correct an onboard anomaly, fault, or failure is exclusively a ground responsibility. These plans may be as simple as specification of a mode change, or as complex as major hardware reconfiguration or FSW code modification. In many cases, canned solutions are stored onboard for execution in response to an onboard trigger or ground command, but creation of the solution itself was done by ground system personnel, either in immediate response to the fault, or (at times) many years prior to launch, in anticipation of the fault. And even where the solution has been worked out and validated years in advance, a conservative operations philosophy has often kept the initiation of the solution within the ground system. So at GSFC, although future technical improvements in onboard computing power and artificial intelligence tools may allow broader onboard independence in fault correction, major changes in operations management paradigms will be needed before we see more widespread migration of this functionality onboard.

2.1.10 Downlinked Data Archiving

Archiving of downlinked telemetry data (including, in some cases, distribution of data to users) is rather straightforward and highly standardized. This ground system function will become steadily more automated via COTS tools.

2.1.11 Engineering Data Analysis/Calibration

Traditionally, nearly all spacecraft engineering analysis and calibration functions (with the exception of gyro drift-bias calibration and, for the small explorer (SMEX) missions, magnetometer calibration) have been performed on the ground. These include attitude-sensor alignment and polynomial calibrations, battery depth-of-discharge and state-of-charge analyzes, communications-margins evaluations, etc. Often the work in question has been iterative and highly manually intensive. Some progress has been made toward further automating of at least portions of these tasks, yielding reduced production costs. It appears at this time to be less significant, from a purely cost basis, as to whether this functionality is performed onboard or on the ground.

2.1.12 Science Data Processing/Calibration

Science data processing and calibration have been nearly exclusively a ground system responsibility for two reasons. First, the low computing power of radiation hardened onboard computers (OBCs) relative to that available in ground systems has limited the degree to which science data processing can be performed onboard. Second, the science community generally has insisted that all the science data be brought to the ground. Their position arises from a concern that the data might not be processed as thoroughly onboard as it might be on the ground, and that the science data users often process the same data multiple times using different algorithms, calibrations, etc., sometimes years after the data were originally collected.

Given the science customers' strong views on this subject, independent of potential future advances in radiation hardened processing capabilities, it would be ill-advised to devise a mission concept that relies exclusively on such onboard autonomy features. A more appropriate approach would be to offer these features as options to users, thereby allowing them to take advantage of cost-saving opportunities as they deem appropriate. One can envision a dual scenario where missions not only would send back the raw data for the science community, but also would process it onboard to permit the exercise of onboard autonomy through which the spacecraft might spot potential TOOs and take unplanned science observations without having to wait for possibly (likely) untimely instructions from ground control.

2.2 Flight Software

Although highly specialized to serve very precise (and often mission-unique) functions, FSW must simultaneously satisfy a broad range of competing needs. First, it is the FSW that provides the ground system an interface with the flight hardware, both engineering and science. Since spacecraft hardware components (including SIs) are constantly being upgraded as their technologies continue to advance, the FSW elements that communicate to the hardware components must regularly be updated as well. Fortunately, as the interface with the ground system has largely been standardized, the FSW elements that communicate to the ground remain largely intact from mission to mission. It is in fact the ability of the FSW to mask changes in flight hardware input/output

(I/O) channels that has provided the ground system a relatively stable environment for the development of standardized COTS products, which, in turn, has enabled dramatic reductions in ground system development costs.

Second, it is the responsibility of the FSW to respond to events that the ground system cannot deal with because of the following:

- 1. The spacecraft is out of contact with the ground
- 2. The response must be immediate
- 3. Critical spacecraft or payload issues are involved, or
- 4. The ground lacks key onboard information for formulating the best response

Historically, the kinds of functions allocated to FSW for these reasons were ones such as the attitude control subsystem (ACS), safemode processing and transition logic, fault detection and correction, target acquisition logic, etc.

Third, the FSW can be used to streamline (at least in part) those (previously considered ground system) processes where an onboard, autonomous response is cheaper or more efficient. In many of these cases, routine processes may first be performed manually by operations personnel, following which automated ground software is developed to reduce costs. After the automated ground process has been fully tested operationally, the software or algorithms may then be migrated to the flight system where further cost reductions may be achievable.

Fourth, the process may be performed onboard in order to reduce demand on a limited resource. For example, downlink bandwidth is a valuable, limited quantity on most missions, either because of size/power constraints on spacecraft antennas/transmitters, or because of budget limitations on the size of the ground antenna. In such cases, FSW may be used to compress the output from payload instruments or prune excessive detail from the engineering telemetry stream to accommodate a smaller downlink data volume.

As can be seen from even casual consideration of these few examples, the demands placed on FSW have a widely varying nature. Some require high precision calculation of complex mathematical algorithms. These calculations often must be performed extremely quickly and the absolute time of the calculation must be accurately placed relative to the availability of key input data (here, we are referring to the data-latency issue). On the other hand, some FSW functions must process large quantities of data or must store and manage the data. Other functions must deal with intricate logic trees and orchestrate realtime responses to anomalies detected by self-monitoring functions. And because the FSW is the key line of defense protecting spacecraft H&S, all these functions must be performed flawlessly and continuously, and for some missions (due to onboard processor limitations), must be tightly coupled in several processing loops.

The following is a list of the traditional FSW functions:

- 1. Attitude determination and control
- 2. Sensor calibration
- 3. Orbit determination/navigation (traditionally orbit maneuver planning has been a ground function)
- 4. Propulsion
- 5. Executive and task management
- 6. Time management
- 7. Command processing (target scheduling is traditionally a ground function)
- 8. Engineering and science data storage and handling
- 9. Communications
- 10. Electrical power management
- 11. Thermal management
- 12. SI commanding
- 13. SI data processing
- 14. Data monitoring (traditionally no trending)
- 15. FDC
- 16. Safemode (separate ones for spacecraft and payload instruments)

2.2.1 Attitude Determination and Control, Sensor Calibration, Orbit Determination, Propulsion

Often in the past, attitude determination and control, sensor calibration, orbit determination/navigation, and propulsion functions have resided within a separate ACS processor because of the high central processing unit (CPU) demands of its elaborate mathematical computations. As OBC processing power has increased, this higher cost architecture has become more rare, and nowadays, a single processor usually hosts all the spacecraft bus functions. Attitude determination includes the control laws responsible for keeping the spacecraft pointing in the desired direction and reorienting the spacecraft to a new direction. Currently, at GSFC, onboard attitude sensor calibration is limited to gyro drift-bias calibration (and for some spacecraft, a coarse magnetometer calibration).

Orbit determination may be accomplished by measurement (global positioning system (GPS) for example), solving the equations of motion, or by use of an orbit propagator. Traditionally, orbit maneuver planning has been the responsibility of the ground, but some experiments have been performed migrating routine stationkeeping-maneuver planning onboard, e.g., Earth Observing-1 (EO-1). Regardless of whether the orbit-maneuver planning is done onboard or on the ground, the onboard propulsion subsystem has responsibility for executing the maneuvers via commands to the spacecraft's thrusters, which also at times may be used for attitude control and momentum management.

2.2.2 Executive and Task Management, Time Management, Command Processing, Engineering and Science Data Storage and Handling, Communications

Command and data handling (C&DH) includes the executive, time management, command processing, engineering- and science-data storage, and communication functions. The executive is responsible for coordinating and sequencing all the onboard processing, and separate local executives may be required to control lower level processing within a subsystem. The command processor manages externally supplied stored or realtime commands, as well as internally originated commands to spacecraft sensors, actuators, etc. Again, depending on the design, some command management may be under local control.

The C&DH also has management responsibility for engineering- and science-data storage, in the past via tape recorders, but nowadays via solid state storage. Depending on the level of onboard sophistication, much of the bookkeeping job may be shared with the ground, though the trends are toward progressively higher levels of onboard autonomy. Telemetry uplink and downlink are C&DH responsibilities as well, though articulation of moveable actuators (such as high gain antenna (HGA) gimbals) as well as any supporting mathematical modeling associated with communications (e.g., orbit prediction) are typically the province of the ACS.

2.2.3 Electrical Power Management, Thermal Management, SI Commanding, SI Data Processing

Critical H&S functions like spacecraft electrical power and thermal management are usually treated as separate subsystems, though the associated processing may be distributed among several physical processor locations (or located in the spacecraft bus processor) depending on the design of the flight system. This distribution of subfunctionality is particularly varied with regard to SI commanding and data processing, given the steadily increasing power of the SIs' associated microprocessors. Currently, any onboard processing that is associated with a spacecraft's Optical Telescope Assembly (OTA) falls within the context of the SI functions, though as OTA processing becomes more autonomous with the passage of time, it could well warrant independent treatment.

2.2.4 Data Monitoring, Fault Detection and Correction

The processing associated with data monitoring and FDC is even more highly distributed. Typically, the checking of individual data points and the identification of individual errors (with associated flag generation) are done locally, often immediately after the measurement is read out from its sensor. On the other hand, fault correction is typically centralized, so responses to multiple faults can be dealt with in a systematic manner.

2.2.5 Safemode

The last item, safemode, may include several independent subfunctions, depending on the cost and complexity of the spacecraft in question. Typical kinds of safemode algorithms include Sun acquisition modes (to maintain power positive, maintain healthy thermal geometry, and protect sensitive optics), spin-stabilized modes (to maintain attitude stability), and inertial hold mode (to provide minimal perturbation to current spacecraft state). Usually, the processing for one or more of these modes is located in the main spacecraft bus processor, but often in the past, there has been a fall back mode in a special safemode processor, the attitude control electronics (ACE) processor, in case the main processor itself has gone down. In addition to its safemode responsibilities, the ACE was the interface with the coarse attitude sensors and actuator hardware, obtaining their output data, and providing command access. The individual SIs themselves also have separate safemode capabilities, executed out of their own processor(s). Anomalies causing the main bus processor to become unavailable are dealt with via a special uplink-downlink card, which, in the absence of the main processor, enables continued (though limited) ground communication with the spacecraft.

2.3 Flight vs. Ground Implementation

Increasing levels of onboard autonomy are being enabled by increases in flight data system capacities (CPU, I/O, storage, etc.), as well as by the new approaches/structures for FSW design and development (object-oriented design, expert systems, remote agents, etc.). In particular, operational activity categories that previously were virtually the private domain of the ground systems (such as P&S, engineering data analysis and calibration, and science data processing and calibration) now provide exciting opportunities for shifting responsibility from the ground to the flight component in order to take advantage of the strengths inherent in a realtime software system in direct contact with the flight hardware.

The key advantages possessed by the flight component over the ground component are immediacy, currency, and completeness. Only the flight component can instantly access flight hardware measurements, process the information, and respond in realtime. For example, for performance of basic spacecraft functions such as attitude control and thermal/power management, only the FSW has direct access in realtime to critical information needed to define the spacecraft's operational state, as well as the direct access to the spacecraft actuator hardware required to create and maintain the desired state. The FSW is also the only component of the integrated flight/ground operational system with full-time access to all relevant information for applications such as fault detection and SI target acquisition.

By contrast, in the past, the advantage of the ground over the flight segment has been the larger, more powerful ground computers that (for example) have enabled the ground system to execute extremely intricate schedule optimization algorithms using highly complex predictive models. However, as the power of the flight computers continues to grow with time, a partial shift of even some of these traditional ground monopolies may be justified to take advantage of the realtime information exclusively available onboard. In fact, as these hardware differences between the two platform environments narrow, the distinction between flight-based vs. ground-based may begin to blur somewhat, bringing with it the potential for more mission-customized allocation of functions between flight and ground systems.

Flight Autonomy Evolution

As new ideas surface for implementing advanced autonomous functions onboard spacecraft, the extent to which spacecraft already possess autonomous capability is often not fully appreciated. Many of these capabilities, in fact, have been in place for so long that they have become absorbed within the flight software (FSW) infrastructure, and as a result, typically are not even considered when FSW autonomy is discussed.

Another aspect of flight autonomy not often formally recognized is that the current state of flight autonomy is actually the product of an implicit process driven by the users and developers of FSW. Each autonomous function in place onboard NASA GSFC spacecraft has been developed either in response to the needs of the users of spacecraft, both the science users and the flight operations team (FOT), or in response to FSW development team insights into how their product can be made more useful to its customers. Because of the rightfully conservative nature of all three groups (scientists, FOT, and FSW developers), the pace of autonomy introduction tends to be measured, evolutionary, and targeted to very specific needs and objectives, rather than sweeping and revolutionary.

Also, the budget process, which typically targets funds to the performance of individual missions rather than allocating large research and development (R&D) funds for the development of generic functionality for future missions, tends to select against funding of major change and select for funding of incremental change. As mission budgets have steadily shrunk, funds available to mission project managers must be dedicated more to flight-proven autonomy functionality applicable to meeting immediate mission needs, as opposed to being used for risky, breakthrough autonomy concepts that might greatly reduce costs of both the current and other missions.

To provide a somewhat more balanced perspective on this issue, the evolving role of flight autonomy in spacecraft operations will be described within the context of uncrewed space missions from the following perspectives:

- 1. The reasons for providing flight autonomy and what autonomous capabilities have been developed to support these needs
- 2. The general time frame in which these capabilities have been developed
- 3. Possible future trends in flight autonomy development

While the material in this chapter relates particularly to science missions (e.g., astronomical observatories, communications satellites, and satellites that observe the surface of the earth), it should be applicable as well to the development of missions that conduct robotic explorations of planets, moons, asteroids, etc. Many of the philosophies, cultures, budgetary constraints, and technologies span across all groups and agencies developing and flying any type of uncrewed mission. Brought to the fore by *crewed* missions, however, are numerous different considerations and constraints under which it makes little sense to try to design crewed assets with autonomy. Consequently, crewed missions are beyond the scope of this book.

3.1 Reasons for Flight Autonomy

Flight autonomy capabilities typically are developed at NASA Goddard Space Flight Center (GSFC) in order to

- 1. Satisfy mission objectives
- 2. Support spacecraft operations
- 3. Enable and facilitate efficient spacecraft operations

The object of a science mission, of course, is to perform the science for which the spacecraft has been constructed in an accurate and efficient manner. As the lifetime of a spacecraft (and the mission as a whole) is limited both by onboard hardware robustness and budget allocations, it is crucially important to optimize science data gathering as a function of time. And since all science observations are not of equal importance, the optimization strategy cannot be simply a matter of maximizing data bit capture per unit time. So a GSFC science mission must be conducted in such a manner as to support efficient, assembly line-like collection of data from routine, preplanned observations while still permitting the flexibility to break away from a programmed plan to exploit unforeseen (in the short term) opportunities to perform time-critical measurements of ground-breaking significance.

Reconciling these inherently contradictory goals requires a subtle interplay between flight and ground systems, with carefully traded allocations of responsibility in the areas of schedule planning and execution. Typically, the superior computing power of the ground system is utilized to optimize longand medium-range planning solutions, while the quick responsiveness of the flight system is used to analyze and react to realtime issues. In the next section, we will discuss what autonomous capabilities have been developed to support the flight system's role, and how they enable higher levels of efficiency and flexibility. To achieve all of its mission objectives, a spacecraft must maintain nearnominal performance over a minimum lifetime. This day-to-day maintenance effort implies the presence of an onboard infrastructure supporting routine activities such as command processing and resource management, as well as performance of these activities themselves. Furthermore, since a spacecraft's lifetime may be compromised by onboard hardware failures, latent software bugs, and errors introduced operationally by the ground system, the infrastructure must include the capability to safeguard the spacecraft against further damage or loss from these causes. Therefore, the flight system must routinely monitor spacecraft health and, on detection of an anomaly, either fix the problem immediately so that the mission can continue, or configure the spacecraft so that (at a minimum) it remains in a benign, recoverable state until the problem can be analyzed and solved by operations support personnel.

In addition, to accomplish its mission objectives as economically as possible, the entire system (spacecraft platform and payload, flight data system, and ground system) must be developed not only within increasingly stringent cost constraints, but must also be designed so as to make the conduct of the mission as inexpensive as possible over the entire mission lifetime. To achieve this, the flight system must be designed to carry out spacecraft activities accurately, efficiently, and safely, while at the same time performing these activities in a manner that reduces the complexity and mission-uniqueness of the flight and ground system, and facilitates the reduction of FOT staffing during routine operations.

In the following sections, each of these three major drivers for flight autonomy will be examined relative to more specific objectives or goals to be achieved and the means by which these goals are achieved. A summary of this breakdown is provided¹ in Table 3.1.

3.1.1 Satisfying Mission Objectives

GSFC spacecraft mission objectives can be grouped in three major classifications:

- 1. Science execution
- 2. Resource management
- 3. Health and Safety maintenance

Put briefly, these objectives encompass what must be done to perform science efficiently, what onboard resources must be managed in support of science execution, and what precautions must be taken to safeguard the spacecraft while these activities are being performed. The flight system has a critical

¹In this table and frequently in this book, the simple term "ground" will be used to mean "ground system" or "ground operations" in reference to personnel and spacecraft control capabilities on earth.

Reason	Objective	Means
Satisfy	Efficient science execution	Stored commanding
mission		Autonomous pointing control
objectives		Autonomous target acquisition
_		Onboard science data packaging
	Efficient resource	Manage computing power
	management	Manage internal data transfer
	5	Manage time
		Manage electrical power
		Manage data storage
		Manage telemetry bandwidth
		Manage angular momentum
		Manage propulsion fuel
	Health and safety	Monitor spacecraft functions
	maintenance	Identify problems
		Institute corrections
Support	Command validation	Verify transmission accuracy
infrastructure		associate command with appl.
		Verify arrival at application
		Check content validity
		Verify successful execution
	Request orchestration	Stored commands
	-	Realtime ground requests
		Autonomous onboard commands
		Event-driven commanding
	Efficient resource	See list above
	management	
	H&S maintenance	See list above
Efficient	Access to S/C systems	Commanding infrastructure
spacecraft		Model parameter modification
operations		FSW code modification
	Insight into S/C systems	Optimized telemetry format
		Multiple telemetry formats
		Telemetry filter tables
	Lifecycle cost minimization	Remove ground from loop
		Break couplings between functions
		Exploit S/C realtime knowledge

 Table 3.1. Reasons for flight autonomy

role to play in each of these areas, both by reliably and predictably executing the ground system's orders and by autonomously reacting to realtime events to provide enhanced value above and beyond a "simple-minded" response to ground requests. In the following subsections, each of the three objectives will be discussed in a general fashion, with some specific examples cited to illustrate general concepts and principles. Later, these topics will be touched upon again in a more detailed fashion when current and possible future flight system autonomous capabilities are discussed.

Flight Autonomy Enablers of Efficient Science Execution

For science that is predictable and can be scheduled in advance (e.g., science that is characteristic of a LEO celestial-pointer spacecraft), the objective is to pack as much science into the schedule as possible with minimal time overhead or wastage. Traditionally, it is the responsibility of the ground system to solve the "traveling salesman" problem by generating an error-free schedule that optimizes data collection per unit time elapsed. Although schedule generation is the most complex part of the problem, the ground cannot cause its schedule to be executed effectively without the cooperation and support of several autonomous flight capabilities.

Command Execution

First, the flight system must execute the activities specified by the ground at the required times. Traditionally, the ground-generated schedule has defined both the activities and their execution times in a highly detailed manner. An activity can be viewed as a collection of directives that specify the steps that must be performed for the activity to be completed successfully. The directives themselves are decomposed into actual flight hardware or software commands that cause the directives to be executed. Depending on the sophistication of the flight system, the decomposition of the directives into commands may be done by the ground system and uplinked in detail to the spacecraft, or may be specified at a much higher level by the ground system, leaving the decomposition job to the flight system. In practice, most missions share the decomposition responsibility between ground and flight systems, trading onboard storage resources vs. onboard processing complexity.

However the decomposition issues are decided, the collection of directives and commands are uplinked and stored onboard until executed by the flight system in one of the following three ways: absolute-timed, relative-timed, or conditional. Note that this discussion is limited to stored commanding, as opposed to other approaches for commanding the spacecraft in realtime. The way the FSW orchestrates the potentially competing demands of stored commanding with realtime commanding – the commanding originating externally from the ground as well as internally from within the flight system – will be touched upon in later sections dealing with FSW infrastructure.

Absolute-timed commands include an attached time specified by the ground defining precisely when that command is to be executed. This approach is most appropriate for a ground scheduling program that accurately models both planned spacecraft activities as well as ground track and external events or phenomena. It allows an extremely efficient packing of activities, provided no serious impacts occur due to unforeseen events. Major anomalous events would break the plan and potentially invalidate all downstream events. In such cases, science observations (especially for LEO celestial-pointers) might have to be postponed until the ground scheduling system can re-plan and intercept the timeline. Alternatively, for events only affecting the current observation, the spacecraft could simply skip the impacted observation and recommence science activities (supported by any required engineering activities, such as antenna slewing) at the start time of the next observation. This latter alternative can be particularly appropriate for earth-pointers, where the spacecraft's orbit will automatically carry it over the next target. For such cases, resolving a problem in the scheduled timeline can simply involve "writing off" the problem target and reconfiguring the spacecraft and SI so that they are in the correct state when the orbit ground track carries them over the next target. Minor anomalous events can be handled by padding the schedule with worst-case time estimates, thereby reducing the operational efficiency, or uplinking potentially extensive "what if" contingency scenarios, increasing demands on onboard storage.

Use of relative-timed commands reduces somewhat the accuracy demands on ground-system modeling. The ground system still specifies an accurate delta-time (with respect to a spacecraft or orbital event) for executing the activity, but the flight system determines when that key event has occurred. Although the treatment of timing issues differs in the two cases (absolute- vs. relative-timed), the treatment of the activity definition (i.e., how the activity is decomposed into directives and commands) would remain the same.

By contrast, conditional commanding requires the flight system to make realtime decisions regarding which directives or commands will be executed, as well as *when* they will be executed. When conditional commanding is employed, the ground specifies a logic tree for execution of a series of directives and/or commands associated with the activity, but the flight system determines when the conditions have been met for their execution, or chooses between possible branches based on observed realtime conditions. For current GSFC spacecraft, conditional commanding typically is used for detailed-level commanding within a larger commanding entity (e.g., an activity), with timepadding used to ensure that no time conflicts will occur, regardless of what decisions are made by the flight system within the conditional block. These various commanding methods provide an infrastructure enabling accurate and effective execution of ground-specified activities. Traditional applications utilizing this commanding infrastructure include pointing control and SI configuration. Conditional commanding can also enable more flexible onboard planning and scheduling functions than would be achieved through absolutetimed commanding, permitting selective target scheduling and autonomous target rescheduling (e.g., High Energy Astronomical Observatory-2's (HEAO-2's) very flexible onboard scheduling scheme driven by its target list).

Pointing Control

The second major enabler of science execution is autonomous pointing control. Without autonomous spacecraft pointing control, efficient spacecraft operations would be impossible. In fact, the survival of the spacecraft itself would be highly unlikely for other than passively stabilized spacecraft. For spacecraft requiring active attitude control, it is the exclusive responsibility of the flight system to maintain the spacecraft at the desired fixed pointing within accuracy requirements, to reorient the attitude to a new pointing (as specified by the ground), or (in the case of survey spacecraft) to cause the attitude to follow a desired trajectory. Put somewhat more simply, it is the flight system's job to point the spacecraft in the direction of a science target and maintain that pointing (or pointing trajectory) throughout the course of the science data collection.

Once the required spacecraft orientation for science activities has been achieved, the SI(s) must be configured to support target identification, acquisition, and observation. Although target identification and acquisition can be performed with the ground system "in the loop" for missions where realtime communications are readily available and stable, typically optimization of operations requires that these functions be conducted autonomously by the flight system. For most missions, routinely supplying realtime science data to the ground is precluded by orbit, geometry considerations, timing restrictions, and/or cost.

As a general rule, the flight system will autonomously identify the science target (sometimes facilitated by small attitude adjustments) and acquire the target in the desired location within the field of view (FOV) of the SI (also, at times, supported by small spacecraft re-orientations). Once these goals have been achieved, the flight system will configure the SI(s) to perform the desired observation in accordance with the activity definition uplinked by the ground. Note that the flight system may even be assigned some of the responsibility for defining the details for how the science observation activity should be performed. This autonomy is enabled by relative-timed and conditional commanding structures. For example, for a LEO spacecraft whose SIs are adversely impacted by energetic particles within the South Atlantic Anomaly (SAA), the flight system could determine start and stop times for data taking relative to exit from and entrance into SAA contours. Or, based on its realtime measurement of target intensity, the flight system could determine via conditional commanding how long the SI needs to observe the target to collect the required number of photons.

Data Storage

Once a target has been successfully acquired and science data start flowing out of a SI, the flight system must store the data onboard in a manner such that unnecessary burdens are not forced on the ground system's archive and science data processing functions. To that end, the FSW may evaluate science data as they are generated. Data passing validation can then be routed to onboard storage, while data failing validation can be deleted, saving onboard storage space, downlink bandwidth, and ground system processing effort. In practice, a requirement of many science missions is that all the raw science data be downlinked; so often this potentially available flight capability is not implemented or exercised. Even for those missions, however, lossless compression of data may be performed onboard (usually in hardware), yielding significant savings in onboard space and bandwidth (as much as a 3-to-1 reduction in data without loss of information content), while at the same time affording the science customer full information, even to the point of backing out the original raw science data. Finally, the flight system can play a valuable role by exploiting the realtime availability (onboard) of both science and engineering data to synchronize time-tagging and even to package data into organized files tailored to the needs of the customer for whom those data are targeted.

Flight Autonomy Enablers of Efficient Resource Management

In addition to these fundamental applications (command execution, pointing control, and data storage) that are the primary components of conducting science, the flight system must also support auxiliary applications associated with managing limited onboard resources, including computing power, internal data transfer, electrical power, data storage, telemetry downlink bandwidth, angular momentum, and rocket/thruster propellant.

Computing Power and Internal Data Transfer

The first two items, computing power and internal data transfer, are managed both through the design of the FSW and realtime monitoring of FSW performance. Traditionally, at its high level design, FSW functions have been carefully scheduled so as to ensure that adequate computational resources are available to permit the completion of each calculation within timing requirements without impacting the performance of other calculations. Although the FSW may often be operated below peak intensity levels, the FSW is designed to be capable of handling worst case demands. Similarly with respect to internal data flows, the bus capacities are accounted for when analyzing the feasibility of moving calculation products, sensor output, and commands through the flight data system. To deal with anomalous or unexpected conditions causing "collisions" between functions from either a CPU or I/O standpoint, the flight system monitors itself in realtime and, in the event of a conflict, will autonomously assign priority to those functions considered most critical. If an acceptable rationing of resources is not proved to be possible, or if the conflict persists for an unacceptably long period of time, the flight system (or a major component of the flight system such as an individual SI) will autonomously transition itself to a state or mode of reduced functionality (usually a safemode) with correspondingly lower CPU and/or I/O demands.

Power Management

The flight system plays a role in electrical power management at three levels: positioning solar arrays (SAs), managing battery charging and discharging, and overall power monitoring and response. For celestial-pointing spacecraft having movable SAs, the FSW will select the appropriate SA position for each attitude that produces the desired energy collection behavior. Usually the position is chosen to optimize power generation, though for missions where over-charging batteries is a concern, the FSW may offset the SA(s) from their optimal position(s). For earth-pointing spacecraft, the FSW will rotate the SA to track the Sun as the spacecraft body is rotated oppositely so as to maintain nadir pointing. The FSW will also autonomously control battery discharging and charging behavior consistent with algorithms defined prelaunch and refined postlaunch by operations personnel. While carrying out these functions on an event-driven basis, the FSW also actively monitors the state-of-charge of batteries. If power levels fall below acceptable minimums, the flight system will autonomously transition itself (or individual, selected components) to a state or mode of reduced functionality (usually a safemode) with correspondingly lower electrical power demands.

Data Storage and Downlink Bandwidth

Onboard data storage utilization and downlink bandwidth allocation typically fall out of trade studies for ground system operations costs. The ground system will plan its observations to ensure that adequate space is available to store any science data collected during an observation. Similarly, FSW development personnel will design the formats of all telemetry structures to ensure that operations personnel have access to key performance data at required frequencies and to guarantee that customers receive their science data packaged appropriately for ground system processing. However, even in these cases dominated by prelaunch considerations, the flight system has its own autonomous realtime role to play. Specifically, the FSW must monitor free space availability on the storage device and, in the event of a potential overflow, determine (based on ground-defined algorithms) priorities for data retention and execute procedures regarding future data collection. It also must continuously construct the predefined telemetry structures and insert fill data as necessary when data items supposed to be present in the telemetry structure are unavailable. Further, to ensure that the necessary link with the ground is maintained to enable successful telemetry downlink, the flight system must appropriately configure transmitters and orient movable antennas to establish a link with the ground antenna.

Angular Momentum and Propulsion

The last two onboard resources, angular momentum (for reaction wheelbased spacecraft) and propulsion subsystem fuel, can be viewed as physical depletable resources, though the first item more strictly describes the physical behavior or state of the spacecraft. For LEO spacecraft, angular momentum (for reaction wheel-based spacecraft) management typically is fully autonomous and is performed via the interaction of magnetic torquer coils with the geomagnetic field. For orbital geometries where the geomagnetic field strength has diminished below useful levels, excess angular momentum must be dumped via a propulsion system of some sort (hot gas thrusters, cold gas thrusters, or ion jets). Where a propulsion system is utilized to dump angular momentum, often the ground's planning and scheduling system will play a role (even dominate) when angular momentum dumping will occur because of safety concerns regarding autonomous thruster commanding. However, even for missions following this conservative operational philosophy, there often will be a contingency mode/state in which autonomous angular-momentum reduction via thruster firing is enabled to deal with inflight anomalies jeopardizing the control and safety of the spacecraft. For spacecraft not using reaction wheels (for example, the future Laser Interferometer Space Antenna (LISA)), angular momentum management is not an issue.

By contrast, management of thruster fuel resources is traditionally almost exclusively a ground responsibility. This allocation of functionality, historically, has been due to the mathematical complexity of orbit maneuver planning and the earlier limited computational power of OBCs. So, if the planning of orbit maneuvers (the activity expending the bulk of the onboard fuel supply) is a province of the ground system, then management of the propulsion subsystem's fuel budget quite logically would belong to the ground as well. Recently however, considerable interest has been generated regarding the feasibility of autonomous performance of spacecraft orbit stationkeeping activities. In its more elaborate form, autonomous orbit stationkeeping may even be performed in support of maintenance of a spacecraft constellation, coordinating the orbital motions of several independent spacecraft to achieve a common goal, also referred to as formation flying. For these applications where planning and scheduling of the orbit maneuvering function itself are moved onboard, migrating management of the fuel resources to the flight system will be necessary as well.

Flight Autonomy Enablers of Health and Safety Maintenance

Although each spacecraft has ideally been designed to support completion of its assigned science program within its nominal mission lifetime, unpredictable, potentially damaging events threatening termination of spacecraft operations inevitably will occur sporadically throughout the course of the mission. Many of these events will develop so quickly that by the time the ground would have recognized the onset of the threat, developed a solution, and initiates a response, conditions would have worsened to the point that loss of the spacecraft is unavoidable. To deal with these highly dangerous potential problems, as well as a host of lesser anomalies, the flight system is provided with an autonomous fault detection and correction (FDC) capability.

The first responsibility of the flight system's FDC is to monitor ongoing spacecraft function. To this end, FSW analysis personnel, in conjunction with systems engineers and operations personnel, identify a rather large number of hardware output items and FSW-computed parameters that together provide a thorough description of the state of the spacecraft. The FSW then samples these values periodically and compares them to nominal and/or required values given the spacecraft operational mode. This comparison may be achieved via simple rules and limit checks, or by running models associated with a state-based system. Regardless of the sophistication of the approach, the result of the procedure will either be a "clean bill of health" for the spacecraft or identification of some element not performing within a nominal envelope.

After identification of the existence of a potential problem, the FSW then autonomously commands an appropriate corrective response. The elaborateness and completeness of the corrective response vary depending both on the nature of the problem and the degree of independence a given mission is willing to allocate to the FSW. Ideally, the level of response would be a precisely targeted correction that immediately restores the spacecraft to nominal function allowing continuance of ongoing science observations. Usually, a complete solution of this sort will be possible only for minor anomalies, or significant hardware problems where an autonomous switch to a redundant component or a transition to an appropriate FSW state may be performed without incurring additional risk. However, for most major inflight problems, the flight system's responsibility is less ambitious. It is usually not tasked with solving the problem, but simply placing the spacecraft in a stable, protected configuration, for example, transitioning the spacecraft (or an SI) to safemode. The spacecraft then remains in this state while ground personnel analyze the problem and develop and test a solution. Once this process has been completed, the flight system is "told" what its job is with respect to implementing the solution, and then proceeds again with conducting the mission once the solution has been installed.

3.1.2 Satisfying Spacecraft Infrastructure Needs

In addition to its direct, active role in achieving overall mission objectives, FSW has a key role to play as the middle-man between the ground system and the spacecraft hardware. To serve effectively in this capacity, the FSW must provide a user-friendly but secure command structure enabling the ground system to make requests of the spacecraft that will be carried out precisely, and yet will not simply be acted on mindlessly in a manner that might put the spacecraft at risk. Further, to ensure the spacecraft is capable of responding to these requests in a timely manner, the FSW performs those routine functions necessary to keep the spacecraft available and in near-nominal operational condition. These functions may be thought of generically as the spacecraft's autonomic system, in much the way that one views a human's respiratory and circulatory system. Since they are so common from spacecraft to spacecraft, and so essential to spacecraft operations, they ironically are often neglected in discussions of spacecraft autonomy. In the following sections, the various elements of this spacecraft operational infrastructure will be described in more detail. Some of this discussion will be a bit redundant with material presented in Sect. 3.1.1, which dealt with how the FSW enables satisfaction of spacecraft mission objectives. Such material is repeated here more in the context of what the flight system must do in order to keep the spacecraft available and responsive to the ground's needs, as opposed to what it does to accomplish what the ground wants done.

Flight Autonomy Enablers of Command Execution: Validation

Earlier, the various stored command types (absolute-timed, relative-timed, and conditional) were described to illustrate how the FSW is able to execute ground requests faithfully, while still exploiting realtime information that was unavailable to the ground system at the time their requests were generated and uplinked – so as to provide a value-added response to the ground's needs. However, to make safe and reliable use of these command structures, the FSW independently validates commands on receipt, validates them again on execution, and monitors their passage through the C&DH subsystem as they make their way to their local destination for execution.

The first step in the process is to verify that a command (or a set of commands) has not been garbled. For this purpose, when the C&DH subsystem first receives a command packet, the C&DH checks the bit pattern in the header and verifies it matches the expected pattern. At the same time, the C&DH examines the command packet at a high level to make sure it recognizes the packet as something an onboard application could execute. Second, when it is time for the command to be executed, the C&DH determines to which application the command should be sent and ships the command out to be loaded into that application's command buffer. The C&DH then looks for a message verifying that the command was successfully loaded into the buffer (i.e., that there was room in the buffer for the command). Third, as the application works its way through the buffer contents, it examines the contents of the individual commands to verify they are valid. It also will check the command itself to verify there are no inherent conflicts to executing that kind of command given the current spacecraft state. Finally, once the application determines that the command may be executed, it carries out its prescribed function or launches it toward its final destination for execution and verifies that it then executes successfully. This multiple-tiered validation process ensures that only valid commands are executed, that they reach their proper destination, and that they are executed properly once they get there.

Flight Autonomy Enablers of Command Execution: Request Orchestration

To this point, discussion of commanding has focused largely on the execution of stored ground requests, primarily relating directly to carrying out the science observing program. However, the FSW must deal not only with this class of commands, but also with realtime ground requests and with commands dealing with engineering and/or housekeeping activities, in many cases originating autonomously from within the FSW itself. In the past, although the FSW was provided with sufficient "intelligence" to keep commands from these various sources from "bumping into" each other, much of this complexity could be managed by the ground. This was especially true for missions where the spacecraft primarily executed absolute-timed stored commands, or for missions like the International Ultraviolet Explorer (IUE) where continuous ground contact (enabled by a Geostationary orbit) allowed the ground to conduct observations via block scheduling and realtime commanding.

However, as more spacecraft take advantage of the benign space environment of earth-sun Lagrange points (e.g., the James Webb Space Telescope (JWST) now in development), the opportunities presented by event-driven commanding orchestrated by a more autonomous onboard scheduling system likely will be increasingly exploited. This, in turn, will place a greater responsibility on the flight system to manage and prioritize commands from these various sources to optimize science data gathering capabilities without jeopardizing spacecraft health and safety.

Flight Autonomy Enablers of Efficient Resource Management

When the ground system generates the spacecraft's science observation schedule, it implicitly makes a series of assumptions regarding the spacecraft state, e.g., that sufficient power is available to operate the hardware required for the observations, that sufficient angular momentum capacity is present to enable the spacecraft to be oriented properly to observe the targets, that enough science data have been downlinked from onboard storage to permit the storage of newly captured science data, etc. These considerations have already been discussed in some detail in Sect. 3.1.1; however, it is worth repeating here simply from the standpoint that some elements of resource management (e.g., computing power and internal data transfer) are so intimately associated with the running of the FSW that they have effectively become part of the spacecraft infrastructure. So, for these resources, one tends to view the job of resource management not as an independent application running on the FSW, but instead as an element of the FSW facilitating the running of applications.

Flight Autonomy Enablers of Health and Safety Maintenance

Previously the FDC capability was described at a high level to illustrate how FSW autonomy is utilized to mitigate health and safety risks that might otherwise lead to onboard failures that could, in turn, result in failure to achieve spacecraft mission objectives. However, FDC can also be viewed as a key component of the spacecraft infrastructure dedicated to maintaining the spacecraft in a suitable state so the ground can schedule its science observations with confidence, with the knowledge that the FSW will be capable of carrying out its directives effectively, reliably, and safely. As with the case of onboard resource management, there are reasonable arguments for viewing FSW FDC capability both as an enabler of achieving mission objectives and as a critical component of the spacecraft infrastructure.

The most important safety check for all spacecraft is to verify that electrical power capacity is adequate to keep the spacecraft alive. Detection of unacceptably low power levels will engender the autonomous commanding of major load-shedding and (usually) transition to safemode for the spacecraft and its SIs. Verifying that no violations of thermal limits have occurred is almost of equal importance as the power checks. Such violations at best may lead to irretrievably degraded science data, and at worst, loss of the SI or even the spacecraft itself due to potentially irreversible hardware failures such as freezing of thruster propellant lines. At a more local level, celestial pointers, in particular, are always very concerned with possible damage to their imaging system and/or SIs due to exposure to bright objects or excessive radiation. Another potentially lethal problem is loss of attitude control. Maintenance of attitude control must be checked both to ensure that the spacecraft is able to acquire and collect data from its science targets, and more importantly to ensure that none of the previously discussed constraints (power, thermal, and bright object avoidance) are violated.

Note that not all these safety checks are performed exclusively onboard. The ground system will normally attempt to ensure that none of its commands knowingly violate any of the constraints described above, and the ground system as well will use telemetered engineering data to monitor the spacecraft state to detect any violations that may have occurred. In practice, the crucial job of maintaining the spacecraft system is distributed between flight and ground systems, but when an offending event occurs in realtime onboard, it is primarily the responsibility of the flight system to be the first to recognize the advent of a problem and to take the initial (although not necessarily definitive) steps to solve the problem.

3.1.3 Satisfying Operations Staff Needs

Unlike more mundance earth-bound situations where users typically are permitted direct, physical access to the hardware they are using, for spacecraft applications the users effectively interface with the spacecraft exclusively

through the FSW, with the exception of the cases of some classes of emergency conditions and some special applications. (For example, if the Main C&DH computer goes down, the ground can communicate directly to the spacecraft via the uplink-downlink card and issue commands such as Turn On Back-up Main C&DH Computer and Switch to Back-up). The FSW provides both access to the spacecraft for commanding purposes and insight into ongoing spacecraft operations via telemetry. Furthermore, because the FSW "sees" in realtime what is happening onboard and can take immediate action in response to what it sees, the FSW often is capable of performing tasks that, if assigned to the ground system and FOT, would be much more expensive to do. These cost savings may arise from lower software costs due to simpler modeling requirements onboard or from effective replacement of human staff hours with FSW functionality. Additionally, where those reductions pertain to replacement of repetitive FOT manual activities, one obtains a cost savings multiplier over the entire mission duration. In the following subsections, each of these three services to the operations team will be discussed in more detail.

Autonomy Enablers of Access to Spacecraft Systems

FSW is the mechanism enabling nearly all access to the spacecraft by the FOT. The FSW provides a commanding infrastructure that translates very precisely what the FOT wants done into appropriate hardware and/or software commands, as discussed in previous sections. By this means, the FSW to a certain degree provides the FOT "hands-on" capability with respect to the various spacecraft hardware elements and subsystems. However, while it creates FOT access to spacecraft systems, the commanding infrastructure (through its validation capabilities) protects the spacecraft from the occasional operational errors that might otherwise lead to irreparable damage to delicate hardware.

The FSW even allows the FOT to modify the operation of FSW functions by changing the values of the key parameters that drive the functions' models. For example, most spacecraft model their own orbital position and the positions of other spacecraft, such as tracking and data relay satellites (TDRS), through the use of onboard orbit propagators. The starting position and velocity from which future positions are calculated can be specified by the ground in a table, updates to which are uplinked as frequently as required in order to maintain ephemeris accuracy requirements. For other missions, ephemeris updates are performed via command structures rather than tables, but the basic process is effectively the same. Also, as a protection against missing a routine ephemeris update (either due to inflight problems or simply a ground operations error), FSW for the medium-class explorers (MIDEX) missions notify the ground if the onboard parameters have grown "stale," eventually causing the FSW to terminate that ephemeris model's processing. And if an erroneous update is made, a continuity check in the FSW (implemented on most missions) can detect the mistake, allowing the FSW to reject the update and continue using the existing parameters until an accurate set is supplied by the ground.

Specifically with respect to missions developed by NASA Goddard, the telemetry monitor (TMON) (and its more recent variant, the telemetry and statistics monitor (TSM)) capability provides a yet more powerful access into spacecraft function. TMON not only allows the FSW to monitor ongoing spacecraft behavior by checking the values of specific telemetry data points, it also includes a command interpreter program that can execute logical statements and act upon them. As the operation of TMON is driven by uplinked data table loads, TMON can be used to add new onboard functionality without modifying the FSW itself, thereby providing the FOT with a way to work around onboard hardware or software problems in a manner less demanding on the FSW maintenance team.

In addition to supplying a means to influence the spacecraft and FSW behavior, the ground can indirectly change the FSW itself. For small or localized changes, a FSW maintenance development team can modify the FSW, uplinking new code that effectively bypasses some existing code elements and substitutes modified, or even entirely new functionality in its place. If the number of modifications becomes excessive, or if the scale of the upgrade is extremely large, a new version of the FSW program may be developed by the maintenance team and uplinked, which, in turn, may be modified as future changes are required. For long duration missions having extensive cruise durations (e.g., Jet Propulsion Lab (JPL) missions that may take years to get on station), the cruise phase often is used to re-write the flight code to compensate for major hardware anomalies experienced after launch, or even to complete development and testing of major functionality not finished prior to launch. By these means (command infrastructure, table modification/command uplink, TMON/TSM, and actual FSW coding changes), the FSW affords the ground a remarkably high capacity for accessing, influencing, and modifying spacecraft behavior in flight.

Flight Autonomy Enablers of Insight into Spacecraft Systems

To support the safe and effective utilization of the access to spacecraft systems afforded by FSW, the FSW must also allow the ground a comparably high level of insight into ongoing spacecraft operations. To this end, spacecraft typically are designed so as to ensure that, catastrophic failures aside, the ground will always receive, at the very least, some minimal level of health and safety telemetry that summarizes the current spacecraft state, along with significant error messages describing what, if anything, has gone wrong.

When operating in a more nominal state, the spacecraft regularly supplies the ground with fairly detailed information concerning the operational behavior of the spacecraft and its various hardware components. Additional telemetry fields are reserved for FSW-calculated products (and FSW intermediate calculation parameters) deemed to be of value by the FOT and other analytical support staff. As engineering and housekeeping downlink bandwidth typically is in short supply due to optimization of communication hardware trades relative to weight and cost, the allocation of telemetry slots and the frequency with which they are reported is usually a rather difficult process as individual subsystem information needs are traded to avoid exceeding communication bandwidth limits.

However, in the event of an inflight anomaly or the scheduling of a special activity such as an instrument calibration, at times the nominally optimized engineering/housekeeping telemetry contents may not be adequate to support the immediate needs of the ground. In such cases, the FOT can utilize a capability provided by the FSW to modify telemetry contents. In the past, this was achieved by "flying" several predefined telemetry formats, with the current one in use being that selected by the ground. In response to changing inflight conditions, the ground could simply command the use of a different telemetry format more appropriate to current needs. The weakness of this approach lay in its inflexibility, i.e., a telemetry format needed to be defined and be onboard already for it to be used. So if some unexpected conditions occurred requiring an allocation of telemetry slots and frequencies different from that supported onboard, not much could be done immediately to address the problem.

By contrast, the filter table approach utilized by recent spacecraft affords much more flexibility. In principle, any desired presentation of telemetryaccessible data points can be achieved. Of course, given a limited bandwidth, increasing the frequency of a given telemetry item can crowd out other important data items, so a careful trade must always be performed before changing filter table settings. However, at least the higher degree of flexibility ensures that any single FSW-accessible data point can, in principle, be viewed by the ground as frequently as desired.

Flight Autonomy Enablers of Lifecycle Cost Minimization

As budgets for overall lifecycle costs steadily decrease, mission planners increasingly look to FSW as a means to reduce continuing costs of operation. Many current onboard autonomous capabilities, though originally implemented to satisfy mission-specific objectives or to promote enhanced spacecraft H&S, in fact also reduce the work load of operations personnel, thereby enabling the reduction of staffing levels without the loss of efficiency or increased risk. For example, for decades, spacecraft have autonomously maintained attitude control, calibrated gyro-drift biases, propagated their orbital ephemeris, managed battery charging, maintained thermal constraints, packaged and stored science and engineering data, checked for limit violations, and (on detection of any violations) either fixed the problem or placed the spacecraft (or localized element) in a benign state pending ground attention. Trying to perform any of these functions with ground personnel "in the loop" would not only be less efficient and less safe, but also be far more expensive than delegating the responsibility to the flight system.

Recently, more elaborate flight-autonomy capabilities have been introduced specifically to reduce operational costs. For example, to break linkages between spacecraft target-pointing and communications (antenna pointing), the Rossi X-ray Timing Explorer (RXTE) mission introduced an autonomous antenna-manager function responsible for selecting the appropriate high gain antenna (HGA) that can be used compatibly with the current spacecraft attitude. This capability not only supported greatly reduced lead times on changing targets to observe a TOO, but also reduced staff efforts (and costs) in scheduling TDRS contact times by eliminating couplings between TDRS scheduling and onboard antenna selection, which often is a factor when optimizing communications contact time.

And for JWST, the use of an onboard event-based scheduler could reduce overhead time (in turn, raising observing efficiency) and reduce both groundsystem modeling costs as well as the need for spacecraft "hand holding." And further in the future, increased onboard processing of science data may not only enable increased capabilities to exploit TOOs detected in real time onboard, it could for some missions also permit a reduced science data downlink volume, with associated operations cost reductions, as the science community gains confidence in the accuracy and reliability on the onboard processed product.

It should be noted that these reductions in operations costs do not themselves come without a cost. The development of new FSW functionality typically is an expensive undertaking, both from the standpoint of coding the new capability and the testing required to ensure that no inadvertent harm is done to the spacecraft. The impact of these software costs, however, is lessened when the new autonomous function is implemented for a long duration mission where the costs can now be traded relative to the, say, 10 years of operational effort that the FSW replaces. Similarly, when several missions can use the new capability, the up-front development costs for the first mission can be seen as a long-term investment yielding savings both on that mission and downstream missions. Hopefully as the expense of developing FSW continues to decline and as greater FSW reuse becomes possible, the trade of continuing operations costs for new FSW autonomy will be an increasingly favorable one.

3.2 Brief History of Existing Flight Autonomy Capabilities

In the previous sections, the reasons for developing flight autonomy and the flight autonomy capabilities that were developed in response to those needs were discussed in some detail. In the following sections, those flight autonomy capabilities will be grouped in accordance with the general time periods in which they were developed at GSFC and examined relative to the contributions they made for specific spacecraft on which the FSW functionality was flown. General time periods reviewed are the 1970s, 1980s, 1990s, and 2000s.

3.2.1 1970s and Prior Spacecraft

During the 1970s, NASA made the first attempt to standardize onboard flight data systems with the creation of the NASA Standard Spacecraft Computer (NSSC), versions I and II. The NSSC-I, a derivative of that flown on the Orbiting Astronomical Observatory-3 (OAO-3) in 1972 and IUE in 1978, was first flown on the Solar Maximum Mission (SMM) (Fig. 3.1) in 1980 (originally scheduled for launch in 1979). Compared to modern flight computers, the NSSC-I was slow, had very limited memory, was cumbersome when performing mathematical functions due to its small word size and lack of floating point arithmetic, and was awkward to program due to the exclusive use of assembly language. However, it was extremely reliable and was used successfully to support the onboard needs of many missions, from SMM in 1980 to the HST payload in 1990.

The NSSC computers and other OBCs with comparable capabilities such as those used on the HEAO series were employed successfully in the 1970s to support a basic foundation of spacecraft autonomy, including stored commanding, telemetry generation, FDC, orbit propagation, and pointing control. Stored commanding capabilities included the (now) standard set of absolutetimed, relative-timed, and conditional commands, as discussed previously. A degree of FDC (for constraints such as bright object avoidance and minimum power levels) also was present in the form of limit checks on key onboard



Fig. 3.1. Solar maximum mission (SMM) (image credit: NASA)

sensor measurements, with autonomous mode transition capabilities to familiar safemodes, such as Sunpoint.

In the area of pointing control, for the case of HEAO-2, pointing control and ground attitude determination accuracy (using star tracker data) were required to be good to 1 and 0.1 arcmin, respectively – a very demanding requirement for the time. And the SMM fine Sun sensor was so well calibrated that its attitude could be controlled autonomously to 5 arcsec with respect to the Sunline. Attitude control was achieved using onboard closed-loop proportionalintegral-derivative (PID) control laws (including feed-forward specification of environmental torques) and Kalman filters (for optimization of attitude-error determination and gyro-drift bias calibration). So fundamentally, the control approaches used on these spacecraft from the 1970s were quite similar to those currently used on modern spacecraft in support of spacecraft slews, target acquisition, angular momentum management, and maintenance of attitude during science observations.

Originally SMM also possessed the rudiments of an autonomous target identification and acquisition capability that presaged the more elaborate capability implemented in HST (as discussed later), which, in turn, was the next step on the road to a fully autonomous TOO response capability. Specifically, when SMM's SI detected a Solar flare, the data were processed onboard, the flare's location was determined, and the spacecraft was autonomously reoriented to observe the phenomenon. This feature enabled a far quicker response than would have been the case if the data processing and commanding responsibility resided in the ground system, allowing time-critical measurements to be made during the early stages of the flare's duration. The Orbiting Solar Observatory-8 (OSO-8), launched in 1975, also could steer its payload platform independently to expedite acquisition of its short-lived targets.

The evolving nature of flight autonomy can be seen within the decade of the 1970s itself just by noting the significant increase in pointing independence between HEAO-1 and HEAO-2. Specifically, HEAO-1 (launched in 1977) relied on the ground to provide it periodic attitude reference updates (every 12 h) based on ground attitude determination. Just 2 years later, HEAO-2 already possessed the capability to compute its own attitude reference update, based on ground-supplied guide-star reference information, a capability also implemented in SMM's ACS for autonomous control of roll about the Sunline. Furthermore, HEAO-2 could autonomously sequence through a weekly target list, adjusting the order of the targets so as to economize on the use of thruster fuel used in momentum dumping, a remarkable degree of independence even relative to the 1990s missions discussed later.

Examination of FDC also shows a dynamic quality. Although many spacecraft flew hard-coded limit checking and response code, SMM's statistical monitor performance function provided additional flexibility. It allowed operations personnel to specify additional FSW parameters to be monitored autonomously onboard beyond those specified in the at-launch flight code, without making a modification to the code itself. This function was itself improved upon in the 1980s with the introduction of TMON, a telemetry monitor that also supported an autonomous onboard response capability.

3.2.2 1980s Spacecraft

Relative to the 1970s, the period of 1980s saw the launch of larger, more expensive, and more sophisticated spacecraft. Many of these spacecraft (e.g., HST and the Compton Gamma Ray Observatory (CGRO)) were supposed to launch in the mid- to late-1980s, but in actuality launched in 1990 because of delays due to the loss of the shuttle Challenger.

In the case of HST, a more powerful OBC (the DF224) enabled the development of more elaborate pointing-related mathematical algorithms as well as a wider variety of safemode options (supported by a larger number of FDC checks) than had been present on previous spacecraft. In particular, use of HST's fine guidance sensors (FGSs) required the development of rather complex (by onboard standards) mathematical algorithms to command the FGSs and process their data. In fact, the processing demands of the FGS functionality was so high that the HST FSW's 10 Hz processing rate was created for and exclusively dedicated to this purpose. The FGS guide-star acquisition algorithms were themselves extremely powerful, exploiting the full commandconstruct repertoire to achieve the intricate branching/looping logic needed to optimize the probabilities for acquiring the guide stars essential for performing HST's science.

The very existence of the 10 Hz processing rate points to an additional noteworthy aspect of HST's autonomy capabilities that often is taken for granted, namely its executive function. For the sake of simplicity, most FSW development efforts try to limit the number of tasks to as few as possible, usually one or two. However, because of HST's unique computational, precision, and timing demands, the HST pointing-control subsystem (PCS) software required five distinct processing rates, namely, 1,000 Hz for the executive, 40 Hz for primary PCS control laws and gyro processing, 10 Hz for FGS processing, 1 Hz for star tracker processing, ephemeris modeling, and FDC, and 1/300 Hz for the minimum energy momentum management control law. Just the management of these very different, and often competing, tasks demonstrated a significant degree of executive autonomy.

HST's FSW also displayed a high level of autonomy in acquiring science targets through "conversations" between the NSSC-I computer supporting SI commanding and processing, and the DF224 computer responsible for spacecraft platform functionality. For example, for science observations where the target direction was not known to a sufficient level of accuracy to guarantee acquisition in an SI's narrow FOV, the DF224 could initiate (through stored commanding) a small scan of the region of the sky surrounding the estimated target coordinates. The passing of the target through the FOV of the SI would then trigger a sudden increase in SI intensity measurements, which would then be noted by the NSSC-I. On completion of the scan, the
NSSC-I could then request (through a limited interface between it and the DF224) that the spacecraft attitude be returned to that pointing. The DF2244 would then create and initiate a realtime slew command satisfying the NSSC-I's attitude change. As with the SMM autonomous target acquisition feature, HST's capability provided a far quicker response than would have been the case had data processing and commanding responsibility resided in the ground system, a very important consideration given the high cost of HST operations and the extraordinary demand for HST observing time.

By contrast, because of its lower pointing accuracy requirements, the Extreme Ultraviolet Explorer (EUVE), also originally scheduled for the 1980s but actually launched in 1991, did not require a level of sophistication in its ACS subsystem as high as that required on HST. However, it did provide a higher level of flexibility with respect to onboard data monitoring and limit checking. EUVE's telemetry monitoring capability (referred to as TMON, and also flown on CGRO and the Upper Atmosphere Research Satellite (UARS)) permitted the user to select, after launch, specific data points to monitor. In addition, TMON provided a limited logic capability to respond onboard to detected spacecraft conditions (observed via limit checks on data or flag checking) through autonomous generation of commands. EUVE's FSW also included a separate statistics monitor program (called SMP) that was later combined with TMON and flown on MIDEX spacecraft as the TSM program (see Sect. 3.1.3). Finally, EUVE possessed an extremely user-friendly tabledriven limit-checking/response system that has served as the model for later missions in the explorer series.

As an early predecessor of true event-driven operations, EUVE utilized an Orbit Time Processor (a table-driven task) that allowed its FSW to define orbit-based events, a variation on the relative-time-based commanding discussed earlier. Occurrence of the event could then trigger a relative time sequence (RTS), a task, or set an event flag that in turn could be monitored by a running task. The EUVE FOT employed this enhancement to the standard stored commanding infrastructure to re-phase the timing of EUVE's survey mode, which operated within a third of an orbit duty cycle. EUVE also used its Orbit Time Processor to send dusk/dawn commands to the survey instrument.

Although HST's FDC capabilities are not as flexible as EUVE's, HST checks for a much wider spectrum of anomalous conditions, with a larger range of autonomous responses. For example, HST provides four distinct software safemodes: inertial hold, a multistaged Sunpoint, zero-gyro (derived from Sunpoint), and spin-stabilized. Also, in response to guide-star reacquisition problems associated with radiation hits on its fine guidance electronics (FGE) following SAA entrances, HST's FSW developers have implemented an FGE memory-refresh function that restores key FGS commanding parameters to their latest values prior to the SAA entrance. Note that many of HST's FDC capabilities were added postlaunch in response to problems experienced inflight, which not only illustrates the power of FSW to solve unanticipated operational problems that can be dealt with no other way, but also provides

a strong argument in favor of creating and maintaining a strong FSW maintenance team capable of responding to those operational problems when they do occur.

3.2.3 1990s Spacecraft

Relative to the 1970s and 1980s, the 1990s witnessed major hardware and infrastructure advances that enabled greater onboard capabilities. The flight computers were more powerful, with larger memories, and were faster, enabling more sophisticated algorithms and models. Floating point arithmetic and higher level languages (such as C, C++, and Ada) allowed FSW code to be written more like comparable ground system code. For example, objectoriented design concepts can now be used to make flight code more re-usable, and in the long run, potentially cheaper. Thanks to high capacity, lightweight, and cheap solid state storage devices, larger amounts of science data may be stored onboard and packaged more conveniently (with respect to end-user needs) without undue concern for added overhead space costs (although in practice this gain has been largely offset by corresponding increases in SI output data volume). More sophisticated operating systems are now available to handle the masses of data and manage the more elaborate computations. The cumulative result of this technological progress has been to enable a series of new individual flight autonomy capabilities targeted to the needs of specific missions, as well as to support the development of entirely new FSW concepts.

To meet demanding time requirements for TOO response, the RXTE FSW included three new flight autonomy capabilities: onboard target quaternion computation, target quaternion validation, and the antenna manager. The first two enable a science user to specify simply the target's right ascension and declination, and whether there are any special roll coordinate needs. The FSW then takes this targeting information expressed in the natural "language" of the user, transforms it appropriately (i.e., into quaternion format) for use in slewing to and acquiring the target, quality-assures the attitude vs. Sunangle avoidance, and then slews the spacecraft to point to the target at the ground-specified time.

A new RXTE autonomy capability was proposed as a post-launch update, but could not be funded, which would have greatly enhanced RXTE's already superb TOO response time. If RXTE's all sky monitor (ASM) detected the signature of a possible gamma ray burster (GRB), the ASM FSW could have determined the celestial coordinates of the potential TOO. After verifying that those coordinates had not previously been observed, the ASM could then have communicated the GRB celestial coordinates to the OBC, which could then have utilized RXTE's existing capabilities to compute and validate the new target quaternion. Next, the FSW could have autonomously determined the right time to break away from currently scheduled observations, slew to that target, and then generate the appropriate SI configuration commands so that observations by RXTE's proportional counter array (PCA) could be made. Finally, using onboard SAA contour information and spacecraft ephemeris, the FSW could have determined the time at which those PCA observations could commence relative to earth occultation and SAA entrance times. Although this capability was not flown on RXTE, it has been implemented successfully and is the key to satisfying the rapid TOO response time requirement, for the Swift mission (see Sect. 3.2.4).

The third item among the list of autonomy features implemented prelaunch – the antenna manager – enabled de-coupling of science and communications scheduling. For missions where HGA selection is preplanned by the ground, a change in target attitude due to inclusion of a TOO can induce changes in HGA commanding for that target observation period and even future target observation periods downstream. Further, for ground algorithms that optimize TDRS switchover times and HGA selection choices in order to maximize total TDRS contact time, relatively small changes in the attitude profile can cause major changes in desired TDRS schedule. By contrast, RXTE's antenna manager allowed the FSW to determine in realtime which was the best HGA to use to close the link with the scheduled TDRS spacecraft, based on the FSW's realtime knowledge of its attitude and the relative orbital positions of RXTE and TDRS (derived from onboard orbit models). Also, knowing the TDRS schedule, RXTE's FSW could autonomously determine when HGA transmitters should be turned on and when playbacks from the solid state recorder (SSR) should start. These autonomous features were used routinely with great success until a transponder failure eliminated the two-HGA capability.

RXTE and the Tropical Rainfall Mapping Mission (TRMM) also provided enhanced flexibility in telemetry formatting via their telemetry filter tables. In practice, the same amount of planning effort (the most laborious part of the job) would be required to make major changes to its standard telemetry configurations (identified as operationally necessary prelaunch) as would be the case for earlier approaches to telemetry formatting, but once determined, the modified telemetry formatting could be implemented via a simple table uplink as opposed to a FSW change.

RXTE and TRMM also flew a more sophisticated version of TMON, called TSM (originally developed for the Solar Anomalous and Magnetospheric Particle Explorer (SAMPEX) mission), which supported all the functionality of TMON, i.e., monitoring telemetry points, performing limit checking, and initiating stored command sequences and associated event messages on limit failure. In addition, TSM maintained statistical data for each monitor point and accepted FSW reconfiguration commands. Statistical information generated includes telemetry-point current value and time, minimum and maximum values seen with associated times, average value, and number of times the monitor point has been seen. TSM was particularly useful to the RXTE mission as a means to deal with star tracker problems experienced inflight without having to implement major changes in the flight code itself.

Also in the area of diagnostic functionality, Landsat-7 (launched in 1999) experimented with a high level FDC capability, in addition to flying a more traditional one. The FDC functions for most GSFC spacecraft have a one-to-one quality, i.e., a trigger is received and an associated response is executed. The potential problem with this approach is that a higher level problem (for example, a failure in the ACE) can corrupt output data from several ACS sensors that could be misinterpreted as the individual failures of all those ACS sensors, potentially resulting in unnecessary autonomous switches to their redundant components. To protect against a problem of this kind, Landsat-7 implemented a Boolean logic function that would examine all error flags generated at a component level and, by comparing the error flag pattern to a set of patterns maintained onboard defining the signature of higher level problems, deduce the true cause of the current anomaly and respond accordingly. By setting the counters associated with the triggers for the higher level failures to lower values than for the counters associated with the component level failures, the Landsat-7 FSW would be able to switch out the higher level hardware element before the cascade to redundant components commenced.

At a somewhat more detailed level, SAMPEX incorporated a new autonomous calibration function, which also has been flown on other spacecraft in the small explorer (SMEX) series. SAMPEX possesses the capability to calibrate its magnetometer coupling constants (relative to the magnetic torquer bars) inflight, relieving the FOT of the burden of collecting the necessary engineering data, processing it, and uplinking the modified calibration parameters to the spacecraft.

Lastly, some very interesting new ideas have been implemented at JPL. Because of the long cruise periods until their spacecraft achieve their mission orbits or swing-bys, JPL has the luxury of experimenting with their FSW after launch and even making wholesale changes (or simply completing the original coding effort) after launch. Their deep space missions also, by their very nature, may require more autonomy than is typical of GSFC missions. Because of the long communications-delay times inherent in a deep space mission and because of the time critical aspects associated with celestial flybys, JPL has been experimenting with autonomous target identification and acquisition functions that are more elaborate than those flown at GSFC. At a more fundamental structural level, the New Millennium Program's Deep Space One (DS1) FSW was initially designed with Remote Agents having responsibility for multitask management, planning and scheduling, and model-based FDC [99,108]. In practice (due to schedule conflicts), the mission was flown using a more conventional FSW implementation, but the Remote Agent-based version was activated briefly for test purposes.

3.2.4 Current Spacecraft

A number of interesting new autonomy capabilities have been flown on GSFC spacecraft launched in the 2000s. First, the recent development of quaternion

star trackers has provided spacecraft like the Willsinson Microwave Anisotropy Probe (WMAP) (launched in 2001), a true "Lost in Space" capability. Previously, the limited star catalogs flown on GSFC spacecraft, along with the simple star identification algorithms utilized in the FSW, required that fairly accurate *a priori* spacecraft attitude knowledge be available onboard for reliable fine attitude updates to be performed. Now, however, star trackers are available that use more extensive internal catalogs and more powerful star identification algorithms to provide quaternion information to the FSW without previous attitude knowledge. This new autonomy capability both supports ongoing science observing and streamlines recovery from safemode entry. These new star trackers also output the change in attitude, providing a direct back-up and sanity check to the primary body-rate data supplied by the gyros.

Second, the earth observing spacecraft (EOS) Aqua (formerly EOS-PM, launched in 2002) has implemented an autonomous communication capability referred to as "Call 911." When a serious anomaly occurs on the spacecraft, a stored command sequence reconfigures the communications downlink (from the spacecraft to the TDRS system (TDRSS) to the ground station) and broadcasts an unscheduled multiple access (MA) message via the TDRS Demand Access capability. The message is forwarded from White Sands to the EOS Aqua control center, triggering an alarm that unexpected telemetry has been received. The telemetry provides a status message describing the anomaly. The ground can then be ready for contingency commanding at the next scheduled ground contact, or declare an emergency and schedule TDRSS S-band single access (SSA) contact time.

Third, on the Swift spacecraft (launched 2004), the key to the rapid TOO response to detected GRBs was a capability considered previously as a postlaunch update to the RXTE's FSW (see Sect. 3.2.3). When Swift's survey instrument (the burst alert telescope (BAT)) detects the signature of a possible GRB, the FSW determines the celestial coordinates of the potential TOO. After verifying that those coordinates had not previously been observed, the FSW communicates the GRB celestial coordinates to the OBC, which then computes and validates the new target quaternion. The FSW autonomously determines the right time to break away from currently scheduled observations, "swiftly" slews to that target, and generates the appropriate SI configuration commands so that high precision observations by Swift's narrow field instruments (NFI) (the X-ray telescope (XRT) and UV/optical telescope (UVOT)) can be made. Swift also (via TDRSS) can respond to TOO alerts identified by other observatories. This new autonomy function is a very significant first step in the direction of "smart" SIs controlling the science mission and all resources required to perform the science mission, as opposed to the traditional operational approach in which the spacecraft/ground system controls the mission and configures the SIs to perform the observations.

Fourth, a capability was considered for the Triana mission (launch postponed indefinitely for budgetary reasons) that would have utilized science data to point Triana's imaging instrument at the earth. The FSW for the Epic SI used to image the earth could have been used to process the science data in order to derive the earth centroid. The centroid data would then be communicated to the spacecraft platform FSW for use in improving the accuracy of the spacecraft's pointing toward the earth center, or a region offset from the center. This basic autonomy capability will, however, fly on the Solar Dynamics Observatory (SDO) (scheduled for launch in late 2009). SDO's guide telescopes (providing precision-pointing support to its science instruments) will supply data to SDO's FSW, which will then compute a Sun centroid to support direct autonomous pointing of SDO's science instruments at the Sun without realtime ground-processing of science data.

Fifth, an experiment in spacecraft formation flying was performed in 2001 using the EO-1 (launched in 2000) mission and the existing Landsat-7 mission. Landsat-7 was a passive participant, simply executing its normal science mission. EO-1, equipped with a global positioning system (GPS) receiver to measure the EO-1 orbital coordinates in realtime and an orbit propagator supplying predictive Landsat-7 orbital coordinates, maintained approximately a 1 min separation between its orbit and the Landsat-7 orbit. This experiment successfully demonstrated an important capability that can be used by future earth science constellation missions to synchronize science data taken at different local times and to use images gathered by the "lead" spacecraft over the target to optimize science instrument configuration on the trailing spacecraft, or establish for the trailing spacecraft that the target is "socked-in" so that advance preparations for viewing the next target can begin.

3.2.5 Flight Autonomy Capabilities of the Future

Future GSFC missions are expected to advance current onboard capabilities significantly in the areas of planning and scheduling and FDC. JWST (and several other missions currently under development) have proposed the use of onboard event-driven scheduling to exploit the benign thermal environment of the L2 Lagrange point (Fig. 3.2). An observation plan execution (OPE) function would enable the spacecraft to move through its observation schedule on an as-ready basis, rather than pausing to hit absolute time points dictated by traditional fixed-time scheduling approaches. On the other hand, if anomalous conditions occurred that precluded observing the desired target (for example, guide stars not being acquired), the OPE function would simply move on to the next target on the list without further loss of observing time. So the use of the OPE function should produce some gains in overall observing efficiency. Further, by taking advantage of realtime knowledge onboard concerning the spacecraft's angular momentum, the OPE function could intelligently plan when to perform necessary angular momentum dumps with minimal impact to science observations.

Other major enhancements in FSW capabilities are likely to be driven by the needs of the interferometry missions proposed to study star formation,



Fig. 3.2. Lagrange points (not to scale)

planet formation, etc. In the past, spacecraft were launched mounting relatively smaller science instruments that were pointed at their science targets by pointing the entire spacecraft (e.g., HST). Some science instruments were equipped with swivels that allowed the science instrument to be pointed independently and, in the case of survey SIs, the swivel could be rotated continuously to map out a swath of the sky (e.g., RXTE and Swift). On a few missions, the survey function was carried out without a swivel by continuously spinning and precessing the entire spacecraft (e.g., WMAP). But for interferometry missions, whose performance capabilities are driven by the length of their baseline, in effect a small spacecraft bus supports a very large science instrument (on the order of many tens of meters long). And for some interferometry missions currently on the drawing boards, to achieve even larger baselines, the science instrument is a composite object that is an amalgam of the individual light collecting capabilities of many individual detector spacecraft whose data are consolidated within a hub spacecraft. For these missions, preproposal planning usually concentrates on developing a feasible design for the science instrument, paying less attention to the spacecraft bus whose design needs are often assumed to be satisfiable by an existing Rapid Spacecraft Development Office (RSDO) spacecraft design. This is a major paradigm change from GSFC's earlier approach of paying equal or greater attention to the spacecraft bus design during the early planning phases.

The most interesting developments in flight autonomy may be those features required to support formation flying and spacecraft constellations. Such missions will demand a much heightened degree of spacecraft self-awareness and self-direction, as well as an awareness of the "outside world." Until now, a spacecraft has needed to be knowledgeable regarding outside "entities" to the extent that it needed to use them. For example, to use a TDRS spacecraft to communicate with the ground, a spacecraft had to know both its own ephemeris and that of the TDRS spacecraft. But for a constellation of spacecraft (for example, for a distributed interferometry mission) to maintain the collective grouping needed to achieve their overall mission objectives, one or more (potentially even all) of the constellation will have to possess key knowledge of all near-neighboring (or possibly all) constellation members in order to synchronize orbital positions, SI configurations, onboard data processing and communication schedules, etc.

The LISA mission is a particularly high technology example of this kind of constellation mission. LISA, the first space-based attempt to detect gravitational radiation, will consist of three spacecraft maintaining a triangular formation, with each leg of the triangle being 5 million kilometers in length (Fig. 3.3). The constellation will be located within the earth's orbit about the Sun, about 20° "behind" the earth. Each spacecraft in the constellation will mount two lasers. Each laser will be directed toward a cube (called a proof mass) floating "drag free" within a containment cell housed within one of the other spacecraft. So each spacecraft mounts two lasers and two cubes, and each triangle leg is formed by two laser beams. The first beam is directed



Fig. 3.3. Laser beam exchanges between the three laser interferometer space antenna (LISA) spacecraft

from the master spacecraft to a "slave" spacecraft, which phase locks its laser to the incoming beam and directs its reply back to the master. The master then mixes the incoming light with a small fraction of its original outgoing light to produce an interference pattern that can be processed to determine changes in distance between the free-floating proof masses good to better than the size of an atom. Once ground software processing (utilizing as input data from all 6 laser links) has eliminated the many noise effects and other perturbations that can mask the desired signal, the remaining information can be used to detect the presence of gravity waves (whose existence is predicted by Einstein's General Theory of Relativity, but has not as yet been directly detected), which when passing through the antenna will cause the proof masses to move apart by an amount comparable to the sensitivity of the measuring apparatus.

3.3 Current Levels of Flight Automation/Autonomy

To better understand where those new automation/autonomy opportunities may reside, it is useful to associate the items on the previous list of operational activities with rough estimates of the activity's current level of flight autonomy: "high", "medium", "low", or "not applicable". The annotated list appears in Table 3.2.

Activity 2, command loading, is the ground activity responsible for uplinking data to the spacecraft. It is already a mostly automated process, and within the next 10 years, will likely be a fully autonomous ground process. Downlinked data capture and archiving (activities 6 and 10) also have been automated and will probably be fully autonomous ground processes in the

Activity	Current flight automation/ autonomy level
1. Planning and scheduling	Low
2. Command loading	n.a.
3. Science schedule execution	Medium
4. Science support activity execution	Medium
5. Onboard engineering support activities	High
6. Downlinked data capture	n.a.
7. Data and performance monitoring	Medium
8. Fault diagnosis	Low
9. Fault correction	Low
10. Downlinked data archiving	n.a.
11. Engineering data analysis/calibration	Low
12. Science data processing/calibration	Low

 Table 3.2. Operational activities with rough estimates of current level of flight autonomy

67

next decade. Part of data capture is data validation, which is a necessary part of the flight/ground "conversation" required for onboard science data storage management.

As all three of these areas are concerned nearly exclusively with processing data within the ground system itself, there is no real role for the flight system to play in expediting the process directly. However, there are indirect supporting functions such as onboard packaging of data prior to downlink and initiation of the communications link between flight and ground where the flight system could play a larger role. Routine onboard operations of this sort are subsumed under activity 5, onboard engineering support activities, as will be discussed later.

The remaining nine operational areas all offer significant opportunities for an expanded, autonomous flight presence. Those areas labeled "low" (activities 1, 8, 9, 11, and 12) currently are largely ground dominated, but could be at least partially migrated onboard to produce overall system cost and/or efficiency gains.

The "medium" activity areas (3, 4, and 7) already are performed onboard, but either there will be room for expanded functional scope (potentially replacing ground effort) or the ground system typically would have to generate some support products to simplify current onboard processing. So cost/efficiency gains potentially can be realized either by ground-to-flight migration or by introducing entirely new functionality to the flight system. Finally, the activity area 5 labeled "high" is today already fully autonomous onboard, but new functionality could be introduced to produce improved system performance.

Ground Autonomy Evolution

Having focused on automation and autonomy in space-based systems in the previous chapter, attention is now directed toward ground-based systems, particularly the automation of spacecraft control centers. We describe a strategy for automating NASA ground-based systems by using a multiagent system to support ground-based autonomous satellite-subsystem monitoring and report generation supporting mission operations. Over the last several years, work has progressed on developing prototypes of agent-based control centers [2,36,86,134]. With the prototypes has come an improved understanding of the potentials for autonomous ground-based command and control activities that could be realized from the innovative use of agent technologies. Three of the prototypes will be described: Agent-based Flight Operations Associate (AFLOAT), Lights Out Ground Operations System (LOGOS), and Agent Concept Testbed (ACT).

4.1 Agent-Based Flight Operations Associate

AFLOAT was a prototype of a multiagent system designed to provide automated expert assistance to spacecraft control center personnel. The overall goals of AFLOAT were to prototype and evaluate the effectiveness of agent technology in mission operations.

The technical goals of AFLOAT were to:

- 1. Develop a robust and complete agent architecture
- 2. Address the full spectrum of syntactic and semantic issues associated with agent communication languages (ACLs)
- 3. Develop and evaluate approaches for dealing with the dynamics of an agent community which supports collaborative activities
- 4. Understand the mechanisms associated with goal-directed activities
- 5. Develop a full range of user-agent interface capabilities including the development and use of user modeling techniques to support adaptive user interfaces and interactions

The following discusses the structure, architecture, behaviors, communication, and collaboration required to coordinate the activities of the agents in supporting unattended mission operations in a satellite control center.

4.1.1 A Basic Agent Model in AFLOAT

An agent is a computer-based autonomous process that is capable of goaldirected activities [17,45] and can accomplish tasks delegated to it with minimal reliance on human intervention. As it is goal-directed, it can allow the user to specify simply what he/she wants [38], leaving to the agents how and where to get the information or services. Each agent is also able to participate in cooperative work as an associate with humans or as part of a community of cooperating agents.

The agent architecture used in AFLOAT provided appropriate structural elements and behavior to support basic requirements for adaptive reasoning [172]. An agent is adaptive to the extent that it can respond to short-term and long-term changes in its operational environment, deal with unexpected events and reason opportunistically, maintain a focus of attention among multiple goals, and select and perform appropriate actions.

Structural Elements of an AFLOAT Agent

Each agent in AFLOAT had three structural components (Fig. 4.1):

- 1. An inter-agent communication interface
- 2. A monitor
- 3. A knowledge base

The inter-agent communication interface was responsible for validating the inter-agent semistructured language format, sending outgoing messages, receiving incoming messages, and broadcasting messages to other agents. The monitor was responsible for monitoring interactions between agents, incoming and outgoing messages, and the state of the agent, and maintaining a history of the agent's actions. An agent's history of past-actions supports the agent's learning from experience when presented with new tasks. Each agent's knowledge base consisted of three elements:

- 1. A strategist, or a decision-theoretic planner
- 2. Problem-specific context descriptor
- 3. A set of procedures or rules for domain-dependent actions

The strategist was responsible for planning and scheduling the actions that an agent must perform to achieve its goal. The problem context descriptor defined specific attributes of each request to ensure that each agent's attention was focused on the problem at hand. Problem solutions were modeled as domain-dependent procedures. The internal models maintained functions (such as managing access to the skills of each agent or maintaining its message buffer) that were private to each agent and not accessible to external agents.



Fig. 4.1. Architecture definition for a domain-independent knowledge-based (a.k.a. deliberative) agent in Agent-based Flight Operations Associate (AFLOAT)

The external models module maintained global functions that were accessible to other agents. Both types of models were used to maintain a set of actions necessary to achieve the agent's goals. The actions were stored as either rules or procedures. In AFLOAT, the strategist was implemented as a Strategy-Schema [172] that maintained each agent's subgoals for each request it received. The problem context descriptor was modeled as a Context-Schema to hold the features of a specific request, such as attention-focusing information, default knowledge, and standing orders. Request-specific procedures were modeled as Procedure-Schemas.

Behavioral Elements of an Agent in AFLOAT

Each agent in the AFLOAT architecture had six high-level behavior characteristics similar to Laufmann's "action-oriented" attributes [103]. The attributes or capabilities are:

- 1. Autonomy
- 2. Learning

- 3. Migration
- 4. Persistence
- 5. Communication
- 6. Cloning/spawning

Autonomy/semiautonomy is the ability of an agent to respond to a dynamic environment without human intervention, thereby improving the productivity of the user. When presented with a request, it can use its own strategy to decide how to satisfy the request. Each agent is capable of a type of learning that enables it to more responsively interact with its user community over a period of time. Learning also enables the agent to keep abreast of changes in its operational environment. The dynamic behavior of the agents is triggered by either command or event-driven stimuli. Migration is the ability of an agent to relocate to other nodes to accomplish its tasks. This ability can support load balancing, improve efficiencies of communication, and provide unique services that may not be available at a local node. Persistence is the ability to recover from environmental crashes and support time-extended activities, thereby reducing the need for constant polling of the agent's welfare by the user and providing better use of the system's communication bandwidth. A communication ability provides an agent with the mechanisms for supporting agent-agent and user-agent interactions either through an ACL or a domainspecific natural language. Spawning is an agent's ability to create other agents to support the parent agent, thereby promoting dynamic parallelism and faulttolerance. It is our opinion that these capabilities are necessary for building autonomous satellite control centers.

The agent architecture described above is generic enough for use in automating the operations of the control center of any spacecraft. The Explorer Platform's (EP's) [137] satellite control center was selected as a domain to test the feasibility of AFLOAT. Figure 4.2 depicts the interactions between different components of the EP satellite control center. The elements shown in the diagram are similar to those found in a typical satellite operations control facility. A taxonomy of the EP subsystems and data extraction system is also shown in Fig. 4.2. The diagram describes the interactions between the physical model (i.e., elements above the mnemonics database) and the logical model (i.e., elements below the mnemonics database) of the components of the EP system.

The main goal of the AFLOAT project was to implement a multiagent architecture that could interface directly with sources of data from a satellite and process and reason with the data to support autonomous operation of the control center. This was achievable with the aid of a data server agent that could interface directly with satellite telemetry and provide the information as mnemonics to other specialist agents. Due to operational restrictions, magnetic tapes were used to transfer satellite data to a workstation for processing by a data server agent. Even with that restriction, it was possible to demonstrate the feasibility of an automated agent-based satellite operations control center.



Fig. 4.2. A taxonomy of Explorer Platform (EP) satellite system and data extraction process

4.1.2 Implementation Architecture for AFLOAT Prototype

An architecture definition for AFLOAT is shown in Fig. 4.3. The multiagent system (MAS) employed direct communication between agents without a mediator. All external requests, from either a user or a remote client, entered the AFLOAT system through the Interface Services Agent (ISA), which then forwarded them to appropriate agents. The Systems Services Agent (SSA) maintained a database of agents' names, skills, and locations (i.e., TCP/IP socket address). It also monitored the health and status of each agent, and provided essential resources for agent migration.



Fig. 4.3. AFLOAT architecture definition

The architecture provided two classes of specialist agents – coordinator specialist and regular specialist. The coordinator specialist handled complex requests requiring the participation of three or more regular specialist agents to process. The coordinator agent had access to the system's global skillbase and possessed the capability to assemble a group of specialist agents, decompose the request into smaller tasks, and delegate the requests to them. Upon completing their tasks, the specialist agents returned their results to the coordinator, which in turn assembled them and freed up members of the agent group to return to their original states. Each specialist agent also had the capability to collaborate with other agents to process a task.

After accepting a request, a specialist agent examined it to determine whether it would require the services of other agents. If another agent's skill was required to support the task, it sent a message to the SSA requesting the location of the agent. After receiving a response from SSA, it formulated a request and sent a message to the other specialist agent. The other specialist agent would process the task and return a response (result) to the specialist agent. Each agent depicted in Fig. 4.3 existed in a C-language Integrated Production System (CLIPS) [46] process with a persistent socket connection to the SSA for monitoring its health and safety.

Approaches for Addressing Multiagent Architectural Issues in AFLOAT

To successfully develop AFLOAT as a MAS, the following four architectural issues had to be addressed:

• An approach was established for describing and decomposing the tasks that gave the coordinator specialist agents and regular specialist agents the capability to describe and decompose tasks.

75

- A format was defined for interaction and communication between agents that employed a semistructured message format for defining a language protocol for the agents.
- A strategy was formulated for distributing controls among agents. The control strategy initiated by either the coordinator agent or the regular specialist agent was driven by the requirements of the request or the task at hand. In certain situations, control was distributed among the agents, and in other situations, the coordinator agent assigned all tasks to a group of specialist agents in the form of a semicentralized control framework.
- A policy for coordinating the activities of agents was employed. Our design allowed the SSA to maintain a directory of the skills of all the agents and dispense information on the location of other agents upon request. In addition, SSA monitored the status of each agent and would reactivate, clone, or migrate them when necessary.

4.1.3 The Human Computer Interface in AFLOAT

A major component of AFLOAT was the User Interface Agent (UIA). The UIA was ultimately responsible for supporting dialogs and interactions with outside users of the agent system. In reality, the UIA was a community of agents, each with specific tasks. The UIA had a user agent that supported multimodal interaction techniques. As an example, the user could communicate with AFLOAT via typed text or spoken language. There was a Request Analysis Agent that checked for ambiguities in the user's request, checked spelling, filtered superfluous words, and performed pattern recognition and context-dependent analysis. A major component of the UIA was the User Modeling Agent. This agent was responsible for developing user profiles, classifying users, dynamically adapting to user behaviors and preferences, and resolving ambiguities. The Results Management Agent was responsible for interaction with the community of domain specialist agents, collection of results of work done by the domain specialist agents, integration of results, and notification of users. The UIA also supported a local request server that provided the UIA user-environment management, common services such as e-mail and printing, and results-display support.

AFLOAT also provided an operational domain-restricted natural language interface. The domain restriction is a requirement both for keeping the problem tractable and for performance reasons. This interface allowed users simply to make requests in sentence form. For the grammar component of the natural language interface, a "semantic" grammar was used. This can be defined as a grammar in which the syntax and semantics are collapsed into a single uniform framework [5]. This grammar looks like a context-free grammar except that it uses semantic categories for terminal symbols. There are several benefits of using a semantic grammar, the main one being that there is no need for separate processing of semantics. Semantic processing is done in parallel with syntactic processing. This also means that this method is very efficient, since no time is spent on a separate processing step. There is also the benefit of simplicity. It is not much harder to build a semantic grammar than a syntactical one. All requests either in the form of a natural language or structured queries were submitted to AFLOAT through one of the two interfaces labeled 1 in Fig. 4.3. All requests were automatically converted to an ACL to enable inter-agent communication and collaboration.

4.1.4 Inter-Agent Communications in AFLOAT

Inter-agent communication in AFLOAT was based on the following assumptions:

- 1. Agents in AFLOAT communicated through an asynchronous message passing mechanism (i.e., asynchronous input and output messages)
- 2. A common message structure was maintained for all agents
- 3. Communication between agents was achieved through TCP/IP-based sockets

The format of the ACL employed in AFLOAT was an enhanced version of that proposed by Steve Laufmann [103] for "coarse-grained" agents, to which we have added "performatives" proposed by Finin and group for the Knowledge Query Manipulation Language (KQML) [24, 81], and enhanced to meet specific requirements of the domain of spacecraft mission operations. The format for an ACL message in AFLOAT was as follows:

msg-id: A unique identifier composed of hostname, a random number, and system time separated by dashes (e.g., kong-gen854-13412.35).

user-id: The user-id of the person who originally submitted the request. sender: The name of the agent that the message is coming from.

- receiver: The name of the agent that the message is going to (in general). Asterisk (*) is used when we want the agent that is receiving the message to decide to whom the message should finally go.
- respond-to: The name of the agent to which the response to this request should go. It does not apply to response messages.
- reply-constraint: This is used for time constraints (e.g., ASAP, soon, whenever)
- language: The programming language that the message expects. This is especially important when the performative is "evaluate" and the string passed in the *input-string* slot is just evaluated. This would only apply to interpreted languages (e.g., CLIPS, shell scripts, etc.).
- msg-type: The message type (request, response, status, etc.).
- performative: The basic task to be performed (e.g., parse, archive, generate).
- recipient: The name of the user interface agent of the person to receive the results of the request.
- result-action: The type of action to invoke on the result message (display, print, etc.).

- domain: The domain is the system the message is dealing with. The domain in this case is usually "EP."
- object-type: The type of the main object that the message refers to. This is a level in the object hierarchy (e.g., system, subsystem, or component).

object-name: The actual name of the object given in the user request.

- object-specifiers: Any words from the user request that give more detailed information about which object is being referred to. This covers specification of a number of objects that are being referred to (e.g., "number 1," "this," "any").
- parameters: Any information that adds detail about the task to be performed. For example, when monitoring or reporting on the solar arrays, this field would specify temperature, performance, etc.
- action-start: The date/time at which the action being referred to in the message is to begin.
- action-duration: The length of time the action is to last.
- function: An actual function that is to be called as a result of the message. If the *input-string* field is not empty, it contains parameters that are to be passed to the function.
- input-string: This slot contains any data or text information required by other slots of the message. If the function slot is populated, this slot would contain input parameters. For a "parse" performative message, this slot would contain the actual sentence submitted by the user.

Each ACL message format had a message header and a message body. The message header consisted of the message attributes from *msg-id* through *result-format*. The rest of the message was the body of the message. An example of an ACL message generated from a user's natural language input is shown below. A user's input is the value stored in the *input-string* at the bottom of the message. Only pertinent values of the attributes of the message need to be included. Because the performative for this message is "parse," the SSA would use the information in its skill base to route the request to the natural language parser for processing. An example of a REQUEST-MESSAGE in an ACL format is the following:

```
(msg-id, gen001)
(user-id, john)
(sender, umbc-ui)
(receiver, loral-coord)
(respond-to, umbc-ui)
(reply-constraint, ASAP)
(language, CLIPS)
(msg-type, request)
(performative, parse)
(recipient, john-uia)
(result-action, nil)
(result-format, nil)
```

```
(domain, "EUVE spacecraft")
(object-type, nil)
(object-name, nil)
(object-specifiers, nil)
(parameters, nil)
(action-start, nil)
(action-duration, nil)
(function, nil)
(input-string, "monitor the health and safety of the
spacecraft's batteries")
```

Start-Up and Activation of a Community of Domain Specialist Agents in AFLOAT

To initiate the prototype, a UNIX process loaded the System Services Agent (SSA), labeled as box 3 in Fig. 4.3. The SSA then loaded each of the agents depicted in boxes 2–11 and established a persistent socket connection with each of them. The links were persistent to enable the SSA to monitor the status of the nine agents. In addition to monitoring the status of the agents, the SSA stored the location and skills base of other agents and provided appropriate resources to support the agents' migration requirements. To support fault tolerance, the ISA had the capability to monitor the status of the SSA, and to restart it if it died. Communications from clients (i.e., external interfaces – either user interfaces or remote clients) needed to be registered at the ISA. This was necessary to relieve the processing load of the SSA. The agents numbered 2 and higher constituted the AFLOAT server. All the agents communicated when necessary via TCP/IP sockets. At run time, the User Interface Agent (UIA) was loaded. Users submitted requests to AFLOAT from a remote web client and received responses/results locally.

4.2 Lights Out Ground Operations System

LOGOS [175, 177, 181, 183] was a proof-of-concept system that used a community of autonomous software agents that worked cooperatively to perform the functions previously undertaken by human operators who were using traditional software tools, such as orbit generators and command-sequence planners. The following discusses the LOGOS architecture and gives an example scenario to show the data flow and flow of control.

4.2.1 The LOGOS Architecture

For reference, an architecture of LOGOS is shown in Fig. 4.4. LOGOS was made up of ten agents, some of which interfaced with legacy software, some



Fig. 4.4. Lights out ground operations system (LOGOS) agent architecture

performed services for the other agents in the community, and others interfaced with an analyst or operator. All agents could communicate with any other agent in the community, though not all of the agents were required to communicate with other agents.

The System Monitoring and Management Agent (SysMMA) kept track of all of the agents in the community and provided addresses of agents for other agents requesting services, similar to the ISA in AFLOAT. Each agent when started had to register with SysMMA to register their capabilities and to obtain addresses of other agents whose services it needed.

The Fault Isolation and Resolution Expert (FIRE) agent resolved satellite anomalies. FIRE was notified of anomalies during a satellite pass. It contained a knowledge base of potential anomalies and a set of possible fixes for them. If it did not recognize an anomaly or was unable to resolve it, it then sent the anomaly to the user interface agent to be forwarded to a human analyst for resolution.

The User Interface Agent (UIFA) was the interface between the agent community and the graphical user interface that the analyst or operator used to interact with the LOGOS agent community. UIFA received notification of anomalies from the FIRE agent, handled login of users to the system, kept the user informed with reports, routed commands to be sent to the satellite, and performed other maintenance functions. If the attention of an analyst was needed but none was logged on, UIFA would send a request to the PAGER agent to page the required analyst.

The VisAGE Interface Agent (VIFA) interfaced with the Visual Analysis Graphical Environment (VisAGE) data visualization system. VisAGE was used to display spacecraft telemetry and agent log information. Real time telemetry information was displayed by VisAGE as it was downloaded during a satellite pass. VIFA requested the data from the Genie Interface Agent (GIFA) and Archive Interface Agent (AIFA) agents (see below). An analyst could also use VisAGE to visualize historical information to help monitor spacecraft health or to determine solutions to anomalies or other potential spacecraft problems.

The Pager Interface Agent (PAGER) was the agent community interface to the analyst's pager system. If an anomaly occurred or other situation arose that needed an analyst's attention, a request was sent to the PAGER agent, which then paged the analyst.

The Database Interface Agent (DBIFA) and the AIFA stored short-term and long-term data, respectively, and the Log Agent (LOG) stored agent logging data for debugging and monitoring purposes. The DBIFA stored information such as a list of the valid users and their passwords, and the AIFA stored telemetry data.

The GenSAA/GIFA interfaced with the GenSAA/Genie ground station software [65], which handled communications with the spacecraft. Gen-SAA/Genie was used to download telemetry data and maintain scheduling information and to upload commands to the spacecraft. As anomalies and other data were downloaded from the spacecraft, GIFA routed the data to other agents based on their requests for information.

The Mission Operations Planning and Scheduling System (MOPSS) Interface Agent (MIFA) interfaced with the MOPSS ground-station planning and scheduling software. MOPSS kept track of the satellite's orbit and when the next pass would occur and how long it would last. It also sent out updates to the satellite's schedule to requesting agents when the schedule changed.

4.2.2 An Example Scenario

An example scenario illustrating how the agents would communicate and cooperate would start with MIFA receiving data from the MOPSS scheduling software informing MIFA that the spacecraft would be in contact position in 2 min. MIFA would then send a message to the other agents to wake them up, if they were sleeping, and let them know of the upcoming event. The advance notice allowed them to do some preprocessing before the contact. When GIFA received the message from MIFA, it would send a message to the GenSSA Data Server to put it into the proper state to receive transmissions from the control center.

After receiving data, the GenSSA Data Server would send the satellite data to GIFA. GIFA had a set of rules that indicated which data to send

to which agents. As well as sending data to other agents, GIFA also sent all engineering data to the archive agent (AIFA) for storage, and sent trend information to the visualization agent (VIFA). Updated schedule information was sent to the scheduling agent (MIFA) and a report was sent to the user interface agent (UIFA) to send on to an analyst for monitoring purposes. If there were any anomalies, they were sent to the FIRE agent for resolution.

If there was an anomaly, the FIRE agent would try to fix it automatically by using a knowledge base containing possible anomalies and a set of possible resolutions for each anomaly. To fix an anomaly, FIRE would send a spacecraft command to GIFA to be forwarded on to the spacecraft. After exhausting its knowledge base, if FIRE was not able to fix the anomaly, it would forward the anomaly to the user interface agent, which then would page an analyst and display the anomaly on the analyst's computer for action. The analyst would then formulate a set of commands to send to the spacecraft to resolve the situation. The commands would then be sent to the FIRE agent so that it could add the new resolution to its knowledge base for future reference. The commands then would be sent to the GIFA agent, which in turn sent them to the GenSAA/Genie system for forwarding on to the spacecraft.

There were many other interactions between the agents and the legacy software that were not covered above. Examples include the DBIFA requesting user logon information from the database, the AIFA requesting archived telemetry information from the archive database to be sent to the visualization agent, and the pager agent sending paging information to the paging system to alert an analyst of an anomaly needing his or her attention.

4.3 Agent Concept Testbed

The motivation behind ACT was to develop a more flexible architecture than LOGOS for implementation of a wide range of intelligent or reactive agents. After developing the architecture, sample agents were built to simulate ground control of a satellite constellation mission as a proof of concept. The following discusses the ACT agent architecture and gives an operational scenario using the satellite constellation proof of concept.

4.3.1 Overview of the ACT Agent Architecture

The ACT architecture was a component-based architecture that allowed greater flexibility to the agent designer. A simple agent could be designed by using a minimum number of components that would receive percepts (inputs) from the environment and react relative to those percepts. This type of simple agent would be a *reactive agent*.

A robust agent could be designed using more complex components that allowed the agent to reason in a deliberative, reflexive, and/or social fashion.



Fig. 4.5. Agent Concept Testbed (ACT) agent architecture

This robust agent would maintain models of itself, other agents in its environment, objects in the environment that pertain to its domain of interest, and external resources that it might utilize in accomplishing a goal. Figure 4.5 depicts the components for a robust agent. The depicted components gave the agent a higher degree of intelligence when interacting with its environment.

The ACT agent architecture was capable of several types of behaviors. Basically, "agent behavior" refers to the manner in which an agent responds to some sort of stimulus, generated either externally (outside the agent) or internally (within the agent). We have identified four basic classes of behaviors that agents can realize. These are:

- Social: Social behaviors refer to behaviors shared between/among agents. The ACT architecture supported two types of social behavior: social behavior triggered by another agent and social behavior triggered by the agent itself. In each of these cases, the agent utilized ACL messages to solicit help or to coordinate the behaviors of other agents.
- Proactive: This type of behavior is stimulated in some way by the agent itself. For our agents, there was one type of proactive behavior that was to be supported, self motivating. Self-motivating behaviors are triggered by built-in or intrinsic goals.
- Reactive: Reactive behaviors are those that require "no thinking." These behaviors are like built-in reflexive actions that are triggered by events in

the agent's environment or by other agents. When detected, the agent responds immediately with a predetermined action.

Deliberative: This type of behavior is perhaps the most difficult and interesting. At the highest level of abstraction, this type of behavior involves the establishing of a hierarchy of goals and subgoals, the development of plans to achieve the subgoals, and the execution of the planned steps to ultimately accomplish the goal that started the process of deliberation in the first place.

4.3.2 Architecture Components

Components

A component in the agent architecture is a software module that performs a defined task. Components when combined with other software components can constitute a more robust piece of software that is easily maintained and upgraded. Each component in the architecture can communicate information to/from all other components as needed through various mechanisms including a publish-and-subscribe communication mechanism, message passing, or a request for immediate data.

Components may be implemented with a degree of intelligence through the addition of reasoning and learning functions. Each component needs to implement certain interfaces and contain certain properties. Components must implement functionality to publish information, subscribe to information, and accept queries for information from other components or external resources being used by the component. Components need to keep track of their state and to know what types of information they contain and what they need from external components and objects.

The following describes the components in the ACT agent architecture.

Modeler

The modeling component was responsible for maintaining the domain model of an agent, which included models of the environment, other agents in the community, and the agent itself. The Modeler received data from the Perceptors and agent communication component. These data were used to update state information in its model. If the data caused a change to a state variable, the Modeler then published this information to other components in the agent that subscribed to updates to that state variable. The Modeler was also responsible for reasoning with the models to act proactively and reactively with the environment and events that affected the model's state.

The modeler could also handle what-if questions. These questions would primarily originate from the planning and scheduling component, but could also come from other agents or from a person who wanted to know what the agent would do in a given situation or how a change to its environment would effect the values in its model.

Reasoner

The Agent Reasoner made decisions and formulated goals for the agent component through reasoning with received ACL messages and information in its local knowledge base, as well as with model and state information from the Modeler. This component reasoned with state and model data to determine whether any actions needed to be performed by the agent to affect its environment, change its state, perform housekeeping tasks, or perform other general activities. The Reasoner would also interpret and reason with agent-to-agent messages received by the agent's communications component. When action was necessary for the agent, the Reasoner would produce goals for the agent to achieve.

Planner/Scheduler

The planner/scheduler component was responsible for any agent-level planning and scheduling. The planning component formulated a plan for the agent to achieve the desired goals. The planning component was given a goal or set of goals to fulfill in the form of a plan request. This typically came from the Reasoner component, but could be generated by any component in the system.

At the time that the plan request was given, the planning and scheduling component acquired a state of the agent and system, usually the current state, as well as the set of actions that could be performed by this agent. This information would typically be acquired from the modeling and state component. The planning and scheduling component then generated a plan as a directed graph of steps. A step is composed of preconditions to check, the action to perform, and the expected results from the action (post condition). When each step was created, it was passed to Domain Expert components/objects for verification of correctness. If a step was deemed incorrect or dangerous, the Domain Expert could provide an alternative step, solution, or data to be considered by the planner.

Once the plan was completed, it was passed back to the component that requested the plan (usually the Reasoner). The requesting component then either passed it on to the Agenda to be executed or used it for planning/whatif purposes.

Agenda/Executive

The Execution component managed the execution of steps and determined the success or failure of each step's execution. Output produced during a step's execution could be passed to an Effector or the Reasoning component. The Agenda and Executive worked together to execute the plans developed by the Planner/Scheduler. The agenda typically received a plan from the Reasoner, though it could receive a plan from another component that was acting in a reactive mode. The agenda interacted with the Execution component to send

the plan's steps in order for execution. The agenda kept track of which steps were being executed, had finished executing, were idle, or were waiting for execution. It updated the status of each step appropriately as the step moved through the execution cycle. The agenda reported the plan's final completion status to the Planner and Agent Reasoner when the plan was complete.

The Executive would execute the steps it received from the Agenda. A step contained preconditions, an action, and possible postconditions. If the preconditions were met, the action was executed. When executions finished, the postconditions were evaluated, and a completion status was generated for that step. The completion status was returned to the agenda, which allowed for overall plan evaluation.

The execution component interacted with the agenda in the following way. The agenda sent the first step to the execution component. This woke up the Executive. The component then began executing that step. The Executive then checked to see if another step was ready for execution. If not, the component would go back to sleep until it received another step from the agenda.

The Modeling component would record state changes caused by a step execution. When a plan was finished executing, the Agenda component sent a completion status to the Reasoning component to indicate that the goal established by the Reasoner had been accomplished. If the Agent Reasoning component was dealing with data from the environment, it could decide either to set a goal (for more deliberative planning) or to react quickly in an emergency situation. The Reasoner could also carry on a dialog with another agent in the community through the Agent Communication Perceptor/Effector.

A watch was also attached to the Executive. It monitored given conditions during execution of a set of steps and the consequence if the condition occurred. Watches allowed the planner to flag things that had to be particularly looked out for during real-time execution. They could be used to provide "interrupt" capabilities within the plan. An example of a watch might be to monitor drift from a guide star during an observation. If the drift exceeds a threshold, the observation is halted. In such a case, the watch would notify the Executive, which in turn would notify the Agenda. The Agenda would then inform the Reasoner that the plan had failed and the goal was not achieved. The Reasoner would then formulate another goal (e.g., recalibrate the star tracker).

Agent Communications

The agent communication component was responsible for sending and receiving messages to/from other agents. The component took an agent data object that needed to be transmitted to another agent and converted it to a message format understood by the receiving agent. The message format that was used was based on Foundations of Intelligent Physical Agents (FIPA) [110]. The message was then transmitted to the appropriate agent through the use of a NASA-developed agent messaging protocol/software called Workplace [7]. The reverse process was performed for an incoming message. The communications component took the message and converted it to an internal agent object and sent it out to the other components that had a subscription to incoming agent messages. The communications component could also have reactive behavior, where for a limited number of circumstances, it could produce an immediate response to a message.

Perceptors/Effectors

Percepts received through sensors, communication with external software/systems, and other environmental entities were received through a Perceptor component. These percepts were passed from the Perceptor to the Modeling component, where a model's state was updated as needed.

The Perceptors were responsible for monitoring parts of the environment for the agent. An example might be a subsystem of a spacecraft or recurring input from a user. Other than agent-to-agent messages, any data received by the agent from the environment entered through Perceptors. An agent might have zero or more Perceptors, where each Perceptor received information from specific parts of the agent's environment. A Perceptor could just receive data and pass it on to another component in the agent, or it might perform some simple filtering/conversion before passing it on. A Perceptor might also act intelligently through the use of reasoning systems. If an agent was not monitoring a part of the environment, then it would not have any perceptors (an example of this would be an agent that only provides expertise to other agents in a certain area, such as fault resolution).

The Effector was responsible for effecting or sending output to the agent's environment. Any agent output data, other than agent-to-agent messages, left through Effectors. Typically the data coming from the Effectors would be sent from the executive that had just executed a command to the agent's environment. There could be zero or more Effectors, where each Effector sent data to specific parts of the agent's environment. An Effector could perform data conversions when necessary and could even act intelligently and in a proactive manner when necessary through the use of internal reasoning systems. As with the Perceptors, an agent might not have an Effector if it did not need the capability of interacting with the environment.

Agent Framework

A software *framework*, into which the components were "plugged," provided a base functionality for the components as well as the inter-component communication functionality. The framework allowed components to easily be added and removed from the agent while providing for a standard communications interface and functionality across all components. This made developing and adding new components easier and made component addition transparent to existing components in the agent.

The communications mechanism for components was based on a publishand-subscribe model, with a direct link between components when there was a large amount of data to be transferred. Components communicated to each other the types of data that they produced when queried. When one component needed to be informed of new or changed data in another component, it identified the data of interest and subscribed to it in the source component. Data could be subscribed-to whenever they are changed or on an as-needed basis. With this mechanism, a component could be added or removed without having to modify the other components in the agent.

4.3.3 Dataflow Between Components

This section gives an example of how data flowed between components of the architecture. In this example scenario, a spacecraft's battery is discharging. Figure 4.6 shows a timeline and the flow of data between components. The following is the scenario:

1. The agent detects a low voltage by reading data from the battery via a Perceptor. The Perceptor then passes the voltage value to the Modeler, which has subscribed to the Perceptor to receive all percepts.



Fig. 4.6. Scenario of data flowing between agent components

- 2. When the Modeler receives the voltage from the Perceptor, it converts the voltage to a discrete value and updates this value in the model. In this case, the updated voltage value puts it below the acceptable threshold and changes the model's voltage state to "low." This change in state value causes a state change event and the Modeler now publishes the new state value to all components that have subscribed to changes in this state variable. Since the Reasoner has subscribed to changes in this state variable, the low voltage value is sent to the Reasoner.
- 3. In the Reasoner, the low voltage value fires a rule in the expert system. This rule calls a method that sends the Planner/Scheduler a goal to achieve a battery voltage level that corresponds to fully charged.
- 4. When the Planner/Scheduler receives the goal from the Reasoner, it queries the Modeler for the current state of the satellite and a set of actions that can be performed (this set may change based on the health of the satellite).
- 5. After receiving the current state of the satellite and the set of available actions from the Modeler, the Planner/Scheduler formulates a list of actions that need to take place to charge the battery. It then sends the plan back to the Reasoner for validation.
- 6. The Reasoner examines the set of actions received from the Planner/ Scheduler and decides that it is reasonable. The plans are then sent to the Agenda.
- 7. The Agenda then puts the action steps from the plan into a queue for the Executive.
- 8. As the Executive is ready to execute a new step, the agenda passes the plan steps one at a time to the Executive for execution.
- 9. The Executive executes each action until the plan is finished. At this time, the Executive notifies the Agenda that it has finished executing the plan.
- 10. The Agenda marks the plan as finished and notifies the Reasoner (or whomever sent the plan) that the plan finished successfully.
- 11. After the plan is executed, the voltage starts to rise and will trigger a state change in the Modeler when the voltage goes back into the fully charged state. At this time, the Reasoner is again notified that a change in a state variable has occurred.
- 12. The Reasoner then notes that the voltage has been restored to the fully charged state and marks the goal as accomplished.

4.3.4 ACT Operational Scenario

The operational scenario that was developed to evaluate ACT was loosely based on certain nanosatellite constellation ideas. Figure 4.7 graphically illustrates this scenario. It was based on the idea of a ground-based community of proxy agents (each representing a spacecraft in the nanosatellite constellation) that provided autonomous operations of the constellation. Another



Fig. 4.7. Agent community developed in ACT to test the new agent architecture and community concepts

scenario corresponded to the migration of this community of proxy agents to the spacecraft themselves, to support an evaluation of space-based autonomy concepts.

In this scenario, several nanosatellites are in orbit collecting magnetosphere data. The Mission Operations Control Center (MOCC) makes contact with selected spacecraft according to its planned schedule when the spacecraft (S/C) come into view.

The agents that would make up the MOCC would be:

- Mission Manager Agent: Coordinates the agent community in the MOCC, manages mission goals, and coordinates CMAs.
- Contact Manager Agent (CMA): Coordinates ground station activities (one agent per ground station), communicates with the spacecraft, and sends and receives data, commands, and telemetry.
- User interface: Interfaces with the user to accept commands for the spacecraft and sends data to be displayed.
- MOCC Planning/Scheduling Agent: Plans and schedules contacts with the spacecraft via interface with external planner/scheduler.
- Spacecraft Proxy Agents: There is one proxy agent for each spacecraft in orbit. The agents keep track of spacecraft status, health and safety, etc. The agents will notify the Mission Manager Agent when there occurs an anomaly that may need handling.

Each of the above agents registers with the Ground Control Center (GCC) manager agent. The GCC manager agent notifies the agents when there is an impending contact for their spacecraft, and when another agent is going to

be added to the community; it also provides to the agents the address of the other agents (so agents can pass messages to each other). The following is a spacecraft contact scenario that illustrates how the agents worked with the GCC manager agent:

- Agents register with the GCC Manager Agent at system startup.
- The GCC Planner/Scheduler Agent communicates with the spacecraft Proxy Agents to obtain spacecraft communications-view data. It then creates a contact schedule for all orbiting spacecraft.
- The GCC Manager Agent receives the schedule from the GCC Planner/ Scheduler Agent.
- The GCC Manager Agent informs the CMA about the next contact (when and with which spacecraft).
- The CMA receives notification of an acquisition of signal (AOS) from a spacecraft. The MOCC is now in contact with the spacecraft.
- The CMA executes the contact schedule to download data, delete data, or save data for a future pass.
- The CMA analyzes the downloaded telemetry data. If the telemetry indicates a problem, the CMA may alter the current contact schedule to deal with the problem.
- The CMA performs any necessary commanding in parallel with any data downloads.
- The CMA sends the telemetry to the appropriate spacecraft Proxy Agent for processing.
- The spacecraft Proxy Agent processes the telemetry data and updates the state of its model of the spacecraft from the telemetry received.
- If the spacecraft Proxy determines that a problem exists with the spacecraft and an extended or extra contact is needed, a message is sent to the GCC Planner/Scheduler Agent which will re-plan its contact schedule and redistribute it to the GCC Manager.
- The spacecraft Proxy Agent sends to the Contact Manager any commands that need to be uploaded.
- The Mission Manager Agent ends contact when scheduled.

4.3.5 Verification and Correctness

Whereas AFLOAT and LOGOS demonstrated that typical control center activities could be emulated by a multiagent system, the major objective of the ACT project was to demonstrate that ground-based surrogate agents, each representing a spacecraft in a group of spacecraft, could control the overall dynamic behaviors of the group of spacecraft in the realization of some global objective. The ultimate objective of ACT was to help in the understanding of the idea of progressive autonomy (see Sect. 9.6), which would, as a final goal, allow the surrogate agents to migrate to their respective spacecraft and then allow the group of autonomous spacecraft to have control of their dynamic behaviors independent of relying on ground control. ACT properly emulated the correct interaction between surrogate agents and their respective spacecraft. This "correctness" was determined by comparison between what the surrogate did vs. what a human controller on the ground would have done, in conjunction with what the controllers associated with the other surrogates would have done, to achieve a global objective. This analysis was undertaken more at the heuristic level than at a formal level. The design of the surrogates was realized in a modular fashion in order to support the concept of incremental placement of the functional capabilities of the surrogate agent in the respective spacecraft, until the spacecraft itself was truly agent-based and "autonomous." This particular aspect of the ACT project was heuristically realized, but not rigorously (formally) tested out.

The use of formal methods has been identified as a means of dealing with this complex problem. Formal approaches were previously used in the specification and verification of the LOGOS system [56,118,119,124]. A formal specification in Communicating Sequential Processes (CSP) highlighted a number of errors and omissions in the system. These, and other, errors were also found by an automated tool [57–59, 111, 112], which implemented an approach to requirements-based-programming [52]. For more information on formal verification of agent-based systems, see [123].

Technology

Core Technologies for Developing Autonomous and Autonomic Systems

This chapter examines the core artificial intelligence technologies that will make autonomous, autonomic spacecraft missions possible. Figure 5.1 is a pictorial overview of the technologies that will be discussed. The *plan* technologies will be discussed first followed by the *act* and *perceive* technologies, and finally technologies appropriate for testing.

It is difficult to make definitive statements on the functionality, strengths, and weaknesses of software systems in general, since designers have tremendous latitude on what they do. This chapter explains and discusses the attributes seen in a majority of the systems described. It should be understood that exceptions may exist.

5.1 Plan Technologies

The planning portion of the autonomy cycle is responsible for examining the environment and choosing appropriate actions in light of the goals and mission of the system. Sometimes this choice requires interactions with other systems. Planners are a central technology in all computerized planning, and many techniques have been developed to support planning, such as formal collaboration languages, evidential reasoning, and learning techniques. The rest of this section will discuss these techniques and planning in general.

5.1.1 Planner Overview

A defining characteristic of an autonomous system is the ability to independently select appropriate actions to achieve desired objectives. Planner systems are the software component commonly used to achieve this capability. Work on software planners goes back to 1959 and, over the intervening years, many types of planners have been developed.

Figure 5.2 shows a high level view of a planner and its context in a system architecture. All planners begin with a set of initial mission objectives that are



Fig. 5.1. Collaborative autonomy technologies



Fig. 5.2. Planner architectures

specified as goals to the planner. These goals are analyzed in the light of the planner's view of the environment and a database describing what it is capable of doing. After analysis, the planner chooses a set of actions to perform and these actions are sent off for execution. Results from the execution of these actions are fed back to the planner to update its view of the environment. If the goals have been achieved, success is reported back to the higher level system. If some problem has occurred during execution, error recovery is attempted, a new plan is created, and the cycle repeats. If no other plan can be created, the failure is reported to the higher level system.
This description is generic and it leaves many design decisions unanswered. For example, planners differ in how they describe the goals they are to achieve, the environment they execute in, and their database of potential actions. Each of these descriptions is given in a computer-based language. This language is very important since it defines what the system is capable of doing and how it will do it. It also defines what the system cannot do. If the language used is too limited, it may not be able to describe some aspect of the domain, potentially limiting the planner's ability to handle some situations.

Planners differ in the speed at which they come to decisions. Some planners are slow and deliberative, while others are quick and reactive. In general, the slow deliberative planners make plans that are more globally optimal and strategic in nature. The reactive planners tend to examine the environment and choose from a highly constrained set of plans. They are tactical in nature and work well in rapidly changing environments where the time for slow careful choice is not available. In many real world systems, either the planner is designed to handle both deliberation and reactivity, or two separate planners are integrated together.

All robust planners must deal with the failure of an action during execution. Some planners have low level strategies on hand and when a failure occurs, they immediately attempt to repair the plan. Others have a database of alternative ways to achieve an objective, and when an attempt fails, they analyze the current environment and choose another plan. In domains where the environment changes very rapidly relative to the planner decision time or where action failure is a regular occurrence, the planner may take the possibility of failure into consideration during plan creation. These systems give preference to robust plans that can help recover from likely failures even if the robust plan has a higher cost in resources than the alternatives.

Some planners convert higher level tasks into low level actions just before execution. By using this strategy, they commit fewer resources to any one plan and can quickly react to changes in the environment. Unfortunately, the plan that is ultimately executed will often be suboptimal, particularly if two or more tasks are competing for the same resource. Other planners map the toplevel goals into a complete series of small actions that take place in a timesequenced manner. The advantage of this approach is that a more globally optimal plan can be created. Its disadvantage is that a failure occurring in one step can cause the rest of the plan to be abandoned and re-planned. This is computationally expensive and time consuming. Re-planning can also cause problems if the domain looks forward into the plan and begins the commitment of resources based on the expected plan. In these domains, the replanning step must use repair strategies in an attempt to maintain most of the original plan.

Planners must be sensitive to an action's cost in resources. In computer domains, such as software agents, actions have small costs and plan selection can usually ignore resource issues. In other domains, like spacecraft, some actions commit resources that cannot be replenished (such as propellant). When resources have a very high cost, a failure during an action can threaten the whole mission. For planning in this type of domain, the costs and recovery strategies must be carefully chosen before action is taken and resources committed.

Having reviewed planner technology and a number of design choices facing a planner, we now will describe several common planner technologies.

5.1.2 Symbolic Planners

Symbolic planners are systems that represent their goals and plans as a series of symbolic assertions, instead of numbers, fuzzy quantities, or probabilities. Symbolic planners have been used in many domains.

Figure 5.3 depicts a symbolic planner. The plans of a symbolic planner are stored in a centralized database. Each plan has a set of preconditions and a list of operations to perform. The planner uses the preconditions to determine when it can use a plan. These preconditions can specify environmental constraints, the availability of resources, and potentially whether other plans have been executed before this plan. The list of operations in the plan may be primitive operations to perform, or they can be additional plan components. These plan components become additional subgoals that need to be examined by the planner. Often a plan instantiated during the planning cycle is called a *task*.



Fig. 5.3. Symbolic planner

While symbolic planners vary in detail, they generally start the planning activity by examining the goals in light of the current environment and then break the goals into a series of tasks. These tasks are themselves examined and broken into simpler subtasks. This iterative refinement process continues until a point is reached where the tasks define a series of steps at an appropriate level of abstraction for plan implementation. The plan is executed and feedback is generated on the success of the tasks. The higher level system is signaled when objectives have been met. If a failure occurs in one or more of the tasks, the planner modifies its plan and the cycle is repeated. If the planner exhausts all of its options and still the objectives have not been met, the planner signals to the higher level that it has failed.

Symbolic planners use many different strategies for choosing among the potential plans. They often spend large amounts of computer resources generating plans and selecting the best ones. They have difficulties in situations where the decision cycle time is short, or where actions and failures are not deterministic.

5.1.3 Reactive Planners

Reactive planners are specifically designed to make rapid choices in time critical situations. They can be designed around either symbolic or numeric representations.

Figure 5.4 shows the structure of a reactive planner. Reactive planners begin by evaluating the available plans in light of the current context and then choosing the most appropriate. These plans are simple in nature and are designed to execute immediately. Once the plan is selected and being executed, the planner monitors the current situation and only changes the plan if the



Fig. 5.4. Reactive planner

goal is modified or the plan succeeds or fails. A reactive planner can quickly assess the updated situation and switch to a new plan. The selection process is kept fast by limiting the number of potential plans that have to be examined. Sometimes this is accomplished by the higher-level control system, explicitly listing all plans the reactive planner needs to examine. On other systems, the set of plans is automatically constrained by indexing the plans on their applicable context. At each decision step, only the context-appropriate plans will have to be examined.

Reactive planners spend little time choosing the next plan, and therefore, are appropriate for time critical situations. They are often used in the low-level control of robot and robot-like systems, which have hard real-time deadlines. Because reactive planners only examine a limited number of choices, they will often make good local decisions and poor global ones. Another component of the overall system architecture is usually made responsible for optimizing the global objectives.

5.1.4 Model-Based Planners

Model-based planners use models to analyze the current situation and create their plans. They represent this information symbolically with goals described as desired changes in the environment.

Figure 5.5 shows the structure of a model-based planner. These systems give careful focus to the process of updating their internal models. As part of world-model update, these planners may create detailed submodels of the internal status of sensors and actuators, and using their model, can extend the direct measurements to determine the status of internal, but not directly measured, subsystems.

Once an accurate model of the platform and world has been created, the model-based planner examines the goals, and using its models, creates a plan



Fig. 5.5. Model-based planning

to achieve the objectives. Model-based systems can be extremely effective at creating plans in the face of one or more damaged subsystems.

One can argue that many planners are model-based since their database of potential actions implicitly defines the system and its capabilities. While there is some merit to this argument, it misses the fact that this implicit knowledge is usually generated by human programmers and may not be complete. These systems are, therefore, artificially limited in what they can plan. A model-based planner, using its deep understanding of the system it controls, can come up with novel approaches to achieve the objectives. This is most important when the system is damaged. Model-based systems can use their model to work around the damage. In other approaches, the human programming must explicitly define the strategy to use when failure occurs, and if they did not make a plan for the failure, the system will respond suboptimally.

The advantages of model-based approaches are that they only need the model description of the system being controlled and the rest is automatically determined. Their disadvantage is that reasoning on models can be very slow and the necessary models are often hard to construct and possibly incomplete. Model-based approaches are often used in support of other forms of planners.

5.1.5 Case-Based Planners

Case-based planners are systems that represent their planning knowledge as a database of previously solved planning cases. Case-based planners exploit the idea that the best way to solve a new problem is probably very similar to an old strategy that worked. Their cases can be a mixture of symbolic and numeric information. Case-based planners are a class of case-based reasoning (CBR) technology [1,116,138].

Figure 5.6 shows the structure of a case-based planner. When a casebased planner is given a new goal, it attempts to find a similar case in its case database. The cases are indexed on the goals being solved and the environmental conditions when they were solved. If similar cases are found, they are extracted, adapted to the current situation, and analyzed to see whether they will achieve the goal being pursued. Often the case will be simulated in the current environment to determine its actual behavior. This simulation can discover defects in the plan so that repair strategies can be applied to customize the plan for the current situation. The new case is simulated and the cycle repeated. Once a plan has met all the necessary requirements, it is executed. The results of execution are examined, and if it did not achieve the objective, new cases are retrieved and the cycle is repeated. Cases that are successful may be placed in the database for future planning activity. In this way the case-based planner learns and becomes more proficient.

Two difficulties that arise in case-based planners pertain to the methods used to index the cases and the reasoning component of the system. How cases are indexed determines whether appropriate cases will be found when they are needed to solve a problem. If the indexing is too restrictive, an appropriate



Fig. 5.6. Case-based planning

case may not be retrieved for examination and the knowledge it represents may be lost. The system will have to start the planning with a less optimal case. If the indexing is too loose, a large number of inappropriate cases will be supplied and each will have to be examined and eliminated. This makes the resulting system slower and less responsive. How to index the cases is a central problem in case-based systems, and in a real sense, can determine its success.

The second difficulty is the reasoning component of a case-based planner. This component is responsible for determining whether the plan will work and repairing it if the plan can be repaired. Ultimately, if no plan is similar to the current situation, the reasoning component must create a whole new plan. These are difficult components to construct.

Many case-based planners use a conventional symbolic or a model-based planner as the reasoning component of the case-based planner. The advantage to this approach is that the resulting system has all of the capabilities of the conventional symbolic or a model-based planner with a case-based planner's ability to learn new cases. In effect, the case-based database generated is used to cache the knowledge of the conventional system. This makes the resulting system more responsive over time.

5.1.6 Schedulers

Planning systems and scheduling systems work in the same domain, solve very similar problems, and use the same core technologies. The primary difference is that scheduling systems are more focused on detailed accounting of resources, and they tend to generate complete and detailed plans.

5.2 Collaborative Languages

Most forms of collaboration require the collaborating agents to communicate using a shared language. This language specifies the kinds of information that can be exchanged between the agents, and in a real way, determines what one agent can and cannot communicate to another. Well designed languages allow the agents to say what is necessary to solve the problem. Poor languages limit communication and support less than optimal solutions.

Computer systems support many kinds of communication. The simplest communication, from a computer science point of view, is an exchange of messages whose content is an internal program data structure. The data structure can be a simple primitive type like a number or string, or the structure can be complex connections of records. This approach to communication is easy to implement, and the communicating parties can immediately interpret the content and meaning of the messages they exchange. However, this method of communication has a limited expressive capability.

The other approach is to create a general purpose formal language. These systems construct a computer language that is general in nature and able to express a wide range of concepts. The first advantage of this approach is flexibility, since these languages can express and collaborate on richer sets of problems. Also, being a formal language, it can be documented and used by different and disjoint teams building collaborative agents. Even though they share no code or common ancestry, they can collaborate using this common language. The Defense Advanced Research Projects Agency (DARPA) knowledge sharing effort (KSE) has developed a very capable formal language for intelligent systems to exchange information. In practice, both types of communication are used in collaborative agents.

5.3 Reasoning with Partial Information

Intelligent behavior can be loosely broken down into problem solving (planning) and reasoning on evidence. Several technologies are used to reason on partial information. This section will describe two common techniques: fuzzy logic and Bayesian reasoning.

5.3.1 Fuzzy Logic

Fuzzy Logic was developed by Lotfi Zadeh for use in control systems that are not easily converted into mathematical models. He wrote a paper in the 1960s called "Fuzzy Logic," which described this new technology. Acceptance in the United States was initially slow and may have been hindered by the word "fuzzy" in its name. Despite its slow start, it is now being used in a wide range of control systems.

Fuzzy Logic works well in domains that are complex and where knowledge is contained in engineering experience and rules of thumb, rather than mathematical models. Fuzzy Logic allows the creation of set-membership functions that support reasoning about a particular value's degree of membership in that set. Figure 5.7 shows three fuzzy sets that define Cold, Warm, and Hot. A specific temperature, say 100°, has different membership in each of these sets. Combiners are used to connect fuzzy expressions together to generate compound expressions. Rules can be constructed that reason on these memberships and perform some action. For example, "if EngineTemp is Hot and the DeltaV is NecessaryDeltaV then Stop Burn" would reason on the current engine temperature in relationship to the necessary delta V and determine whether it should terminate burn. Other rules could deal with a hot engine and an insufficient delta V, or could stop the engine when the delta V is optimal.

Despite their name, Fuzzy Logic systems are time invariant, deterministic, and nonlinear. They are computationally efficient. Once rules have been developed, they can be "compiled" to run in limited computational environments. Unfortunately, how and when to use fuzzy combination rules are not well understood or systematic.

Engineers might be drawn to this technique because it is easy to get an initial prototype running and incrementally add features. If the domain is primarily engineering rules of thumb, the use of this technique is defensible. If, however, good mathematical or other models of the domain can be



Fig. 5.7. Fuzzy sets

constructed, even at the cost of more up-front effort, the resulting system should be more robust. Serious consideration should be given to whether the initial ease of fuzzy construction outweighs the robustness of more formal modeling techniques.

5.3.2 Bayesian Reasoning

Bayesian reasoning is a family of techniques based on Bayesian statistics. Its strength is the firm foundation of statistics on which it rests (hundreds of years). It has a well-developed methodology for mapping real world problems into statistical formulations, and if the causal model is accurate and the evidence accurately represented, Bayesian systems give scientifically defensible results.

In simple terms, Bayes' rule (or Bayes' theorem) states that the belief one accords a hypothesis upon obtaining new evidence can be computed by multiplying one's prior belief in the hypothesis and the likelihood that the evidence would appear if the hypothesis were true. This rule can be used to construct very powerful inferential systems.

Bayesian systems require causal models of the world. These models can be developed by engineers using their knowledge of the system, or in some cases, by analyzing empirical data. These models must be at the appropriate level of detail for the system to make correct inferences. Multiple implementation strategies have been developed to reason with Bayesian models once they have been created. Figure 5.8 shows two methods for representing a Bayesian model. Bayes nets encode the state information as nodes and causality as links between the nodes. Figure 5.8a shows a simple example. More recent work has developed systems that can reason on the probabilistic equations in a symbolical manner similar to Mathematica. An example data set can be seen in Fig. 5.8b.

Many implementations of Bayesian systems make a simplifying assumption that the supplied evidence is completely independent from other



Fig. 5.8. Bayes network (a) and corresponding symbolic representation (b)

evidence. This makes Bayesian model generation and analysis much simpler. Unfortunately, the independence rule is often violated in real domains and distorts the final results. This effect has to be carefully controlled.

5.4 Learning Technologies

Learning techniques allow systems to adapt to changing circumstances such as new environments or failures in hardware. They also allow systems to respond more rapidly to situations they have previously explored and to which they have found good solutions. While many techniques exist, this section will focus on artificial neural networks and genetic algorithms.

5.4.1 Artificial Neural Networks

Artificial neural networks are a learning technique based loosely on biological neural networks. Figure 5.9 is a pictorial representation of a neural net. Each "neuron" is a simple mathematical model that maps its inputs to its outputs. It is shown in Fig. 5.9 as a circle. The input situation is represented as a series of numbers called the input vector that is connected to the first layer of neurons. This layer is usually connected to additional layers finally connecting to a layer that produces the output vector. The output vector represents the solution to the problem described by the input vector.

Neural networks are trained using a set of input vectors matched to output vectors. By iterating through this training set, the network is "taught" how to map this and similar input vectors into appropriate output vectors. During this training, the network determines what features of the input set are important and which can be ignored. If the training set is complete and the network can be trained, the resulting system will be robust.



Fig. 5.9. Artificial neural network

Neural networks have been used successfully in many domains. However, one difficulty with neural networks is that it is not always easy to understand what they have learned since the learned information is locked up in internal neuron states. A related difficulty is that there is no way to determine whether the training set was complete enough to have the network learn the correct lessons. When faced with a novel situation, the network may choose the wrong answer. The final difficulty is that controlled adaptation can be good, but too much learning can cause a previously stable system to fail.

5.4.2 Genetic Algorithms and Programming

Genetic Algorithms are learning techniques based loosely on genetic reproduction and survival of the fittest. They use either numeric or symbolic representations for the definition of the problem.

Figure 5.10 is a pictorial representation of a genetic algorithm reproduction cycle. Processing begins by randomly creating the initial pool of genes. Each gene represents one potential solution and each is evaluated to determine how well it solves the problem. The best solutions from the current generation are paired and used to make new genes for the next generation. The bits from the two parents' genes are mixed to create two new genes. The new genes enter the gene pool to be evaluated on the next cycle. Since the best genes are continually being mixed together, the genes in later generations tend to be better at solving the problem. This cycle continues until a strong solution is found.

Genetic programs operate in a manner similar to genetic algorithms with the modification that each gene is actually a small program. Each reproduction cycle mixes parts of each parent program.



Fig. 5.10. Genetic reproduction cycle

5.5 Act Technologies

The action portion of the autonomy cycle is responsible for implementing the choices made by the plan portion of the cycle. The actions interact directly with the operating environment and must be customized to the problem domain.

The most obvious action technologies are the actuators used by robots and immobot agents. Actuators are devices that are directly coupled in the operation environment and are able to make some change when activated. Examples include opening a valve, moving a drive train, or starting a pump.

Being physically connected to the real world, actuators are subject to the wear and tear due to age and potential damage caused by the environment. Wear and tear can cause actuators to have complex failure pathologies, and robust systems are designed to detect these failures and recover. Since actuators modify the operating environment, designers must ensure that actuator actions do not have negative side effects for the agent or the environment.

Communication is a different form of action. Agents use it to share world views, negotiate options, and direct subordinates. Usually, the communication is transmitted over a network. The messages can be sentences in a formal language or can be the exchange of computer data structures.

In software agents, all actions are represented as some form of communication to the subsystems being controlled by the agent. These communications cause the controlled system to begin some internal process or modify a process that is underway. Examples of these actions would include sending email to the owner, beginning a database search, canceling a buy order, or sending a work request to another agent.

5.6 Perception Technologies

Perception is the activity of sensing and interpreting the operating environment. Like the action technologies, perception technologies are tightly coupled to the problem domain.

5.6.1 Sensing

Sensing is the process where some attribute or attributes of the environment are measured. Many types of sensors exist, from the simplest switch to complex multi-spectral image detectors, and they use a wide range of technologies to measure the environment.

In many cases, the sensor will take several steps to convert the sensed attribute into electronic signals that can be interpreted by computers. For example, a microphone converts the sound energy moving in the air into tiny vibrations in a mechanical system. These mechanical vibrations are then converted into electrical energy usually using piezo-electric or electromagnetic devices. In some applications, the sensor is integrated directly into the actuator and helps the actuator achieve its desired effect. An example is the angle sensor in a servo system.

Sometimes the sensor and actuator are the same device. Electric motors, for example, convert electrical energy into mechanical rotation (actuator), but they can also be used to convert mechanical rotation into electrical energy for measurement (sensor). Another example is communication. The sending agent is performing an action and the receiving agent is sensing. When the receiving agent responds, the roles reverse.

5.6.2 Image and Signal Processing

To utilize the information supplied by the sensor, some sensor-specific computation is necessary. This computation is used to clean up the information supplied by the sensor and to apply processing algorithms that detect the attribute being measured.

Some sensors require little processing. A simple mechanical switch is either on or off and its interpretation should be easily understood by the agent. However, mechanical switches bounce when their position is changed, and this bounce causes a rapid series of on/off signals to be generated before the switch settles into its new position. The computer systems interpreting signals from switches are fast enough to detect these bounces and take them into consideration so as not to misinterpret the environment. Therefore, mechanical switches are de-bounced using a small electrical circuit or software. This is a simple example of signal processing.

Other sensors require massive amounts of signal processing. Radar and sonar systems work in environments with a large amount of background noise. If the signals are not carefully processed, the background noise will be incorrectly categorized as an item of interest to the system. Complex, sensorspecific algorithms are used to eliminate the background noise, interpret the results, and find target items in the signal.

In a similar manner, imaging detectors require large amounts of postprocessing to clean up device specific noise and interpret the data and find target items of interest to the agent. In many situations, the algorithms are layered, where lower level algorithms clean up the individual pixels and higher level algorithms merge multiple pixels together to interpret what is in the image.

Image and signal processing can be computationally very expensive. Sometimes this computational expense will justify the use of special purpose processors whose only purpose is to perform the image or signal processing.

5.6.3 Data Fusion

The ultimate goal of perception is to supply the agent with an accurate representation of the working environment. In most real world systems, perception requires the agent to integrate multiple sensor data into a single consistent view of the world. This is often done by having each sensor perform its sensorspecific processing and then merge the sensor output [91,192]. The techniques used vary between domains and the types of sensors being used.

One advantage of data fusion is that it can supply a more accurate understanding of the environment. This is because the data missing from one sensor can be filled in by another. Also, when multiple potential interpretations are possible in the data from a single sensor, information supplied from another sensor can help remove ambiguities.

Data fusion is a complex subject and many design decisions must be made when designing a system. Some of the important issues are:

- How is the sensed data represented? (pixels, numbers, vectors, symbols, etc.)
- How are the different sensor representations merged into the common view?
- How is conflicting information handled? Does one viewpoint win or are all views represented?
- Can uncertainty be represented with appropriate weightings?
- Can information from one sensor be used to fill in holes in the information from another?
- What is the level of granularity in the data (pixels, symbols, etc.) and how will the discrepancies between data granularity and model granularity be handled?

5.7 Testing Technologies

The development of robust agent systems requires testing. Testing supports the detection of errors in implementation where the system implementers overlooked or misunderstood a requirement, or just made a mistake. But another major purpose of testing is to uncover requirements that are missing in the original specifications when the system was designed.

A complete testing plan requires testing at each stage of the development effort and involves numerous strategies. Ultimately, the actual system (hardware and software) should be tested in a realistic environment that is capable of exercising both the nominal situations and many error cases. While this level of testing is important, it can be very expensive. As the cost of computation declined, it became possible to develop very realistic testing environments that exist solely in a computer without physical test hardware. This section will limit itself to software testing relative to cooperative autonomy.

5.7.1 Software Simulation Environments

Software simulation environments are based on the idea that it is possible to use computerized simulations to model accurately not only the system being tested, but also its operating environment. While they cannot replace conventional testing of real systems, software simulations have many advantages. The two testing paradigms (i.e., real-world testing and testing in simulation) can be mutually complementary, and each can be used to cross-check the other.

Each of the testing paradigms has strengths and weaknesses that vary in different domains and problem situations. This set of relationships must be thoroughly considered and understood relative to the strengths and weaknesses of the development team. Only then can a cost-effective decision be reached as to the degree to which the development team utilizes each testing paradigm. NASA mission development teams have a long history of the use of both, and have a high level of corporate knowledge regarding the tradeoffs between the two paradigms for testing space mission systems. The expectation of the degree of cost-effectiveness of testing in a given mission development is strongly related to experience of this kind.

The first advantage, in certain situations, of software simulation environments is that the testing cycle time is shorter. Building real hardware and a physical simulation environment is, in some cases, very expensive. In contrast, again in some circumstances, once the software environment is set up, building models of the test system and its environment is relatively easy to do. This allows experiments in all levels of cooperative autonomy with only modest expense (depending on the problem domain).

As development proceeds, systems will need to be expanded, changed, or replaced. This, in many cases, may be very expensive to do when using physical test hardware. Software simulations, however, may allow these changes to be made quickly and cheaply. Indeed, in many situations, the costs are so low that several different solutions can be developed and tested without concern that the losing solutions will be thrown away.

Software testing environments also allow testing of system components in isolation. This allows components to be developed and tested without the need and cost of building a full model of the system. After testing, the promising components can be further developed and integrated into a complete system for further testing.

In a similar manner, software testing environments allow testing of different levels of fidelity. To perform a quick test of a new idea, it is not usually necessary to build a complex, high fidelity model of the whole system. A simple, low fidelity model will usually tell the designer whether the idea merits further development and testing at a higher fidelity.

Software testing environments usually support superior debugging tools. Simulated systems can give the developer an ability to examine the state of the simulated hardware and software at any point in the execution. If an unexpected event occurs, the simulation can be stopped and examined in detail until the cause is determined. This reduces the testing time and allows the developer to have a better understanding of how well the system works.

Finally, software testing environments, in general, are repeatable and can be designed to systematically test for faults. Real systems must face many environmental challenges and can fail in many different ways. Creating multiple environmental challenges with real hardware can become very time consuming and labor intensive. Software simulations can, in many circumstances, quickly and simply create challenges and test for faults. Also, depending on how the model is created, the software system may automatically create failure scenarios not imagined or tested by human engineers on real hardware. Often it is the unexpected faults that cause the most damage.

Crucial in deciding whether to adopt the approach of simulating the hardware and the environment is the question of whether the development team has the necessary understanding of the hardware and the environment. To design and implement simulation software that will have sufficient fidelity (i.e., that will be true to the real world) is, in many cases, fraught with difficulties. This issue must be considered with dispassionate realism and with the recognition of the failure of some past development efforts in both government and private-sector projects, due to an over-simplified view of the hardware, the environment, and the dynamics of the combination, and due to an overly optimistic view of the capabilities of the development team.

The most important design issues when creating software testing environments are:

- The goals of the simulation
- The level of fidelity required
- The types of debugging desired
- Required interactions with other simulations or software/hardware components

Though simulation software can exercise the flight software, the test software must always remain physically separate from the flight software so that it does not accidently get activated when deployed.

The rest of this section will examine the common techniques used to create software testing environments.

5.7.2 Simulation Libraries

Simulation libraries consist of a set of program library routines that the test component calls to emulate an appropriate effect in the test environment. The library routines emulate the effect of the calls, interact with the environment, and return appropriate information to the calling component. This type of environment works well for systems with well-defined control interfaces.

Simulation libraries are the simplest testing environments to build because they only require the normal software tools and programmer knowledge to develop. Additionally, any software language and environment can be used for the simulation development. Simulation libraries tend to have low fidelity, particularly in the timing of hardware actions. Again, the simulation software would have to be kept separate from the flight software.

A good example of this technique would be testing a robot control system on a robot simulation. The control program would make the same calls it would on the actual hardware and the simulated calls move the robot in the environment and return the same sensor results the real system would return. If the simulation is good enough, the control system can then be placed in the real robot and run.

5.7.3 Simulation Servers

A more advanced technique is to use a separate simulation server for components being simulated and have the component being tested interact with the simulation server using a local area network. The primary advantage of this approach is that the simulation and test component run on different hardware, and therefore, each can have free access to needed resources. The simulation can be at an arbitrary and appropriate level of fidelity. If a high fidelity simulation is desired, very fast hardware can be used to simulate it. The software being tested can be run in an environment very close to the real system with the network interfaces being the only potential difference. If, for example, the planned hardware for the control software being tested is an unusual flight computer, then that exact computer can be used. It does not affect the simulation server's hardware.

A simulation server has an additional advantage: it can usually accommodate multiple simulated entities in the same simulated environment – a big advantage when testing cooperative autonomy systems. Though it is somewhat uncommon, software library techniques can be used in simulations accommodating multiple simulated entities.

5.7.4 Networked Simulation Environments

Another software simulation technique uses a network of computers to simulate multiple entities in a single shared simulation environment. The focus in designing these systems is on efficient protocols that allow an accurate simulation in the shared environment. The primary advantage of this approach is that it allows very large and complex systems to be simulated.

The distributed interactive simulation (DIS) system funded by DARPA and the military is an example of a networked simulation environment [6, 40, 48]. This very successful project built a training system that allowed hundreds of simulated entities (tanks, planes, helicopters, missiles, etc.) to interact in a realistic battlefield simulation. The thrust of this effort was originally to have human-controlled simulated hardware engage in interactions. The continuations of DIS have developed in many areas. Work relative to cooperative autonomy simulations is intended to bring an understanding of how real hardware and simulated entities can interact together in the same simulated/real world. Such an understanding would support the development and testing of real spacecraft or robots interacting with simulated ones in mutual cooperation.

Agent-Based Spacecraft Autonomy Design Concepts

In this chapter, we examine how agent technology might be utilized in flight software (FSW) to enable increased levels of autonomy in spacecraft missions. Again, as stated in the Preface, our discussion relates exclusively to uncrewed assets (robotic spacecraft, instrument platforms on planetary bodies, robotic rovers, etc.) or assets that must be capable of untended operations (e.g., ground stations during "lights-out" operations). The basic operational functionality discussed in Chap. 2 is accounted for and allocated between flight and ground, and between agent and nonagent FSW. Those technologies required to enable the FSW design are touched upon briefly, and new autonomous capabilities supportable by these technologies are described. Next, the advantages of the design from a cost-reduction standpoint are examined, as are the mission types that might benefit from this design approach.

For each design concept, the appropriate distribution of functionality between agent and nonagent components and between flight and ground systems is discussed. For those functions assigned to remote agent processing, enabling technologies that are required to support the agent implementation are identified. Further, we determine for which mission types each design concept is particularly suitable and describe detailed operational scenarios depicting the remote agent interactions within the context of that design concept.

6.1 High Level Design Features

The general philosophy of this FSW concept is that conventional nonagent software will encompass all H&S onboard functionality. This FSW segment, referred to as the "backbone," will also contain functionality directly supporting H&S needs, such as slewing and thruster control, and will be directly accessible by realtime ground commands in the "normal" manner. The Remote Agent software will embody mission support functionality, such as planning and scheduling, as well as science data processing. To achieve its goals, a Remote Agent may access backbone functionality through a managed bus whose

data volume will be restricted to avoid potential interference with critical backbone processing. Should any or all Remote Agents "fall off" the bus or simply cease to operate, the backbone will not be impacted, though science observations could well be temporarily degraded or terminated.

The FSW backbone processing is time-driven. Each function within the backbone receives a well-defined "slice" of processing time. All backbone functionalities are scheduled to begin at a fixed time relative to the start of a processing cycle and will complete within one or more time cycles. A list of backbone functions is provided below, with more detailed descriptions following:

- 1. Safemode
- 2. Inertial fixed pointing
- 3. Ground commanded thruster firing
- 4. Ground commanded attitude slewing
- 5. Electrical power management
- 6. Thermal management
- 7. H&S communications
- 8. Basic fault detection and correction (FDC)
- 9. Diagnostic science instrument (SI) commanding
- 10. Engineering data storage

6.1.1 Safemode

Safemode is the key enabler of higher-level FSW functions. Safemode guarantees that no matter what may happen during the conduct of the mission, be it a hardware malfunction (for redundant H/W components) or a FSW abnormality, as a last resort (provided the problem is detected), the spacecraft can enter a state where no further permanent damage will be done and spacecraft H&S can be maintained. The allocation of this function to the FSW backbone is clearly essential.

Return from safemode is contingent on diagnosing the underlying cause of the problem and implementation of corrective action, either selection of a "canned" solution or creation of a new solution. Currently, return from safemode is within the purview of the flight operations team (FOT), but, with a sufficiently advanced onboard FDC, the capability could be shared. Some spacecraft are provided with multiple levels of safemode, but nearly all Goddard Space Flight Center (GSFC) spacecraft have a sun-pointing safemode to guarantee power and SI safemodes to protect delicate SIs. In the past, most GSFC spacecraft also had a hardware safemode in case the onboard computer (OBC) went down.

6.1.2 Inertial Fixed Pointing

Although the safemode function discussed above guarantees maintenance of spacecraft H&S, recovery from safemode and interception of the science

schedule can be a lengthy process during which precious observing time will be lost, potentially irretrievably depending on the mission lifetime and the nature of the science. So, a lesser response to anomalies than safemode entry is highly desirable.

An inertial fixed pointing mode serves this purpose. For celestial pointing spacecraft, being inertially fixed is effectively being in observing mode without a science target and without making use of fine error sensor or SI data. The pointing accuracy and stability will, therefore, not meet mission requirements, but will be adequate to facilitate re-initiation of the science program. And virtually all spacecraft, independent of mission type, require an inertial fixed pointing mode to support sensor calibrations, such as gyro scale factor and alignment. For these reasons, and since ground controllers may need to command the spacecraft to an inertial mode either during the immediate postlaunch checkout phase or in response to unusual spacecraft performance during mission mode, the FSW backbone will require the capability to transition to and maintain the spacecraft at an inertially fixed pointing. Note that typically the following associated functions are part of and subsumed under an inertial hold mode:

- 1. Gyro data processing and drift bias calibration
- 2. Attitude control law(s) for fixed pointing
- 3. Reaction wheel commanding and momentum distribution
- 4. Actuator commanding for momentum management

6.1.3 Ground Commanded Slewing

Just as ground controllers need access to an inertial fixed pointing mode in response to an anomalous condition on-orbit, the need may arise to slew the vehicle attitude to a different orientation. Control of nominal spacecraft slewing activities in support of science execution or sensor/instrument calibration could be allocated to an appropriate Remote Agent. But some basic slewing capability must also reside within the FSW backbone to ensure that in the event of the onset of an anomaly originating from within (or at least affecting) the agent itself, the backbone retains the capability for slewing the spacecraft to a needed orientation, such as sun pointing.

6.1.4 Ground Commanded Thruster Firing

Similarly to slewing, control of nominal thruster firing could be allocated to an appropriate Remote Agent. Nominal use of the propulsion subsystem might well be highly automated, coupled with autonomous onboard planning and execution of orbit stationkeeping maneuvers. However, the need will still exist to provide the ground controller a "back door" into the propulsion subsystem to enable ground commanding of emergency angular momentum dumps or orbit changes. For that matter, the FSW backbone itself may need to command

the thrusters to damp out unusually high booster separation rates. So, some access to and control of thruster firing by the backbone is critical.

6.1.5 Electrical Power Management

Even in conventional FSW designs, electrical power management is somewhat isolated from other FSW functionality to protect against "collateral" damage in the event of a processor problem or software bug. Some functionality within the electrical power subsystem may even be hard-wired. Although managing SI requests for use of electrical power might rightly reside inside a Remote Agent, management of the most critical spacecraft power resource must reside within the FSW backbone.

6.1.6 Thermal Management

As with electrical power management, thermal management in conventional FSW designs is already set apart from other FSW functionality for the same reasons. Also for the same reasons, it also must reside within the FSW backbone.

6.1.7 Health and Safety Communications

Although nominal communications should be highly autonomous to support efficient downlink of the massive quantities of data generated by modern SIs, in the event of a spacecraft emergency, ground controllers must be guaranteed direct access to any information stored onboard to help them understand and correct any problem. They must also be capable of sending commands to any hardware element involved in the resolution of the crisis. To support these fundamental mission requirements, especially during the early postlaunch checkout phase, full ground command uplink capability (probably via a low volume omni antenna) must be provided both through the FSW backbone and through a command and data handling (C&DH) uplink card in the event the main processor(s) is (are) down. Low rate downlink through an omni must also be provided to guarantee the ground's reception of key status data, as well as (if necessary) more detailed engineering data stored just prior to the onset of a problem.

6.1.8 Basic Fault Detection and Correction

Just as all processing necessary for maintaining the spacecraft in safemode must be contained within the FSW backbone, the key functionality associated with transition into safemode must also reside within the backbone. In particular, basic FDC limit checking and safemode-transition logic should be part of the backbone. On the other hand, more elaborate FDC structures that adapt to new conditions, create new responses, or employ state-of-art methodologies such as state modeling, case-based reasoning, or neural nets should be hosted within an appropriate agent.

6.1.9 Diagnostic Science Instrument Commanding

If a SI experiences an anomaly within its hardware or embedded software, it will be necessary to downlink diagnostic data to the ground for analysis. In that event, as the SI or spacecraft platform may even be in safemode at the time, it is necessary to provide a "bullet-proof" capability for ground control to retrieve that diagnostic data, or even send troubleshooting commanding to the SI to generate additional information needed to solve the problem. To provide ground control a direct route to the SI for these purposes, the associated diagnostic SI commanding should also reside within the FSW backbone.

6.1.10 Engineering Data Storage

Just as communications with ground control could be fully autonomous and handled by Remote Agents (both on the spacecraft and in the ground's lightsout control center), so too could management of science data collected in the course of an observation. On the other hand, the main purpose of engineering housekeeping data is to support fault detection in realtime onboard, as well as anomaly investigations post facto on the ground. To support these ground control efforts, which at times are critical to spacecraft H&S, the backbone should control storage and management of these data.

6.2 Remote Agent Functionality

The Remote Agents have the responsibility for achieving science mission objectives. They are event-driven, independent modules that operate as background tasks, which in some flight hardware architectures would be distributed among multiple processors. These mission-support agents are expected to negotiate among themselves and cooperate to accomplish higher level goals. For example, they might strive to achieve optimal science observing efficiency by coordinating the efforts of individual Remote Agents (potentially distributed between the flight and ground systems) responsible for data monitoring and trending, science and engineering calibration, target planning and scheduling, and science data processing.

To ensure the viability of the spacecraft, the FSW is designed such that individual Remote Agents can unexpectedly terminate functioning without impacting the FSW backbone. To isolate the agents from the backbone, an executive agent controls agent communication with the backbone. The bandwidth available for agent-to-backbone "conversations" is limited, so as not to impact backbone processing. Bandwidth limitation is also used to control agent communications with spacecraft hardware through the backbone.

A list of potential Remote Agent functions is provided below, with more detailed descriptions in the following subsections:

- 1. Fine attitude determination
- 2. Orbit determination (and other reference data)
- 3. Attitude sensor/actuator and SI calibration
- 4. Attitude control
- 5. Orbit maneuvering
- 6. Data monitoring and trending
- 7. "Smart" fault detection, diagnosis, isolation, and correction
- 8. Look-ahead modeling
- 9. Target planning and scheduling
- 10. SI commanding and configuration
- 11. SI data storage and communications
- 12. SI data processing

6.2.1 Fine Attitude Determination

For safemode or inertial-hold purposes, gyros and sun sensors supply adequate information to guarantee acquisition and maintenance of a power-positive spacecraft orientation, or to ensure that the spacecraft will not drift far from its current attitude. So an autonomous onboard capability to determine a fine-pointing attitude (i.e., pointing knowledge good to a few arcseconds) is not critical to H&S. It is, however, essential for mission performance for any precision pointer, be it earth- or sun-pointer.

Given that "Lost-in-Space" startrackers are now available that directly output attitude quaternions, this function, for some missions, has already been realized in an independent hardware unit. For other missions with higher accuracy pointing requirements, FSW (which could be developed with an agent structure) would still be required for interpreting and combining data from the startracker with those from a fine error sensor or SI. The simplest implementation of this capability would involve creating a fine-attitude-determination agent that continuously generated attitude solutions, independent of the need of any other agent for their use. These solutions could then be stored in a data "archive" pending a request by other agents. Calibration data and sensor measurements needed by the attitude-determination agent could be stored in similar archives until requested by the agent. Old data (either input to or output from the agent) could be maintained onboard until downlinked, or if not needed on the ground, simply overwritten periodically.

Orbit Determination (and Other Reference Data)

Typically, accurate spacecraft position information is not required to support safemode processing or emergency communications, so orbit determination is not a function that must be resident in the FSW backbone. For spacecraft in near-earth orbits, global positioning system (GPS) can be used for orbit determination. GPS also provides a time fix, simplifying onboard clock correlation. The GPS solution is purely a hardware one, constituting a fully independent, modular agent.

On the other hand, if future orbit information must be predicted, or if GPS cannot be used (as, for example with a mission at the L2 Lagrange Point), a software solution, often requiring input from the ground, must be used instead. The simplest implementation of this capability would involve creating an orbit determination agent that continuously generated orbit solutions, independent of the need of any other agent. These solutions could then be stored in a data "archive" pending a request by other agents. Similar to attitude data, old data could be maintained onboard until downlinked, or if not needed on the ground, simply overwritten periodically.

In addition to the spacecraft ephemeris already discussed, there often is an onboard need for Solar, Lunar, ground station, and/or tracking and data relay satellites (TDRS) position information as well. These data are usually computed onboard via analytical models and would supplement the spacecraft orbit data already supplied by the orbit agent. In a similar fashion, other reference information such as geomagnetic field strength and South Atlantic Anomaly (SAA) entrance/exit times (for a set of SAA contours) could be supplied by the agent, as required by the mission.

6.2.2 Attitude Sensor/Actuator and Science Instrument Calibration

Currently, very few calibrations are carried out autonomously onboard. For nearly all current GSFC missions, gyro drift biases are calibrated at high cycling rates via a Kalman filter, and for small explorer (SMEX) missions, onboard magnetometer calibration is standard. Neither of these is required to be performed (at least immediately) when in safemode or inertial hold, so they need not be part of the FSW backbone. In the future, however, the dynamic quality of other spacecraft hardware may require more elaborate onboard calibrations, both engineering-related and SI-related. In such circumstances, consolidating all calibration functionality within a single Remote Agent would facilitate interaction with (for example) a data monitoring-and-trending agent or a planning-and-scheduling agent, whether those agents are located onboard or on the ground.

6.2.3 Attitude Control

Other than for responding to ground realtime commands, spacecraft slews are performed in support of scheduled science observations or calibration activities. So, the attitude slew function can safely be assigned to a Remote Agent as long as a basic slewing capability is included in the FSW backbone as well. The agent version could be much more sophisticated than the version in the FSW backbone. For example, it could automatically adjust slew trajectory to avoid known constraint regions, while the backbone version followed a fixed eigenvector path. This same agent would have responsibility for high-precision fixed-pointing in science observation mode, a function also not required by the backbone.

6.2.4 Orbit Maneuvering

As long as direct thruster control is available to the ground via the FSW backbone, it is acceptable from a risk management standpoint to assign a higher level orbit maneuvering capability to a Remote Agent. This application would be responsible for planning routine stationkeeping maneuvers based on ground-developed algorithms.

To illustrate how the capability might be utilized inflight in a Remote Agent context, imagine that an autonomous onboard monitoring-and-trending agent has determined that a geosynchronous spacecraft will leave its orbital box in the next 24 h. That agent would notify the planning and scheduling agent of the need for a stationkeeping maneuver in that timeframe. The planning and scheduling agent would request that the orbit maneuvering agent supply it with a package of thruster commands that will restore the orbit to lie within accepted limits. The orbit maneuvering agent, in part, uses input data supplied by the trending agent, then constructs two sets of thruster commands, one to initiate a drift back to the center of the box and the second to stop the drift when the spacecraft reaches this objective. The orbit maneuver agent also specifies time windows within which both command packages must be executed, and submits its products to planning and scheduling, which schedules their execution.

6.2.5 Data Monitoring and Trending

Although rule-based limit checks for safemode entry need to be contained within the FSW backbone, more elaborate data monitoring and trending functionality could safely reside within a separate Remote Agent. This agent could utilize a statistical package to perform standard data analysis procedures such as standard deviation, sigma-editing, chi-squared, etc. to evaluate the believability of new measurements and/or calculated parameters and to extrapolate predicted values from past data. This agent would also be responsible for providing data products required by FDC to detect the presence of operational anomalies that do not threaten spacecraft H&S, but could impact science data-collection efficiency or success. The output from this agent would also be of interest to application agents such as the orbit-maneuvering agent, as previously discussed. Note that the data monitoring-and-trending agent could also utilize a wide variety of AI products to generate and interpret its results, including state modeling, case-based reasoning, and neural nets, as well as a simulation capability.

6.2.6 "Smart" Fault Detection, Diagnosis, Isolation, and Correction

Some "primitive" fault detection driven by simple rule-based limit checking must be present in the FSW backbone to control entry into safemode and other critical mode transitions and/or hardware reconfigurations. However, that portion of fault detection solely concerned with threats to performance of science observations (as opposed to the hardware itself) would be assigned to a Remote Agent. That same agent could also perform fault diagnosis and isolation and determine the appropriate course of corrective action, which could be either a predetermined "canned" solution or an original solution independently created by the agent. To the extent that fault diagnosis, isolation, and correction exist within the backbone, they should be viewed as canned correlations between simple limit checks and canned corrective actions developed and specified by ground personnel, primarily before launch. This agent may also utilize a wide variety of AI products to generate and interpret its results, including state modeling, case-based reasoning, and neural nets.

6.2.7 Look-Ahead Modeling

The FSW backbone is a time-driven software element that looks at data in realtime and responds in realtime or near-realtime: there is no need for any look-ahead modeling functions within the backbone. On the other hand, many Remote Agent applications (such as planning and scheduling, data trending and monitoring, and orbit determination and maintenance) may require the support of look-ahead models. Examples include ephemerides, solar intensity, wheel speed, and SAA entrance/exit. The model outputs could be generated continuously independent of the need of any other agent for their use. These solutions could then be stored in a data "archive" pending a request by other agents.

6.2.8 Target Planning and Scheduling

Traditionally, FSW has been time-driven, and full planning and scheduling responsibility rested with the ground system. The ground would generate an absolute-timed target list that the FSW would execute precisely as specified. If conditions were inappropriate for the science observation to execute (for example, if no guide stars had been found), the spacecraft would still remain uselessly at that attitude until the pointer in its C&DH FSW moved to the time of the slew to the next target. However, as planning and scheduling capabilities are migrated from the ground to the spacecraft, more flexible responses to anomalies are enabled leading to greater overall operational efficiency at reduced costs. It is these new capabilities, which are unnecessary to H&S maintenance within the FSW backbone, that are the purview of the target planning-and-scheduling Remote Agent. This agent may also utilize a wide variety of AI products, including state modeling, case-based reasoning, and neural nets.

6.2.9 Science Instrument Commanding and Configuration

Other than to support collection of SI diagnostic data and transition into SI safemode, no SI commanding and configuration functionality needs to reside in the FSW backbone. That functionality could be provided by a Remote Agent. The actions of the agent could be driven by receipt of template structures (generated by the planning and scheduling agent) defining what SI configuration is needed and/or what operational usage is desired. Separate agents could be assigned for each SI, or a single Remote Agent could handle the entire job. The agent(s) would also have responsibility for verifying the legality of any directives issued to an SI.

6.2.10 Science Instrument Data Storage and Communications

Although management of engineering data storage, as well as management of transmission of that data and SI diagnostic data, must be the responsibility of the FSW backbone, storage onboard and transmission to the ground of primary product SI data collected during a science observation are purely associated with satisfying science mission objectives and may, therefore, be entrusted to a Remote Agent. With a lights-out control center, one can easily conceive of all the duties associated with storing and transmitting SI data being handled autonomously by Remote Agents, where no general principle dictates whether these agents belong better in the ground system or the flight system. Instead, the decision on their location and relative distribution of responsibilities could, with a generalized flight/ground architecture, be made on the basis of simple convenience on a mission-by-mission basis.

6.2.11 Science Instrument Data Processing

Other than processing of SI data required to monitor the H&S of the individual SIs, no SI data processing need be contained within the FSW backbone. The advantage of assigning this functionality to a Remote Agent is that it enables cooperative behavior with other agents where the coupling of their functionality can yield a greater whole than the sum of their individual functions acting in isolation. For example, following collection of wide field data from a CCD detector in the course of a scan over a region of the celestial sphere, the data could be processed onboard by this agent, which (using a case-based pattern recognition algorithm) could identify point sources appropriate for more detailed study. The agent would report its results to the planning-and-scheduling agent, which would notify the slew agent to execute a return to the specified target coordinates. At the same time, the SI commanding-and-configuration

agent, which prepares the SI, as necessary, for the upcoming observation utilizing plans generated by the planning-and-scheduling agent, would command execution of these plans at the proper time. The data generated would then be processed by the SI data processing agent for storage and later transmission to the ground by that agent. By this means, immediate onboard response removes the need for scheduling revisits at a later date.

6.3 Spacecraft Enabling Technologies

To support the Remote Agent functionality described in the previous subsections, a series of enabling technologies will be required. Many of these technology elements have already been flown on NASA missions, but have not as yet achieved mainstream status. Others are proposed for use on upcoming NASA missions, but are not as yet flight-proven. And others are still purely in the "talking" stage, but could reasonably be expected to be available in the next decade. The following is the list of technology items:

- 1. Modern "Lost-in-Space" star trackers
- 2. Onboard orbit determination
- 3. Advanced flight processors
- 4. Cheap onboard mass storage devices
- 5. Advanced operating system
- 6. Decoupling of scheduling from communications
- 7. Onboard data trending and analysis
- 8. Efficient algorithms for look-ahead modeling

The following subsection discuss these in more detail.

6.3.1 Modern CCD Star Trackers

The Wilkinson Microwave Anisotropy Probe (WMAP) mission (launched in 2000) utilized CCD star trackers containing both calibration and star catalog information, permitting the star tracker itself to output a direct measurement of its orientation relative to the celestial sphere. A simple multiplication (within the attitude control subsystem (ACS) FSW) by the device's alignment matrix relative to the spacecraft body yields the spacecraft's attitude quaternion. This capability enables fine attitude determination "on the fly" without a priori initialization, an important autonomy feature supporting calibrations, target acquisitions, communications, target of opportunity (TOO) response, and smart FDC.

6.3.2 Onboard Orbit Determination

Onboard orbit measurement via GPS is widely in use on commercial satellites and is planned for use on the Global Precipitation Measurement (GPM) mission (scheduled launch date of 2013). So, onboard orbit determination on the fly without a priori initialization is a capability readily available now for spacecraft orbiting in the Earth's immediate vicinity. GPS timing also allows synchronizing the spacecraft clock with ground time.

The future challenge is to develop hardware and algorithms enabling fully independent onboard orbit determination far from the earth. Promising work along these lines is currently in progress, where star tracker measurements of celestial objects, such as the Moon, earth, and other planets could be used to derive spacecraft position information. Position information for the planets and the earth's moon are available now from self-contained analytical models requiring very infrequent (if any) input parameter updates. Independent spacecraft orbit determination is necessary to support other autonomous functionality, such as onboard calibrations, target acquisitions, communications, TOO response, dynamic scheduling, and smart FDC.

6.3.3 Advanced Flight Processor

Nearly all the functionality projected in this design for implementation within a Remote Agent framework can be accommodated by flight processors that have already supported missions. However some capabilities require such abundant computing power (to support massive data processing and/or intricate logic analysis) that the development of a flight-capable, radiation hardened high performance computer may be required. These items may include areas such as onboard SI data processing, dynamic scheduling with lookahead, and smart fault diagnosis and corrective plan creation.

6.3.4 Cheap Onboard Mass Storage Devices

Trends in this area are bringing costs (both dollar and weight) down so rapidly that in the future it is unlikely that onboard storage capacity will be a major limiting factor with regard to flight data processing capabilities. However, as SI data volume production is increasing rapidly as well, available storage will continue to play a dominant role in communications trade issues so long as all SI raw measurements must be downlinked, or lossless compression is utilized.

6.3.5 Advanced Operating System

Although no autonomy function discussed here absolutely requires an operating system with file manipulation capability comparable to a general purpose computer, such a capability would (for example) simplify ground handling of science observation downlink products, as well as the C&DH FSW's own manipulation of data onboard.

6.3.6 Decoupling of Scheduling from Communications

Decoupling of scheduling from communications has been beneficial in simplifying the Rossi X-ray Timing Explorer (RXTE) ground system's planning and scheduling process and in reducing TOO response time. A similar decoupling could also be expected to make the job of onboard scheduling easier and enable a more dynamic and autonomous scheduling process onboard.

6.3.7 Onboard Data Trending and Analysis

To support smart fault detection, diagnosis, and isolation, more elaborate capabilities for onboard data trending and analysis will be required. As discussed previously, a statistical package is needed to perform standard data analysis procedures (e.g., standard deviation, sigma-editing, chi-squared, etc.) in order to evaluate the believability of new measurements and/or calculated parameters and to extrapolate predicted values from past data. The statistical package will also be needed to support planning of new calibration activities.

6.3.8 Efficient Algorithms for Look-Ahead Modeling

Conventional FSW typically is time-driven realtime software that tries to determine the best course of action at the moment as opposed to the optimal decision over a future time interval. This is an appropriate approach for standard FSW applications such as control law processing and fault-detection limit checking, but is not likely to yield acceptable results in areas such as planning and scheduling. For such applications, at least a limited degree of look-ahead modeling will be required to capture the complete information critical to efficient decision making.

6.4 AI Enabling Methodologies

To implement the Remote Agent functionality discussed in the previous section, more sophisticated modeling tools and logic disciplines will be required than the simple rule-based systems that have been employed in FSW in the past. In particular, a smart fault detection, diagnosis, isolation, and correction agent could also utilize a wide variety of AI products to generate and interpret its results, including state modeling, case-based reasoning, and neural nets. "Intelligent" constraint-evaluation algorithms used in planning and scheduling, as well as the data monitoring-and-trending agent, will also need these AI enabling methodologies. The following discusses onboard operations that would be enabled through the use of collaborative Remote Agents.

6.4.1 Operations Enabled by Remote Agent Design

Although the functionality assigned in the previous section to Remote Agents potentially could equally well have been implemented onboard in a more conventional fashion, a major asset of the Remote Agent formulation would have been lost. If onboard implementations of these functions are evaluated as separate entities on an objective cost-benefit basis, for most GSFC missions, one would probably find that more than half of those items could be developed more cost effectively in the ground system and would not significantly reduce operational costs if migrated onboard.

In particular, the following functions looked at in isolation should probably remain on the ground:

- Virtually all calibration, both science and engineering
- All data trending (although limit checking should remain onboard)
- Any smart fault detection, diagnosis, isolation, and correction
- Any look-ahead modeling
- All target planning; nearly all scheduling
- All specification of SI commanding and configuration (execution onboard)
- Nearly all communications decision making
- Nearly all SI data processing

The remainder already are largely onboard autonomous functions, or in the near future, will be likely.

However, an analysis of this kind misses a key aspect of the Remote Agent formulation, namely the ability of agents to engage in "conversations," negotiate, and collaborate. It is that distinctive nature of Remote Agents that makes their collective functionality more powerful than the simple sum of their component parts.

In the subsections that follow, examples are provided illustrating how the cooperative capacity of Remote Agents can yield more sophisticated (and more profitable) performance than they could achieve working alone. Operational capabilities enabled by the agents working together include:

- 1. Dynamic schedule adjustment driven by calibration status
- 2. TOO scheduling driven by realtime science observations
- 3. Goal-driven target scheduling
- 4. Opportunistic science and calibration scheduling
- 5. Scheduling goals adjustment driven by anomaly response
- 6. Adaptable scheduling goals
- 7. SI direction of spacecraft operation
- 8. Beacon mode communications
- 9. Resource management

6.4.2 Dynamic Schedule Adjustment Driven by Calibration Status

For spacecraft with high performance and accuracy requirements but very stable calibration longevity (both science and engineering), a ground scheduling system will be able to schedule science observations well in advance with a high degree of reliability. This confidence is due to the knowledge that if calibration accuracy degrades prior to execution of a key scheduled observation, the degradation will occur gradually and gracefully, leaving the ground scheduling system ample time either to insert a re-calibration activity to restore nominal spacecraft function or to re-schedule any extremely performance sensitive observations for a later time.

However, if calibration stability is extremely dynamic, the "half-life" of the ground scheduling system's spacecraft state knowledge may be significantly less than the lead time on execution of many of the scheduled targets, in which case a prescheduled canned observing sequence will experience many observation failures and poor overall efficiency. An alternative to this is to couple realtime calibration status monitoring directly into planning and scheduling of both science observations and re-calibration activities.

For example, the ground could uplink to the spacecraft a target list having a label attached to each science target defining the level of telescope calibration needed to make the science observation worthwhile. The planning and scheduling agent could then elect to schedule only those list targets (or goal generated targets, as discussed later) compatible with the spacecraft's current state of calibration (as determined by the monitor-and-trending agent). After all such targets are exhausted, the agent could schedule a re-calibration activity (as created by the calibration agent) to bring the spacecraft back up to specifications, following which the remaining list targets could be observed. Alternately, utilizing another label supplying priority information, the presence of any target with a high priority designation could be sufficient to cause planning and scheduling to order a re-calibration activity immediately.

Although this example primarily illustrates the cooperative behavior of planning-and-scheduling, data monitoring-and-trending, and calibration agents, a somewhat lower degree of participation by several other agents (including attitude/orbit determination, attitude/orbit maneuvering, SI commanding-and-configuration, and SI data processing) may arise as a requirement.

6.4.3 Target of Opportunity Scheduling Driven by Realtime Science Observations

For most TOOs, the ground system, due to its access to data from other space and ground observatories, will be best positioned to designate appropriate TOOs for its spacecraft and adjust the observing schedule accordingly. Also, for those TOOs identified by processing SI measurements from the spacecraft itself, as long as those TOOs do not have a short lifetime and as long as spacecraft revisits are not extremely costly or inefficient, maintaining both SI data processing and TOO scheduling responsibility within the ground system is probably a better trade than migrating the capabilities to the spacecraft. However, if the TOO has a very short time duration (relative to turnaround times for ground SI data processing and scheduling), if target revisits are costly, or if the spacecraft is not in regular contact with the ground, advantage can be gained by installing onboard at least limited functionality in these areas.

For example, after execution of high level survey measurements of a general region of the celestial sphere, an SI Remote Agent could process the data collected and identify point targets appropriate for the follow-up work. That agent could then inform the planning-and-scheduling agent of the presence of interesting targets in the immediate neighborhood, which could then schedule and execute those targets immediately (via SI commanding), avoiding the operational overhead of a re-visit following processing of the survey data on the ground. So by coupling the efforts of planning-and-scheduling and SI data-processing Remote Agents onboard, overall system responsiveness can be greatly improved and measurements of some classes of science targets can be performed that would otherwise not be observed. This capability has already been utilized on the Swift mission.

Although this example primarily illustrates the cooperative behavior of the planning-and-scheduling agent, the SI commanding-and-configuration agent, and the SI data-processing agent, additional participation by other agents (including attitude/orbit determination, attitude/orbit maneuvering, and SI data storage) could also be required.

6.4.4 Goal-Driven Target Scheduling

Traditionally, the ground system has uplinked to the spacecraft a fixed, timesequenced target list that the FSW has executed in an absolute time-driven fashion precisely as specified by the ground. This is probably the optimal scheduling approach for missions in which the science targets are easily defined in advance and the spacecraft is in regular contact with the ground. However, for missions more isolated from ground control or where the science targets can only be specified by general characteristics and have limited contact duration opportunities (e.g., asteroid flybys), a goal-oriented targetspecification technique may be essential.

For the case of an asteroid flyby, the SI data-processing Remote Agent could not only perform straightforward data reduction of the measurements, it could also employ sophisticated pattern-match techniques (possibly via casebased reasoning) to identify regions for closer investigation. In this respect, the previously discussed TOO scheduling is just a subset of goal-driven target scheduling. However, goal-driven scheduling might also include the onboard specification of targets without any a priori ground input whatsoever. For example, suppose a percentage of a spacecraft's mission consists of performing survey work in selected regions of the celestial sphere. Knowing those portions of the celestial sphere targeted for survey work, the total percentage of time allocated to surveying, and guidelines regarding how much time should be spent on surveys within, say, a week-long interval, the scheduling agent could identify opportunities to "piggy-back" survey observations following successful ground-specified targets at approximately the same attitude.

6.4.5 Opportunistic Science and Calibration Scheduling

The previous discussion of goal-driven scheduling cited an example that could also be considered opportunistic scheduling, i.e., piggybacking a goal-specified survey observation on top of a specific ground-specified point target. However, opportunistic scheduling need not be goal-driven.

For example, while the spacecraft is executing attitude slews commanded for the purpose of acquiring targets and collecting science as specified by the ground, a calibration Remote Agent could keep a running tally of "miss" distances derived by comparing star-tracker-computed attitudes with gyromeasured angular separations. After collecting an adequate amount of data, the agent could calculate new gyro scale factors and alignments to maintain slew accuracy within performance requirements. The calibration computations would be done in the background on a nonimpact, computational timeas-available basis relative to ongoing science or higher priority engineering support activities. Should representative slew data of a specific geometric variety be lacking, the calibration agent could request that the scheduling agent reorder the target list or consult its targeting goals to see whether a target could be "created" that would enable the collection of necessary engineering data as well. Failing that, the scheduling agent could simply add a slew of the required type to complete the calibration activity, if the FDC agent (using data provided by the monitoring and trending agent) indicated that a recalibration needed to be performed in the near future to maintain spacecraft slewing accuracy requirements.

6.4.6 Scheduling Goals Adjustment Driven by Anomaly Response

This operational capability highlights cooperative behavior between the Remote Agents performing scheduling and fault-correction. Suppose that the review, by the fault-detection agent, of the fault monitoring-and-trending agent's spacecraft-state analysis has generated an error flag, which the faultisolation agent has associated with a solar array drive. Suppose further that the fault-diagnosis agent has determined that continued nominal use of that drive mechanism will lead eventually to failure of the mechanism, and in response, the fault-correction agent rules that the solar array should be slewed to its safemode orientation and not moved from that position until instructed by ground command. At this point, the fault-detection agent could rule that all targets that previously were validated as acceptable and that required a solar array orientation other than the safemode position should now be considered invalid, thereby requiring that the scheduling agent delete those targets and at least compress the schedule, or possibly even generate a new one reflecting the currently degraded hardware state.

6.4.7 Adaptable Scheduling Goals and Procedures

A previous subsection discussed the cooperation of onboard Remote Agents to satisfy a ground-specified goal. In this subsection, we examine how the goals themselves might be modified/specified by the flight system based on inflight experience. Also, ground-specified procedures used for achieving those goals (either science observation or calibration goals) could similarly be modified. The input data for the goal modification process might be obtained from either of two sources, described now in the following paragraphs.

First, one could envision an onboard validation process that applies ground-supplied metrics to the execution of science observations, or for that matter, engineering support activities. Sticking to the science application for simplicity, a monitoring and trending agent could calculate an overall observing efficiency, as well as individual efficiencies associated with the various different types of science (a function of SI, mode, target type, etc.). The fault detection agent might then look for relatively poor efficiency outliers, so the fault diagnosis, isolation, and correction agent could determine which onboard goals or canned scripts/procedures might require enhancement.

Second, to define what changes might be made to those scripts/goals deemed inefficient, an onboard simulation function (under the control of the monitoring and trending agent) could be dedicated to running simulations on other approaches, either canned options provided by the ground or new options independently derived by the spacecraft (via ground-supplied algorithms) by running "what if" simulations in the background on a nonimpact basis.

Although this capability would be primarily useful for missions where the spacecraft can expect to be out of contact with the ground altogether for entire phases of the mission (or during nonscience cruise phases where most of the onboard computing power is idle), it might also prove useful for spacecraft constellations where several spacecraft themselves are cooperating to achieve overall constellation goals (see Chap. 9).

6.4.8 Science Instrument Direction of Spacecraft Operation

This section examines how a smart SI can influence the behavior of spacecraft platform subsystems. It is not unusual to allow SIs to command the ACS to adjust the attitude of the spacecraft to facilitate a target acquisition, or to shift the target from one SI aperture to another. For example, HST used SI specified "peak-ups" to facilitate realtime target acquisitions inside SIs having narrow FOVs. However, the formalism of Remote Agents enables more elaborate interactions.

For example, suppose a spacecraft's SI complement included an instrument with a movable component (say, for survey scanning). Depending on the mass/motion of the component and the mission jitter requirements for fine pointing, the attitude perturbations induced by the component's motions could threaten satisfaction of the jitter requirement. To deal with this potential problem, the agent embodying the SI could inform the attitude control agent of its intended motions in advance, so that the pointing control law could compensate for the disturbance in advance, so as to protect a precisionpointing fixed-boresight SI observation from blurring. If the attitude agent determined it could not compensate for the effect, the fault correction and scheduling agents could be brought into the "discussion" to resolve the conflict, perhaps by giving priority to the precision pointing SI.

A more interesting solution would be for the more demanding SI to "announce" to an SI executive agent when it needed a particularly steady platform, and then have the executive agent prohibit motions by the lower priority scanning instrument during such periods. Similar restrictions would also apply to other moving structures such as gimbaled antennas.

Another application for SI/agent direction of platform behavior is SI calibration. At its least elaborate, it would be straightforward to automate the process of periodically placing an SI in self-test mode (using its internal test source), comparing the output relative to a nominal signal, and passing on any discrepancy information to an SI calibration function. This could be a purely internal SI implementation. A more sophisticated example would involve periodic re-observations of a baseline target for the purpose of checking whether SI performance has remained nominal (and this example would entail agent collaboration with the attitude-control and scheduling agents). This example would also include the option for recalibrating the SI or optical telescope assembly (OTA) if a significant degradation has occurred (where the re-calibration would be the job of the calibration agent).

6.4.9 Beacon Mode Communication

For missions where contact with the ground is expected to be irregular and where down-link requirements are driven by the spacecraft's need to inform the ground only of the results of science observations processed onboard (triggered by event messages from an onboard processing agent), it would be appropriate for the spacecraft to control communications. Similarly, if contact is driven by the spacecraft's need to confer with the ground regarding some problem experienced onboard (also triggered by an agent message), it would be most efficient/least costly for an onboard communications agent to "dial up" the ground's agent. In its extreme form, contact responsibility totally migrated to the spacecraft is referred to popularly as beacon mode, which can include
periodic status summaries as well as the downlink of science measurement results and requests for help. Because the nature of beacon mode is to inform the ground of all onboard "experience" the spacecraft deems to be interesting to the ground, beacon mode operation can involve the interaction of the communications agent with nearly all the other onboard agents, including SI data processing, SI data storage (potentially part of communications), data monitoring and trending, smart fault detection, diagnosis, isolation, correction, etc.

6.4.10 Resource Management

Spacecraft operations concepts usually are driven by constraints imposed by limitations in onboard resources, both of the expendable varieties (for example, fuel or cryogen) and renewable varieties (such as electrical power, computing power, data storage, and telemetry bandwidth). When these resources are in demand by more than one distinct spacecraft subsystem or component, a mechanism must be provided (either ground-based or flight-based) to deal with the inevitable conflicts that will result. Other potential sources of conflict (which may also be thought of somewhat abstractly as resources) are priorities on use of onboard functionality, such as attitude control or thruster utilization/orbit adjustment. A Remote Agent formalism provides a particularly convenient infrastructure for resolving issues arising from overlapping needs.

In this application, one can envision a resources management agent (possibly subsumed under the executive agent) responsible for evaluating all resource requests in excess of a nominal set of limits associated with the set of agent users of those resources. The management agent would evaluate those requests relative to overall resource envelopes to determine whether the request can be honored, whether it must be rejected, or whether the need is high enough in priority that other agents must surrender a portion or all of their individual assets. The management agent would have to perform a similar function if an onboard anomaly unexpectedly reduced the availability of a resource. Due to the broad characterization of what constitutes a resource, resource management potentially can involve the interaction of that agent with all the other onboard agents. The management agent would also ensure that no resource deadlock situations arise.

6.5 Advantages of Remote Agent Design

The Remote Agent design described in the previous sections presents the opportunity for significant overall mission cost savings arising from its enabling of higher operational efficiencies, reduced FSW development costs, and reduced FSW testing costs. In the following subsections, each of these contributing factors is discussed.

6.5.1 Efficiency Improvement

Traditionally, onboard schedule execution has been mostly absolute timedriven. This design choice, while enabling an extremely simple onboard element in the overall schedule execution process, has necessitated the development of a very expensive ground planning and scheduling element run by a large, expensive operations staff. The resulting product, due to the high reliance on the ground system's approximate look-ahead modeling with its associated worst-case time pads, has also been somewhat inefficient and quite limited in its capacity to respond to off-nominal or unexpected events. The following discusses cost savings enabled by an onboard scheduling capability.

Event-Driven Scheduling

Autonomous, event-driven management of onboard activity transitions by Remote Agents has the potential for considerable cost savings. There always has been considerable use of event-driven transition control of onboard processes. For example, onboard logic can control mode transitions such as slewing to inertial hold when body rates and pointing errors have dropped below threshold levels, or the transition from inertial hold to safemode when anomalies are detected requiring immediate, extreme responses. Onboard logic, either existing within FSW or hardwired into the flight hardware itself, has also controlled transitions in sensors/SIs from search activities to tracking activities when the target objects' signal profile warranted it. Still, the use of an event-driven mechanism for managing transitions between scheduled science observations has often been precluded by the complexity of the spacecraft's orbital environment, performance demands imposed by the mission, and the computational limitations of the flight data system.

An example of an event-driven short-term scheduling system is the Adaptive Scheduler originally proposed for use on the James Wells Space Telescope (JWST). In Adaptive Scheduling's simplest form (from an onboard perspective), the ground system would still be responsible for generating an ordered list of desired science targets, but no absolute-time tags would be attached. The FSW then would observe the targets in the specified order, but nominally would trigger the execution of the next observation on the list in response to a FSW event signaling the successful completion of the previous observation. This avoids the waste of potential observing time engendered by the old paradigm's use of worst-case time pads to space out absolute-timed commands that might otherwise "collide" with each other.

The Adaptive Scheduler would also have the capability to insert into the timeline engineering events, as required, when informed by other FSW subsystems that an action needs to be taken. For example, when angular momentum has built up to the point that a momentum "dump" must be performed, the ACS would accordingly inform the Adaptive Scheduler, which would then find

a good point in the timeline to insert the dump, and would order execution of the dump by the ACS at the appropriate time.

Further, in the event of the detection of onboard anomalies, the Adaptive Scheduler could take corrective action, which might involve deleting a target from the list, temporarily skipping a target until it could be observed at a later date, or even more elaborate reordering of the ground-supplied target list. For example, if after the ACS had exhausted all its acquisition options the JWST's fine guidance sensor had still failed to acquire a guide star, the Adaptive Scheduler could command a deletion of that target from the observation list without waste of additional observing time and could order that a slew to the next target on the list be initiated.

For JWST's L2 orbital geometry, an event-driven scheduling system would be simple, efficient, and easily responsive to at least a substantial list of possible anomalies. A LEO orbit would present much more of a challenge due to the complexity of the environment and would probably require the support of an elaborate look-ahead capability.

In one respect, adaptive scheduling is more complex than the absolutetimed paradigm. Because under the old paradigm, the start and end of onboard activities were very well-defined, the process for insertion of realtime commands into the timeline (provided ample uplink opportunities were available) was well-defined, if at times awkward. However, under the new approach, ground operations staff will not always be certain when "safe" opportunities for realtime uplink might present themselves. Therefore, additional "intelligence" would need to be present onboard to enable the FSW to consolidate information and prioritize commands from a wide variety of realtime and preplanned/predicted sources, such as uplinked realtime commands, realtime sensor output, realtime event messages from onboard performance monitoring, and the ground-supplied target list.

The FSW must still support ground uplinks of time-tagged commands along with its list of event-driven activities for those situations where an activity must be performed within a specific time window (for example, a ground-planned orbit stationkeeping maneuver). This need is addressable via a short-term look-ahead capability that would support delaying scheduling any timeline events that would "swamp" the absolute-time window until after the absolute-timed event has completed, and giving the absolute-timed event priority over any routine engineering event that needs to be inserted.

It is apparent just from the Adaptive Scheduler example that considerable gains in efficiency can be realized just by moving from a fully groundprogrammed, absolute time-driven style of operation to a ground-ordered, onboard event-driven style. Exploiting the onboard system's knowledge of realtime information enables a flexible response to on-orbit events that could greatly reduce loss of valuable observation time arising from conservative time-padding in ground look-ahead models. Nonreplaceable onboard resources could be utilized more economically by expending them in response to realtime measured needs, while renewable resources could be allocated more appropriately and with smaller contingency safeguards or margins. Also, a side benefit to this new paradigm is that wherever expensive ground-system look-ahead modeling can be replaced by simple onboard realtime response, cost savings are achieved within the ground system as well.

The example cited in this subsection, the originally proposed JWST Adaptive Scheduler, is a rather modest effort at exploiting the potential offered by an event-driven FSW operation. The Remote Agent FSW formulation described in this section would have much more powerful capabilities. More than just being an event-driven executor of a ground-planned, ground-scheduled (both long-term and short-term) target list, a Remote Agent design provides a framework for migrating short-term scheduling responsibility to the onboard system, greatly increasing the spacecraft's capacity for managing onboard resources more efficiently, reacting to anomalies without loss of observing time, and responding to TOOs while the opportunity presents itself.

Short-Term Scheduling

Also, as discussed for the Adaptive Scheduler's event-driven scheduling, short-term scheduling by the FSW must be able to accommodate realtime commanding originating from the ground system. For onboard short-term scheduling, the complications will be even greater as there is no reason to expect the ground to be cognizant of ongoing onboard activities when the ground's command is sent. To deal with such intermittent interrupts of varying levels of priority, one can envision an onboard "fuzzy logic" reasoning module that juggles an array of priority levels and time windows attached to ordered/unordered target lists, target clusters, ground-originated realtime commands, ground-originated absolute-timed commands with/without windowing, onboard-originated housekeeping/engineering activities, realtime onboard sensor measurements, and externally-originated realtime data, and then deduces which activity (or activities) should be executed next (or in parallel).

Further, one could envision an event-driven Remote Agent engaging in a dialog with more powerful computer resources (and models) on the ground to acquire deeper look-ahead information prior to making its activity transition decisions. The agent could even establish dialogs with other spacecraft to benefit from their realtime measurements. For example, if two earth-pointing spacecraft were flying in formation, the one arriving later over the target could interrogate the earlier one regarding cloud cover and other conditions and make appropriate SI reconfigurations prior to arrival.

6.5.2 Reduced FSW Development Costs

Conventional FSW is highly integrated code, optimized for timing performance and computational efficiency. As a result, long-term FSW maintenance tends to be very expensive (relatively small code updates may require very wide-spread regression testing) and FSW reuse has been rather limited. However, a Remote Agent implementation offers much promise for producing significant reductions in FSW development and testing costs.

Since each agent can be built as a standalone module (consistent with the objects associated with object-oriented design), the development schedule for an agent can be synchronized with the availability of the information needed to define its requirements. As long as their interfaces with the executive agent and other agents with which they need to interact with are well-defined, those agents (for example) associated with hardware components that will only be procured toward the end of the lifecycle can be developed later than those agents whose functionality is well-understood from the start.

As agents are developed, they can easily be added to the system, and if a problem develops with an agent inflight, it can be dropped offline without H&S impact to platform or payload. As this approach to developing FSW matures, it will be possible to build up a library of agents from previous missions, which can be reused economically on future missions once protocols and standards for agent communication have been established, eventually stimulating the creation of generalized COTS products, which will greatly facilitate the reduction of FSW development costs.

Significant reductions in FSW testing costs can also be expected. Since each applications agent is decoupled from direct communication with the FSW backbone and since their communication with the backbone is bandwidth limited by the executive agent, a modification to an applications agent should not normally require full-scale system-level regression testing. The modified agent could instead just be tested at a module level and then added back into the flight system. As an agent can drop offline without impacting the FSW backbone (and its corresponding H&S functionality), less stringent (and less costly) testing standards may be applied to the applications agents than to the backbone and than as currently applied to all conventional FSW. Finally, the similarity of this software architecture to typical ground system architectures should enable (in some cases) initial agent software development in the ground system, with the associated cheaper software development and testing methodologies, with later migration to the flight system following operational checkout (see Chap. 9 for a detailed example of this, called progressive autonomy).

6.6 Mission Types for Remote Agents

In this section, potential advantages of Remote Agents (over a conventional FSW design) are evaluated at a high level relative to a set of characteristic mission types, specifically:

- 1. LEO celestial pointers
- 2. LEO earth pointers
- 3. Geosynchronous-earth-orbit (GEO) celestial pointers

- 4. GEO earth pointers
- 5. Survey missions
- 6. Lagrange point celestial pointers
- 7. Deep space missions
- 8. Spacecraft constellations

The following subsections discuss these in more detail.

6.6.1 LEO Celestial Pointers

For LEO celestial pointers, the intermediate- and short-term scheduling problem is quite complex due to the many time-dependent constraints characteristic of near-earth orbits. The ground system component responsible for that function must be large and expensive, embodying many complex spacecraft and environmental models. But, since the dominant input to optimizing scheduling is not realtime measurements, there is little advantage in migrating short-term planning to the spacecraft, given that communications with the ground can be expected to be regular, frequent, and of long duration. It is even unclear that event-based scheduling would win a cost-benefit trade with conventional ground preprogrammed absolute time-based scheduling, given the need for look-ahead to maintain high scheduling efficiency. An exception to this general statement would be realtime support for TOOs, where the special onboard processing would not schedule the TOO itself, but instead simply make platform and payload housekeeping adjustments, as necessary, to support the change in the plan. So the planning and scheduling area probably is not a productive application for agent formalism for LEO celestial pointers, and similarly, SI commanding and configuration is likely to be limited in carrying out ground instructions, probably via templates stored onboard to reduce uplink data volume.

On the other hand, one can easily imagine additional calibration functions being migrated to the flight system for LEO celestial pointers. For example, as miss-distance data from slews are collected onboard, the current state of calibration could be monitored autonomously and a background task could re-compute the gyro scale factors and alignments (supported by background fine-attitude and orbit-determination background tasks). Whenever fault detection determined that the current calibrations no longer were acceptable, fault correction would direct use of the current "latest-and-greatest" computed values. Implementation of each of these functional areas in the Remote Agent formalism via the design structure described earlier in this section would greatly facilitate the cooperative behavior described above without adding risk to the maintenance of platform and payload H&S via the FSW backbone.

Summarizing, implementing the following additional onboard functions as Remote Agents would be consistent with the lights-out control center philosophy:

- 1. Fine attitude determination (as discussed)
- 2. Orbit determination (as discussed)
- 3. Attitude sensor/actuator and SI calibration (as discussed)
- 4. Attitude control (execute slew requests; fine pointing)
- 5. Orbit maneuvering (plan and execute orbit stationkeeping)
- 6. Data monitoring and trending (as discussed)
- 7. "Smart" fault detection, diagnosis, isolation, and correction (as discussed)
- 8. Look-ahead modeling (probably not required)
- 9. Target planning and scheduling (not required)
- 10. SI commanding and configuration (execute science calibration requests)
- 11. SI data storage and communications (in cooperation with ground agent)
- 12. SI data processing (just for target acquisition)

LEO Earth Pointers

The scheduling problem for LEO earth pointers is much simpler than for LEO celestial pointers. Long duration look-ahead is no longer an issue since the spacecraft's orbit cannot be predicted to a high level of accuracy very far in advance. The planning aspect of the problem, however, is very target dependent, and might vary with time depending on science prerogatives. One can, therefore, imagine a set of templates (with ground tunable parameters) associated with individual targets (or target types) that an onboard scheduling agent could invoke whenever the onboard orbit-determination agent (a GPS receiver coupled with a short-duration orbit propagator) in conjunction with a data monitoring-and-trending agent deemed the target was coming into view. Because fuel available for stationkeeping maneuvers is directly equivalent to mission lifetime, it is unlikely that most NASA LEO earth pointers would expend fuel for large orbit-change maneuvers for the purpose of observing TOOs (short-duration events like volcanoes, for example), though the requirements on TOO response for some non-NASA spacecraft (for example, military imaging spacecraft) may be more demanding. Given that GPS receivers now give the spacecraft itself more immediate access to accurate, current spacecraft orbit data than does the control center, migrating some portion of the shortterm LEO earth-pointer scheduling responsibility to the spacecraft to improve observational efficiency could be justified on a cost-benefit basis for some missions of this type, and migrating routine orbit-stationkeeping maneuvers to the spacecraft could reduce operations costs as well.

As in the case of the LEO celestial pointer in the previous section, one can easily imagine additional calibration functions being migrated to the flight system for LEO earth pointers. There also may be significant advantages in providing an SI data processing capability/agent onboard for the purpose of distinguishing between useful data-taking opportunities (for example, for a Landsat spacecraft, in the no-cloud-cover situation) and unusable conditions (in this case, a full-cloud-cover situation). Summarizing, implementing the following additional onboard functions as Remote Agents would be consistent with the lights-out control center philosophy for LEO earth pointers:

- 1. Fine attitude determination (needed for orienting an imaging subsystem)
- 2. Orbit determination (as discussed)
- 3. Attitude sensor/actuator and SI calibration (as discussed)
- 4. Attitude control (needed for orienting an imaging subsystem; fine pointing)
- 5. Orbit maneuvering (plan and execute orbit stationkeeping, as discussed)
- 6. Data monitoring and trending (as discussed)
- 7. "Smart" fault detection, diagnosis, isolation, and correction (as discussed)
- 8. Look-ahead modeling (needed in conjunction with orbit-trending)
- 9. Target planning and scheduling (scheduling as discussed)
- 10. SI commanding and configuration (templates associated with targets)
- 11. SI data storage and communications (in cooperation with ground agent)
- 12. SI data processing (as discussed)

6.6.2 GEO Celestial Pointers

A GEO celestial pointer's operational constraints are much more straightforward than those for a comparable LEO mission, so much so that it is possible to operate the spacecraft in a "joystick mode," as demonstrated by the International Ultraviolet Explorer (IUE), where blocks of time were assigned to astronomers who could command the spacecraft directly when making their observations. This is not to say that automation is not critically important for enabling full utilization of spacecraft capabilities and performance accuracy at optimal efficiency, but that automation can be implemented on the ground in support of a human operator rather than migrated onboard to provide autonomous function. Since full contact with the ground station is nominally possible at all times and time delays in transferring spacecraft supplied data to the ground station are small, and because FSW development and testing costs are likely to remain significantly higher than those of ground software of corresponding size and complexity, this mission type is probably not a good candidate for an onboard Remote Agent implementation on a cost-benefit basis. However, Remote Agents may find many useful applications within the ground system's lights-out control center.

6.6.3 GEO Earth Pointers

As with the GEO celestial pointer described in the previous section, operational constraints are much more straightforward than those for a comparable LEO mission. And as before, full contact with the ground station is nominally possible at all times and time delays in transferring spacecraft supplied data to the ground station are small. Since FSW development and testing costs are likely to remain significantly higher than those of ground software of corresponding size and complexity, this mission type is probably not a good candidate for an onboard Remote Agent implementation on a cost-benefit basis. Again, Remote Agents may find many useful applications within the ground system's lights-out control center.

6.6.4 Survey Missions

The nominal operation of a survey mission is by its very nature very straightforward and predictable. Planning and scheduling are well-defined for extremely long durations. For a properly designed spacecraft, calibrations should be stable and easily monitored/trended by the ground. SI data need only be collected and dumped to the ground for processing and archiving; no onboard processing would be required, unless contact time and downlink bandwidth/onboard data storage availability is an issue. Again, in a cost-benefit trade, the relatively higher cost of FSW software over equivalent ground software will always be an argument in favor of a ground implementation decision (at least initially, with migration as an option). The one exception might be in the area of fault detection, diagnosis, isolation, and correction (supported by data monitoring and trending), where given the likelihood of nonfulltime contact between flight and ground, an autonomous capability to deal with a wide range of anomalies and still keep the mission going could prove very useful. Otherwise, it would probably be more cost efficient to maintain most potential Remote Agent functionality in the ground system.

6.6.5 Lagrange Point Celestial Pointers

Lagrange points are points of stable or unstable equilibrium relative to the influence of the combined gravitational fields of two celestial objects acting on a third object in orbit about the two objects. For simplicity of presentation, we will restrict our discussion to the sun-earth Lagrange points, though Lagrange points can be defined for any two celestial objects close enough in proximity for both of their gravitational fields to affect a third orbiting object (for example, the behavior of the Trojan asteroids relative to the sun-Jupiter system). For a given celestial system, there will be a total of five Lagrange points, three (L1, L2, and L3) in-line with the two objects and two (L4 and L5) off-axis, as illustrated in Fig. 3.2. Objects at the on-axis points (L1, L2, and L3) are in unstable equilibrium (i.e., objects at these points will drift off in response to perturbations), while those at the two off-axis points are in stable equilibrium.

The simple orbital geometry and benign environmental conditions (relative to those of a near-Earth orbit) make Lagrange-point orbits good choices for spacecraft with sensitive thermal constraints. The absence of occultations facilitates long duration observations of dim objects. In addition, for spacecraft in halo orbits about L1 or L2, the distance to the Earth is as low as a million miles, so communications-bandwidth needs do not demand excessively large antennas or transmitter power.

From a planning and scheduling perspective, the space environment is simple enough (contrary to the conditions at LEO) to enable considerable onboard scheduling autonomy. The absence of complex time-dependent constraints makes event-driven processing a preferred choice (over absolute timedriven, totally preprogrammed processing) for higher operational efficiency, and in fact, has been considered for use by the JWST mission. However, the value of onboard short-term scheduling is highly mission dependent.

Although communications with the ground can be regular, frequent, and of long duration, there still can be almost equally long periods when the spacecraft is out of contact. At these times, a smart fault detection, diagnosis, isolation, and correction capability, in conjunction with onboard data monitoring and trending, would (depending on designed redundancy capacity) allow the spacecraft to "fly through" a failure and continue its mission while out of contact with the ground. However, an equally viable approach would be simply to rely on conventional fault detection with a transition to a robust safemode to guarantee platform and payload H&S until contact with the ground is regained. Currently, the answer to the question of which approach is best is highly mission dependent and can only be determined following a rigorous cost-benefit trade. In the future, the answer will depend on how much progress has been made in standardizing onboard smart fault processing and in reducing FSW development costs relative to comparable ground software costs.

Similar arguments can be made for many of the other functions that might potentially be assigned to Remote Agents. Onboard SI data processing could be very helpful to opportunistic identification, scheduling, and acquisition of science targets (as well as TOO response). It could also enable a reduction in downlink bandwidth by telemetering only processed science products, not raw data, to the ground. But for some missions, the science targets are quite predictable and greater overall scheduling efficiencies may be obtained via a conventional ground scheduling system. And for most missions, the astronomer customer will insist on receiving the raw data, rather than the downstream processed product. So again, an evaluation of the desirability requires rigorous cost-benefit analysis.

Summarizing, implementing the following additional onboard functions as Remote Agents would be consistent with the lights-out control center philosophy:

- 1. Fine attitude determination (needed for target acquisition)
- 2. Orbit determination (not needed)
- 3. Attitude sensor/actuator and SI calibration (need is a function of H/W design)
- 4. Attitude control (execution of slew requests; fine pointing)
- 5. Orbit maneuvering (infrequent; planned on ground)

144 6 Agent-Based Spacecraft Autonomy Design Concepts

- 6. Data monitoring and trending (as discussed)
- 7. "Smart" fault detection, diagnosis, isolation, and correction (as discussed)
- 8. Look-ahead modeling (a function of planning and scheduling autonomy)
- 9. Target planning and scheduling (as discussed)
- 10. SI commanding and configuration (execution of science and calibration requests)
- 11. SI data storage and communications (in cooperation with ground agent)
- 12. SI data processing (as discussed)

6.6.6 Deep Space Missions

Deep space missions are tailor-made for the flexibility and responsiveness enabled by a Remote Agent implementation of an autonomous spacecraft. Being out of contact with the ground for very long periods and with significant radiosignal propagation time for communications, the spacecraft needs to have a greater degree of "self-awareness" not only to maintain H&S, but even to perform its mission efficiently. This is particularly true of a mission like an asteroid flyby where mission-critical observing decisions must be made in realtime. Although very complex deep space missions have successfully been performed by Jet Propulsion Laboratory (JPL) in the past, JPL has rather clearly determined that the key to maintaining their recent downward trend in mission cost is to promote steadily increasing onboard autonomy through the use of formalisms such as Remote Agents. In the DS1 mission, the responsibility of health monitoring was transferred from ground control to the spacecraft [99, 135, 195]. This marked a paradigm shift for NASA from its traditional routine telemetry downlink and ground analysis to onboard health determination.

6.6.7 Spacecraft Constellations

Being recent in development, where deployment for now is largely restricted to communications networks, spacecraft constellations represent a new opportunity for flight autonomy applications and are covered in Chap. 9. Whereas consideration of Remote Agent applications to the previous mission types concerned interactions of spacecraft subsystems with each other or with the ground, for constellations, the scale of interaction expands to conversations potentially between all the various members of the constellation. In a complex "conversation" of this type, just the job of determining which members of the constellation should be included in the conversation, when they should enter the interaction, and when they should drop off can be a thorny problem. More discussion on this topic will be provided later, but to support constellation interactions, a hierarchy of subsystem subagents controlled by agent spacecraft will need to be introduced.

6.6.8 Spacecraft as Agents

While spacecraft share many of the properties of the agents described above, the unique environment in which they operate makes unusual demands on their design. Since spacecraft are mobile, self-contained, and externally focused, they are often viewed as space-based robots, but this is not the complete picture.

While mobile, a spacecraft consumes much less of its time and resources for navigation than does a comparable robot. Navigation usually happens at only a few fixed points in the mission and the spacecraft is focused on other issues the rest of the time. In addition, the external orientation of a spacecraft is primarily for the use of science sensors whose data are usually shipped to earth and not used directly by the spacecraft. Most of the other sensors can be viewed as internally focused, distributed throughout the vehicle, and their purpose is housekeeping or health management of the craft. They perform activities to manage power, manage angular momentum, and keep the craft correctly positioned. This internal focus has led some to argue that spacecraft should be viewed as immobots. Certainly, some immobot technologies should be included in future spacecraft designs.

As autonomous spacecraft become more common, they will find themselves in the role of determining which science goals to pursue based on the current situation. If, for example, a pursued science goal cannot be met because of external events or an internal failure, the spacecraft will choose between the other available goals to maximize the science returned. Analyzing and prioritizing the information returned to humans is a primary area of research in software agents. In addition, software agents are a principal focus in the effort to build cooperative capabilities. In the future, it is likely that groups of spacecraft will work together to achieve larger science goals. These technologies, first worked out in software agents, need to be included in spacecraft designs.

While all of these agent technologies represent elements of the whole picture, NASA has the burden to evaluate them and adapt the technologies for spacecraft use. There are many ways that space-mission agent technologies differ from nonspace agent technologies. Most software agents are ephemeral; their only goal is to acquire, manipulate, and exchange information, and the only resource they consume is computation. Spacecraft are not ephemeral. They exist in the real world and their primary resources are sensors and actuators. Actions consume tangible resources, and many of these resources are irreplaceable.

Actions in a spacecraft are usually costly, and once an action is taken, it may not be reversible. This makes it necessary for planning to factor-in the resources used and future cost carefully if the action cannot be retracted. These issues are usually ignored in software agents. Even most robots and immobots are not as deeply concerned about these issues as are spacecraft. Therefore, the planning and control systems of other agent technologies must be carefully studied and possibly modified for insertion in spacecraft systems.

Software agents assume that communication costs are modest and that communication is rapid. These systems often have gregarious agents that would rather ask for help than work out a solution for themselves. In space applications, most communication is expensive and the timeliness of delivery depends on the distance between the vehicles or between a vehicle and a ground station. While collaboration is desired and necessary to achieve objectives, approaches that are more introspective will be necessary to limit communications and handle speed-of-light delays.

Much of the work on software agents has involved the construction of informational agents. Their primary purpose has been the acquisition of knowledge in support of human goals. While spacecraft both collect information and support human goals, to achieve their objective, they must make many real-world decisions that require types of autonomy not necessary in the construction of informational agents.

When a software agent has a programming bug and fails, it is modified and restarted with little fanfare. This is often true of ground-based robots and immobots. Certain software failures in spacecraft are dangerous and could cause the loss of the whole mission.

Most believe that agent cooperation means communicating with humans or other agents to work toward a common goal. This is not the only mode of cooperation. Two or more spacecraft flying in formation may not be in direct communication with one another, but they still cooperate as the data they collect are merged to form a more complete picture. These simple modes of cooperation can assist in achieving missions with lower costs and risks.

As will be discussed in the next chapter, the potential for cooperative autonomy technologies in spacecraft systems is large, but, as this section has shown, the use of this technology will require considerate and careful design.

Cooperative Autonomy

The philosophy of "faster, better, cheaper" reflected NASA's desire to achieve its goals while realistically addressing the changing environment in which it operated. One outgrowth of this philosophy is the shift from performing science missions using a few complex spacecraft to one where many simple spacecraft are employed. While simple spacecraft are faster and cheaper to build and operate, they do not always deliver better science. There must be offsetting compensation for any loss of power to deliver science value by employing innovative technologies and new methodologies that exploit the use of less complex spacecraft. One such technology is cooperative autonomy.

Cooperative autonomy flows from the study of groups of individuals in terms of how they are organized, how they communicate, and how they operate together to achieve their mission. The individuals, in the context of space missions, may be human beings, spacecraft, software agents, or a mission operations center. From modeling and studying the cooperative organization and the interactions between its members come insights into possible new efficiencies and new technologies for developing more powerful space missions by which to do science more cost effectively.

Cooperative autonomy also creates new opportunities. Its technologies support cross-platform collaboration that allows two or more spacecraft to act as a single virtual platform, and thus, possess capabilities and characteristics that would not be feasible with a single real platform. For example, with small telescope apertures on multiple, coordinated small spacecraft, the resultant aperture of the virtual combined telescope can be much larger than the telescope aperture on any single spacecraft.

This chapter outlines a model for cooperative autonomy. NASA's current mission organization is described and discussed relative to this model. Virtual platforms are also modeled and these models are used to assess the impact that virtual platforms may have on the current NASA environment. Optimizations are suggested that could lower overall operations costs, while improving the range and/or the quality of the science product. The optimizations could be performed in a controlled and staged manner to minimize undesirable impacts and to test each change before the next is implemented.

The technologies for constructing cooperative autonomy systems will be discussed. Some of the technologies enable cooperation between humans and nonhuman entities, while other technologies enable fully autonomous spacecraft. Many computer technologies are necessary, and each is at a different level of readiness for NASA environments.

7.1 Need for Cooperative Autonomy in Space Missions

As indicated above, cooperative autonomy is the study of how humans and computer agents can cooperate in groups to achieve common goals. This section describes some of the challenges faced by NASA in future missions, and informally describes how cooperative autonomy technologies can address these challenges.

7.1.1 Quantities of Science Data

Over the last decade or so, the rate at which science data are collected and returned by spacecraft has risen by several orders of magnitude. This is due, in large part, to advances in sensor and computer technologies. Data handling facilities have been expanded to handle the enormous amounts of incoming data, but the mechanisms supporting science extraction from the data have changed very little. In the current milieu of space missions, cross-mission data analysis is largely ad hoc, with large variations in data formats and media. The cost-effective ability to maintain scientific yield is in doubt without new approaches to support scientists in data analysis, reduction, archiving, retrieval, management, and correlation.

7.1.2 Complexity of Scientific Instruments

The instrumentation available to the scientist has increased, and continues to increase in complexity, making it difficult to utilize fully the resources available. The amount of required documentation increases geometrically with complexity, and so the burden on investigators to become and remain able to effectively use any given instrument could easily get out of hand and become untenable for many investigators. Through experience, it has become clear that tools to map scientist's goals onto the space resources available to achieve those goals are essential.

7.1.3 Increased Number of Spacecraft

The number of active spacecraft has multiplied within the last decade. With the recent focus on small spacecraft, the numbers will soar. To realize the potential of the small spacecraft, NASA must utilize them in associated arrangements, i.e., clusters, formations, or constellations. But to prevent a large growth in ground systems to accommodate the large number, they must be organized into virtual platforms that appear as a single entity to the ground. Virtual platforms, especially those that can be dynamically constructed, will allow science, formerly requiring huge spacecraft, to be performed with flexibility and effectiveness.

7.2 General Model of Cooperative Autonomy

The field of cooperative autonomy studies how autonomous agents should work together to achieve common goals. Autonomous agents can be humans, robots, spacecraft, or even factory shop floors. Cooperation can occur between any members of a group of agents, and their reasons for cooperation vary with their domain and mission.

This section will focus on describing a formal model of cooperative autonomy. Since autonomous agents play a central role in cooperative autonomy, it will begin by describing autonomous agents and discussing their properties. Once this groundwork is laid, cooperation between autonomous agents will be defined and four separate patterns for cooperative autonomy will be outlined. These four patterns are used to clarify the different facets of cooperation and describe individual attributes of cooperating autonomous agents. Most cooperative agents will simultaneously use two or more of these patterns.

7.2.1 Autonomous Agents

All autonomous agents share a set of common features independent of their operating domain and individual skills. These common features are outlined in Fig. 7.1 (based on Fig. 5.1, Chap. 5) and can be briefly summarized as: the ability to make plans, to act upon these plans, and to perceive and internalize the external world.

The most important single feature of an autonomous agent is its ability to plan. Planning is the process of making independent choices based on internal goals and the agent's beliefs about the operating environment (see Chap. 5). The results of this process are plans of actions to perform. Plans are the primary mechanism used by an agent to pursue its goals. Without the ability to plan, an agent would hardly be considered autonomous. To be considered a rational and intelligent agent, it must also adapt its plans to the changing situations in the operating environment. The agent's set of beliefs about the operating environment is commonly called its world model. Some agents also model their own internal state (self models).

Agents must act on the plans they make and these actions change the operating environment. Actions vary greatly from domain to domain and are



Fig. 7.1. The features of autonomous agents. The common features can be seen in the diagram on the left and they are the ability to plan, to act, and to perceive the external world. The feature list on the right describes the unique characteristics of an autonomous agent that define its purpose and necessary skills

always customized to the problem being solved. They are in some sense the final manifestation of the agent's plan, and without them, the agent's choices have no meaning.

The ability to perceive the operating environment is another feature common in all autonomous agents. Its primary purpose is to allow the agent to determine whether previous actions were successful and to detect changes in the operating environment. This information is used to update the agent's world model as well as its self-model and ultimately to allow the agent to adapt its plans in the continuing pursuit of its goals. In some agents, the sensing of the operating environment is a goal in itself and this sensed information is delivered to the agent's superior without interpretation.

While the abilities to plan, act, and perceive are common to all agents, there are many features where they may differ. Some of these features are listed on the right of Fig. 7.1. An agent's purpose and its domain of expertise are the defining aspects of an agent. Other key aspects include: what the agent must do, what knowledge and skills it must have, the actions it must be able to perform, the kinds of perception required, and how it plans.

The domain and purpose dictate the degree of individual identity the cooperating agents must have. In some domains, the cooperating agents have common skills, share control of the resources, and have large overlaps in their world models. Often the world models are managed externally to the agents. Agents working in these domains tend to have a low degree of individual identity and are often interchangeable. They have little or no self-model. An example would be a large public scheduling task where the agents are all working together to create a common schedule. The schedule, which forms a large part of the world model and represents the use and control of resources, is external to the agents performing the scheduling.

In other domains, the individual agents have direct control of resources and often they do not share much of their world model with other agents. Such agents tend to have a high degree of individual identity and when cooperating, the individual agents must negotiate as peers to achieve their common objectives. Their self-model is expansive and it represents the status of all resources and systems under their control. Spacecraft are agents with a high degree of individual identity since they manage unique resources in unusual places. For example, multiple spacecraft attempting to perform joint science would need to coordinate their positions, their orientations, and the times when they need to perform actions such as data collection.

In some domains, some goals can be achieved with only sporadic cooperation, while others require continuous contact as the execution proceeds. The previous spacecraft scenario is a good example of sporadic cooperation to collect the necessary science. High precision formation flying is an example of continuous cooperation since each spacecraft must constantly sense and modify its position in relationship to its neighbors. Some domains have a specific hierarchy of responsibility, with the lower agents subservient to the upper. Human controllers of a spacecraft demonstrate this hierarchical cooperation. The controllers make the high level decisions, which are communicated to the spacecraft as the lower level agent for execution. In other domains, agents are direct peers who work together to come to a common agreement.

Finally, agents differ in how well they learn from experience. Some systems have fixed, prescribed rules to specify how the agent should operate under all known situations. Most current spacecraft fit in this category. Other systems learn as they operate. These systems can adapt to changing environments, and over time, can become more skilled. Human agents are a good model of learning agents. With the attributes of autonomous agents described, it is now possible to examine the behavior of groups of autonomous agents. This is the topic of the next section.

7.2.2 Agent Cooperation

There are a number of software-related aspects of cooperative autonomy that are embodied in agent technologies, which is a broad field. Figure 7.2 (see Chap. 5) shows an overview of agent technologies and the lower level technologies that are used to construct agents.

Computer-based agent technology is an active area of research and its goal is to build computerized autonomous agents that fit within the models defined in Sect. 7.2.1. Agents cooperate with one another in different ways. In the simple case, agents work alone to achieve their goals. These types of agents usually assume that no other agent is in the environment, and their plans can be generated and pursued without concern over interference from others. Sometimes, multiple agents work cooperatively to achieve a common goal.



Fig. 7.2. Cooperative autonomy technologies

When designing these agents, issues such as how the domain is divided between agents, how the agents negotiate among themselves, and the degree to which they cooperate must be resolved.

Cooperation occurs whenever autonomous agents work together to achieve common goals. Cooperation can occur in a variety of ways depending on the domain and goals being pursued. The following sections will describe several different patterns of cooperation and show them in relationship to the Plan, Act, and Perceive cycle described above. These patterns are not mutually exclusive, and often the agents must simultaneously engage in multiple cooperation patterns.

Cooperative Planning

Autonomous agents that work in groups as peers use cooperative planning. They cooperate with one another to reach agreement on the actions each member should perform to achieve their common goals. Figure 7.3 depicts a group of autonomous agents that are engaged in cooperative planning. Each layer of the Plan, Act, and Perceive cycles represents a separate autonomous agent and the dots represent the points where these layers communicate during cooperation. In cooperative planning, the communications are focused on exchanging local views of the world, local and shared goals, actions that can be performed, and shared priorities. During cooperation, each member must reach an agreement on the set of actions that it will perform. In this pattern, the agents exchange information during the planning activity and before action is taken.

This form of cooperation is peer to peer. No single agent needs to be in charge, and yet, each agent participates in a global negotiated solution.



Fig. 7.4. Hierarchical cooperation

This pattern of cooperation occurs on a daily basis in human activities. The ubiquitous "weekly status meeting" is an opportunity for a group of people to share the results of the previous week's activities, discuss individual and group priorities, and to agree on the activities of each member for the coming week.

Hierarchical Cooperation

Hierarchical cooperation occurs when the agents have specific responsibilities and authority. In hierarchical cooperation, the superior agents decide the overall strategy and goals that the subordinate agents are responsible for achieving. The subordinate agents will, in turn, plan and execute based on their local view of the domain, and will report their successes and failures to their superiors.

Figure 7.4 shows one example of hierarchical cooperation. In this example, the superior agent plans the activity of the group and its Act step is to

communicate the plan to the subordinates. The subordinate agent receives the plan as a series of goals. These goals are interpreted from the agent's perspective and a plan is generated and executed. The subordinate agent then senses its local environment to determine how to continue the pursuit of its goals.

Hierarchical cooperation is not peer to peer. While a subordinate can negotiate with its superior by communicating its goals, resources, and constraints, it does not make the choice on what will be done. Once the superior has made these choices, the subordinate must attempt to achieve the goals to the best of its ability.

Hierarchical cooperation can be found in many business organizations. The superior agent in Fig. 7.4 could represent a senior management group and the subordinate agent a department within the company. The senior management group sets goals and priorities for the department. The department manager, usually at department status meetings, distributes these goals in the department. The department manager also brings back the results of the department to the senior management group. In this way, the perceptions of the department become part of the perceptions of the senior management.

Computerized systems allow the connection between hierarchical layers to be much tighter than is possible in human cooperation. In these systems, the two levels are directly connected together with the action of the superior agent directly converted into plans for the subordinate agent. Figure 7.5 depicts this scenario. An excellent example of this type of hierarchical cooperation can be found in many robotic control systems. An upper level agent is a slow and deliberative planner that determines the overall strategic goals. The subordinate agent is a high-speed reactive planner. This planner converts the high level goals into direct robotic actions and responds rapidly to changes in the environment.



Fig. 7.5. Tightly coupled hierarchical cooperation



Fig. 7.6. Cooperative actions

7.2.3 Cooperative Actions

Cooperative actions result when autonomous agents work together to achieve the desired objective. Tightly coupled actions result in a form of cooperation depicted in Fig. 7.6.

Cooperative actions are usually coupled with other forms of cooperation. Satellites flying in formation to collect coordinated imagery are an example of cooperative actions. The individual vehicles must cooperate to keep themselves in the proper position and orientation with respect to each other. However, to achieve this high level cooperation, there must be another lower level cooperation that actually keeps the vehicles in formation (possibly some form of closed loop control).

Cooperative Perception

Autonomous agents use cooperative perception when they need to fuse their individual perception information into a single common perception. This process is commonly called data fusion or multisensor fusion [91, 192]. Many different techniques have been developed that fuse either similar or dissimilar kinds of information into a common model of the perceived world. This form of cooperation is depicted in Fig. 7.7.

A good example of cooperative perception is when imagery is collected on the same region of the earth in different spectral bands. When these different views of the region are merged into a single common view, the resulting data are more than the sum of the parts. Another example is the advisors an executive may talk to before a difficult decision. Each advisor has different perceptions of the situation at hand and the consequences of any particular action. The executive weighs the individual contributions and makes a decision.

This completes the discussion of a formal model for cooperative autonomy. An overview of current spacecraft mission management will now be offered to provide a foundation for discussing the application of cooperative autonomy technologies.



Fig. 7.7. Cooperative perception

7.3 Spacecraft Mission Management

Spacecraft mission management is a complex process involving the coordination of experts from diverse disciplines. This section will outline a spacecraft mission model and describe the attributes of each group in the process. While it mixes some of the nomenclature from the deep space and earth science domains, the model provides a generic view of mission management organizations.

Figure 7.8 is a graphical representation of the management of a typical spacecraft mission. While the organization for a specific mission may vary, generally five different activities make up mission management. The process begins with the Science Planning group, which is responsible for creating the science plan representing the science goals. Using the science goals and spacecraft-housekeeping constraints, the Mission Planning group creates the overall mission plan for the spacecraft. This plan is passed to the Sequence Planning group, which converts the high level plan into a series of individual commands that the spacecraft will perform. This command sequence is passed to the Command Sequencer, which uplinks the commands to the spacecraft and monitors the results. Finally, the spacecraft delivers science data, which go to Science Data Processing to be converted into a form useful to scientists. At various points in the process, telemetry or science data provide feedback to allow one or more of the planning groups to change direction.

This section now describes each group, including its responsibilities and products.

7.3.1 Science Planning

Science planning for typical missions involves a large group of scientists, mission engineers, and instrument engineers with different backgrounds, domains of expertise, and modes of operation. This makes the science planning activity an exercise in collaboration and negotiation. A milestone in this process is the production of the science plan, which details the science goals to be achieved



Fig. 7.8. Spacecraft mission management

in the order of their importance or priority. There is little automation assisting the generation of the science plan. The product of this group is typically a text document that is provided to the science team.

7.3.2 Mission Planning

The Mission Planning group is a team of mission and instrument engineers that converts the science plan into the mission plan. The mission plan is designed to achieve the science goals while ensuring the health and safety of the spacecraft. The nonscience activities include maneuvering and housekeeping, as well as scheduling time on a communications link. Mission planning usually requires one or more domain experts for each spacecraft activity. After planning the mission, these experts monitor the spacecraft for safety and health violations and adapt the mission plan as necessary. This group meets more frequently than the science planning group, because their collaborations are focused on producing a very detailed, usually short-term, mission timeline. The product of this group is either a text document or an electronic schedule of mission activities.

7.3.3 Sequence Planning

An individual or small group of mission engineers produces a very detailed command sequence plan using the mission plan as a guide. This plan specifies all of the commands and communications that take place between the ground and the spacecraft. For a typical mission today, the sequence plan is a detailed timeline for every low-level command to be uplinked to the spacecraft.

7.3.4 Command Sequencer

The Command Sequence software uplinks files of commands to the spacecraft, receives down-linked telemetry and information on spacecraft anomalies, and verifies that the file of commands was uplinked with no errors. In case of an uplink or command failure, commands may just be skipped, or sequence planning may be repeated, or the spacecraft may be put in safemode. If an anomaly occurs, depending on the severity, replanning may have to be done after the anomaly is resolved.

7.3.5 Science Data Processing

Science Data Processing converts the raw data down-linked from the platform into useful science data for dissemination to a wide audience. This processing is conducted on the ground and often involves massive amounts of processing and data storage. The Science Planning group, especially the scientists, constantly monitors the science data produced by the mission. Based on the science data produced, the Science Planning group will sometimes modify the activities and priorities for the mission to produce an updated science plan.

7.4 Spacecraft Mission Viewed as Cooperative Autonomy

In this section, we combine the cooperative autonomy model in Sect. 7.3 with the spacecraft mission model of Sect. 7.3. Figure 7.9 shows the current mission organization, now redrawn from Fig. 7.8 in terms of what it might look like with hierarchical cooperating groups. The hierarchical cooperation pattern is well suited to describing the spacecraft mission organization. The results of this combined model are discussed below.

7.4.1 Expanded Spacecraft Mission Model

At each level of the hierarchy in Fig. 7.9, a group is focused on a specific domain and its purpose is defined by its position in the hierarchy. For example, the domain of the science planning group is concerned with the science aspects



Fig. 7.9. Cooperative autonomy view of spacecraft mission control

of the mission. Their purpose is to produce a science plan that maximizes the science returned by the mission. Each layered autonomous cycle in the science planning section represents one member of the planning team (there are four in the figure).

The diagram shows that these members are simultaneously engaging in each of the possible cooperation patterns. The most important of these is the planning collaboration between members, which sets the science goals for the mission. It is also used to evaluate the science returned and to adapt the science plan to maximize results. The science planning group must also cooperate during the action phase when the members work closely together to generate a single science plan that represents their views. This plan is then communicated to the mission planning group. The relationship between the science and mission planning groups is hierarchical in nature and represents the third cooperation pattern shown. The final cooperation pattern is the data analysis, in which the group must fuse all the data coming from the spacecraft to generate a single coherent view of the science being returned. Sometimes the science being returned will cause the teams to reevaluate their plan and adapt it to collect additional information.

The cooperation found in the mission planning group is similar to that of science planning. One difference is that their goals come from the science planning group, instead of being self-generated. Another difference is that the mission group augments the science activities with additional activities that must be performed for a successful mission. The product of this group is the mission plan. The sequence planning group receives the mission plan, evaluates the spacecraft telemetry, and adapts its plans accordingly.

The cooperative model merges the sequence planning group and command sequencer components of Fig. 7.8 into a single autonomy cycle. The planning element takes the mission plan from the mission planning group and converts it into a series of commands that can be sent to the spacecraft and executed. Once the commands are executed, the spacecraft telemetry reports the status of the craft, and command verification checks to see whether the uplinked command file functioned properly. If it did not, the science planning, mission planning, and sequence generation must work to recover from the problem. If the sequence planning group is made up of multiple members (the figure only shows one), then they must collaborate in building the plan and in coordinating what will ultimately be sent to the spacecraft.

In the above example, the spacecraft has no internal autonomy. It executes the commands that were uplinked in a file and returns the results. Many spacecraft have a limited amount of internal autonomy as shown in Fig. 7.10. This automation is in a hierarchical cooperation relationship with the sequence planning group, and is responsible for converting commands into a series of steps, executing them, and monitoring their completion. It also has



Fig. 7.10. Spacecraft automation

the important role of making sure the actions commanded by the sequence planning group do not jeopardize the spacecraft or damage its instruments. When such a condition occurs, the spacecraft automatically takes control and places itself into a safe mode. This type of automation is critical to spacecraft that operate very far from earth: the delay in long-range communications, even at the speed of light, means receiving status information, sending commands, and receiving confirmations will take a long time and the spacecraft can be lost before ground control has a chance to react to an unexpected condition. Furthermore, the delay means that the changed conditions at the spacecraft will invalidate the commands based on the earlier conditions, rendering control by the ground ineffective and potentially counterproductive. In short, ground control of a very remote space asset (e.g., a rover or a spacecraft) under dynamic risk conditions is, in general, not an option.

7.4.2 Analysis of Spacecraft Mission Model

The first thing to note in Fig. 7.8 is the limited use of automation technologies. There are large amounts of human cooperation, communication, and negotiation, but the only fully automated processing occurs at the lowest level of the hierarchy. This is where commands are sent to the spacecraft and automatically verified. This lack of automation dictates a large staff of expert personnel in each of the mission domains. While this helps to ensure the safety and success of the mission, it also means substantial operational costs.

The large degree of human communication and negotiation also severely limits the speed at which the organization can perform decision making. The time required for human decision making has three major impacts:

- Planning time: The entire planning hierarchy has developed to accommodate the slowness of human deliberations. Decisions requiring long deliberations are accomplished at the top of the hierarchy and are infrequent. The middle tier of the hierarchy is focused upon near-term decision making and the lowest level on those of the immediate future. If human deliberations can be reduced or eliminated at any of the levels, improvements can be made in the time required for the planning process.
- Reaction time: The cornerstone of the planning hierarchy is predictive scheduling. This requires all possible activities to be preplanned, with humans involved in all decision making. In the case of a spacecraft anomaly, the mission planning group must be called in to examine the anomaly and re-plan the short-term activities. While this re-planning does not have a major impact on single platform operations, it can easily lead to nonproductive time when one instrument of a platform fails. An automated mission planner could easily redirect the properly functioning instruments to another activity while the anomalous instrument is being examined.

Iteration time: The long lead time required between science planning and execution on-board the platform restricts the science opportunities of a platform. An alteration of the science plan is required for the mission group to refocus on the near-term activities.

A great deal of informal communications and negotiations happens between the members of the planning groups. Members use a variety of mechanisms to achieve consensus on the high-level science goals. Frequent meetings, e-mail messages, and telecommunications are used, as well as a variety of planning and scheduling tools. Normally the science planning group produces a document that is passed on to the mission planning group. While this is acceptable when the group is composed of humans, it severely limits the possibilities for automation of group activities. The informal nature of group activities imposes severe limitations on the speed at which the group can reach consensus.

A careful analysis of mission cooperation shows that in some cases, the wrong type of human cooperation is being applied to a particular level of the hierarchy. This is especially clear in the mission planning group. Mission planning is primarily a traditional scheduling problem, dealing with optimizing a candidate list of activities, resources, and constraints. Ideally, an automated scheduling system would perform this task and focus on maximizing the science output of the mission. The fact that the planning experts are primarily responsible for the spacecraft's safety suggests that optimizing science output will not be their primary focus. Instead, they spend much of their time focusing on spacecraft safety and health issues. While spacecraft health and safety are vitally important, the mission plan should always be optimized for science output while simultaneously guaranteeing that safety and health goals are met.

While all these features impose limitations and constraints on mission effectiveness, the current mission organization has performed reliably for many NASA missions. The cooperative autonomy model does suggest specific areas where improvements can be made and these will be the focus of the next section.

7.4.3 Improvements to Spacecraft Mission Execution

To increase mission science output in the current environment, it is necessary to insert new technologies that decrease the labor needed in building and managing spacecraft. This section will examine some technologies supporting this goal.

The science planning group is responsible for setting goals and interpreting results. It is one area where it is difficult to eliminate the large investment in human labor. However, some technologies can be inserted that will make planning efforts easier and more efficient. Groupware technologies could assist in the planning cycle by making team communication and idea-sharing more efficient. It would also lower the number of face-to-face meetings required by the staff and allow them to work on the project at times convenient to their schedule. In a similar manner, advanced data fusion, analysis, and visualization packages can assist scientists in interpreting their results. Both these technologies are currently in use and their use should be expanded.

Spacecraft sequence generation creates the series of commands necessary to achieve the objective of the high-level mission plan. The process usually engages software tools to support humans. As stated in Sect. 7.4.1, some spacecraft already have a limited amount of automation in performing a similar function. It would be a reasonably small step to move the sequence generation directly into the spacecraft. This would allow the spacecraft to control its activities and would lower the human staffing requirements.

Between science planning and sequence generation is the mission planning function. This is an area ripe for automation. Mission planning groups already use planning and scheduling software to deal with the more detailed and labor-intensive tasks. By augmenting the existing software, it would be possible to design a system that is completely automated for normal operating conditions, eventually lowering the need for human labor for mission planning. The mission planning system would initially be run on ground-based systems to facilitate monitoring and problem resolution. Figure 7.11 shows this configuration. While not shown in the figure, a human mission manager would probably monitor the mission planning system and would, if necessary, resolve problems.



Fig. 7.11. Cooperative autonomy view of spacecraft mission control



Fig. 7.12. Cooperative autonomy view of spacecraft mission control

Eventually, the mission planning function could be installed on the spacecraft, thereby enabling it to operate fully autonomously. Figure 7.12 shows this final system, where human labor is focused on creating a science plan and interpreting the results. The spacecraft converts the science plan into a mission plan and then converts the mission plan into a series of low level commands. Following execution of the low level commands, the results are evaluated and the collected data are sent to the science team for analysis and interpretation. Possibly, the team will modify the plan and redirect the spacecraft. This model does not eliminate human intervention, since a human mission manager would monitor the system and ensure all problems are properly resolved.

7.5 An Example of Cooperative Autonomy: Virtual Platform

The previous section discussed the cooperative autonomy model in view of current NASA processes and missions. NASA is also pursuing new ways to increase the science return of spacecraft while minimizing the cost of development and operations. One emerging concept is that of virtual platforms.

Virtual platforms involve the instruments of two or more spacecraft to collect data for a common science goal. Many configurations of virtual platforms are possible. In the simplest example, known as formation flying, multiple spacecraft perform their science collection while keeping a fixed position relative to one another. The fixed relative position can, in some cases, be maintained without direct communication between the spacecraft, and the cooperation in such cases is limited to merging the collected data. Simple constellations are groups of identical spacecraft that coordinate their data collection. As in formation flying, merging the data collected by the constellation enables a more complete view of the science. Advanced constellations are able to collaborate during planning phases of the mission, which allows them to allocate tasks to the most suitable spacecraft.

Complex constellations are heterogeneous mixes of different spacecraft. They share the characteristics of simple constellations, but differences in spacecraft sensors allow collections in either multiple spectra (e.g., infrared (IR) and ultraviolet (UV)), or different disciplines (e.g., earth radiation and atmospheric composition). These differences make the resulting data fusion more difficult but allow richer, augmented sets of science data. Further, in such a configuration, older preexisting spacecraft may be used in new ways not planned by the original designers of the spacecraft.

This section will now use the cooperation models previously developed to highlight issues related to the development of virtual platforms.

7.5.1 Virtual Platforms Under Current Environment

For virtual platforms to be effective, mission control must be able to select appropriate spacecraft for data collection and then task them. The current mission management organization, being designed for single platforms, does not scale well when managing multiple platforms. Figure 7.13 shows the mission management structure for a two-spacecraft virtual platform using current management techniques. Since the science planning group sets the goals for the whole virtual platform, the group is shared among all the vehicles of the virtual platform. The group also has the responsibility for fusing the data returning from all spacecraft.

Each vehicle has its own mission planning group. This group is responsible for converting the science plan into a mission plan appropriate for the specific vehicle. This is necessary because each spacecraft will have a different role. It might be possible to share human planners if the cooperating spacecraft were similar and the planning demands were modest.

The sequence generation and monitoring are very specific to each vehicle, because each vehicle has a specific role to play and sequence generation is focused upon platform-specific issues like the battery charge or damaged instruments. It is, therefore, unlikely that the human operators in these roles could easily be shared.

Given the current mission management structure, virtual platforms would make serious demands on NASA. Assume, for example, that the platforms being used have a four member science team and require three mission planners and one command sequence operator. Managing one spacecraft would, therefore, need the efforts of eight team members. Figure 7.13 represents a two-spacecraft virtual platform, which requires twelve team members. The components of Fig. 7.13 that are shaded are those that must be replicated to add additional spacecraft to the virtual platform. Therefore, a ten-spacecraft



Fig. 7.13. The mission management structure for two spacecraft in a virtual platform configuration

virtual platform would require 44 team members. Unless the science is of a very high priority, it is unlikely that NASA would support a large virtual platform system. Some of these issues can be addressed using automation and this will be examined next.

7.5.2 Virtual Platforms with Advanced Automation

Section 7.4.3 discussed how advanced automation could be used to lower the number of team members necessary to manage a single platform. These same techniques can be used to create a virtual platform architecture as shown in Fig. 7.14.



Fig. 7.14. Cooperative autonomy view of spacecraft mission control

As in the previous example, the science planning team is shared between all spacecraft that are cooperating as a virtual platform. This is where the similarity ends. Once the science plan is generated, it is communicated directly to the spacecraft. The top level planning component of each spacecraft negotiates with its counterparts on the other spacecraft to determine their individual responsibilities in the global mission. Once the negotiations are complete, each spacecraft performs its mission and returns results to the science team. In conjunction with the science team, a small staff will be required to monitor the overall virtual platform and address any problems that may occur.

This approach to virtual platforms is very attractive. Using the numbers from the previous example, this architecture only requires five team members, no matter how many spacecraft are involved in the virtual platform. This compares very favorably to the 44 staff members necessary to manage a tenspacecraft virtual platform when the current mission architecture is used.

7.6 Examples of Cooperative Autonomy

Cooperative autonomy requires many different technologies to be synthesized into a functional whole. Aspects of cooperative autonomy can be found in hundreds of projects. This section will outline several projects that incorporate one or more technologies that support the development of cooperative autonomy. The projects were selected to give the reader a cross-section of the technologies available.

New Millennium Program (NMP)

The New Millennium Program (NMP) is a NASA/Jet Propulsion Laboratory (JPL) project that will aggressively demonstrate new technologies for automation and autonomy. Though the primary thrust is technology, the project has scientific goals. NMP will fly a series of deep space and earth-orbiting spacecraft, the first of which was launched in 1998. Some of the software technologies are:

- Model-based reasoning
- Planning and scheduling architectures
- Executive architecture (performs plans)
- Fuzzy logic
- Neural networks

The DS1 spacecraft, launched in 1998, was the first to employ an on-board, autonomous system, AutoNav, to navigate to celestial bodies. About once per week throughout the mission, AutoNav was invoked [114]. The system made navigation decisions about spacecraft trajectory and targeting of celestial bodies with little assistance from ground controllers.

DS1 also used the New Millennium Remote Agent (NMRA) control architecture (Fig. 7.15) [100]. In two separate experiments, the remote agent was given control of the DS1 spacecraft. The remote agent involved an on-board mission manager that used a mission plan comprising high-level goals. A planning and scheduling engine generated a set of activities based on the goals, the spacecraft state, and constraints on spacecraft operations. The plan execution component incorporated a hybrid reactive planner and a model-based identification and reconfiguration system. The reactive planner decomposed the activities supplied by the high level planner into primitive activities, which were then executed. The model-based reasoning component used data from sensors to determine the current mode from the current spacecraft state. If a task failed, the model-based component assisted the reactive planner by using its model to act as a recovery expert and determine possible recovery strategies.



Fig. 7.15. Remote agent architecture

This project is highlighted because it is NASA's showcase for new technology, and the software technologies are truly revolutionary. The project will demonstrate substantial autonomy in space-based missions with the goal to establish a virtual presence in space. Cooperative autonomy technologies [22, 39, 43, 108, 194] could augment the DS1 autonomy architecture and further this goal.

7.6.1 The Mobile Robot Laboratory at Georgia Tech

The Georgia Tech Mobile Robot Laboratory (MRL) has been working on the fundamental science and current practices of intelligent mobile robot systems. The MRL has the goal of facilitating the technology transfer of their research results to real world problems.

Many of the MRL projects should be of interest to those attempting to set up a cooperative autonomy laboratory. The MRL has studied online adaptive learning techniques for robotic systems that allow robots to learn while they are actively involved in their operating environment. This type of learning is intended to be fast, similarity-based, and reactive. The MRL has also studied offline learning where the robot system reasons deeply about its experiences and learns as a result of this analysis. This type of learning is intended to be slower, case-based and explanation-based, deliberative, and goal-oriented.

Georgia Tech has pursued autonomous vehicle research projects supported by the Defense Advanced Research Projects Agency (DARPA). One effort mixed autonomous robot behavior with human controllability. Other research addressed multiagent systems that achieve tasks in the context of hostile environments.

Georgia Tech has developed several software packages that allow users to create robot control architectures for a specific domain and then test the control architecture in a simulated robot environment. One is written in Java and is designed to be portable. The simulation environment is compatible with off-the-shelf robotic hardware and allows the control architecture developed in the simulator to be run directly on a physical robot.

7.6.2 Cooperative Distributed Problem Solving Research Group at the University of Maine

The University of Maine's Cooperative Distributed Problem Solving Research group is centered on determining and devising the features, characteristics, and capabilities that are sufficient to bring about collaboration among groups of autonomous and semiautonomous agents toward the accomplishment of tasks. Their work has involved underwater robots, which share many of the same challenges as spacecraft:

- Operations in hostile environments
- Self-contained operations
170 7 Cooperative Autonomy

- Operations in six degrees of freedom (with neutral buoyancy or weight-lessness)
- Operations with limited communications bandwidth

The group's research has focused on intelligent control for autonomous systems and cooperative task assignments, determining what has to be communicated during cooperative problem solving, and developing a low-bandwidth conceptual language for cooperative problem solving. In one project, a collection of underwater autonomous vehicles collect data in the ocean and can create a 3-D image of the area of interest. Connected through low bandwidth acoustic modems or radio links, the vehicles coordinate their sampling activities and results reporting.

7.6.3 Knowledge Sharing Effort

The Advanced Research Projects Agency (ARPA), now named the DARPA, sponsored the Knowledge Sharing Effort (KSE). KSE developed methodologies and software for the sharing and reuse of knowledge [106] in support of communication between autonomous systems. KSE was broken up into three separate efforts.

- Knowledge Query and Manipulation Language (KQML) is a language for exchanging information and knowledge between agents [24]. It prescribes a set of performatives that represent different types of agent communication actions (like ask or tell). KQML coordinates the exchange of these performatives between agents.
- Knowledge Interchange Format (KIF) is a formal computer language designed for exchanging complex knowledge between agents [44]. It has declarative semantics and allows the agents to exchange information and describe how that information should be interpreted (the definitions of objects, the meaning of relationships, etc.).
- The final effort is the development of deep knowledge bases for domains of interest. These knowledge bases will have definitions for objects of interest, define relationships and concepts, and populate the knowledge bases with important objects.

KSE is highlighted here because it was a long-term ARPA/DARPA project to build formal mechanisms for agent communication. In building their languages and systems, the researchers had to address issues that any group of collaborating agents would have to address.

7.6.4 DIS and HLA

The military has developed a series of high quality training and simulation systems, which involve the fighting soldier in a detailed model of the area of combat. The efforts began in the early 1980s with SimNet [6, 48], which

evolved into Distributed Interactive Simulation (DIS), and then into High Level Architecture (HLA) [18].

These systems created a shared virtual environment where the combat elements (tank, plane, missile, helicopter, etc.) can see themselves and the other combatants. Each soldier sits at a station that controls an element (e.g., a tank position or the cockpit of a fighter) and the soldier's actions cause an appropriate change in the simulation of the element in the virtual world. The soldiers are given a view appropriate to their vehicles and stations and they are able to see the other combatants and the effects of their actions (a missile being fired or a tank turret being rotated). The system has been deemed so good that it has been used to test out new tactics and has been used to assist in the design of new systems by allowing different designs or tactics to be simulated and tested under simulated combat conditions.

Hundreds of individual vehicle simulations can be connected together over a distributed network using specialized protocols running over Internet protocols. These protocols support the efficient exchange of simulation information and allow all participants to experience an appropriate view of the virtual world without requiring an overwhelming amount of computation per station or overloading the network with simulation updates. Work has been directed toward building simulated forces, linking real physical hardware directly to simulated hardware, and building a virtual environment that would allow foot soldiers to engage in simulated combat.

DIS and HLA have been highlighted because they represent the high end of software testing environments for cooperative autonomy. They can support large numbers of simulated objects in a physically distributed environment using Internet protocols. They can also support the integration of real hardware with simulations. If the HLA protocols were modified to meet NASA requirements, the resulting system could allow detailed testing of proposed cooperative autonomy systems, or could allow realistic ground support stations to be integrated into the environment to test new control regimes or to train ground support personnel.

7.6.5 IBM Aglets

Many different technologies have been proposed to support agent-based programming. One system developed by IBM supports the creation of Aglets, which extends Java-based applets to create mobile software agents [82]. The Aglet toolkit helps the programmer develop autonomous agents, which can then be instantiated, cloned, moved to other computation systems, or destroyed. Implemented in Java, the Aglets have an advantage that they can run on any computation platform that supports Java, and they automatically have the many security features provided by Java. The Aglet toolkit does not focus on cooperation between Aglets, but these services could be provided by other Java classes. IBM has made the Aglet toolkit publicly available to support experimentation by others.

172 7 Cooperative Autonomy

Aglets have been discussed as one of the many possible agent architectures. While they have limited services, they are written in Java, and therefore, are portable and extensible. It would be possible to augment Aglets with, for example, KQML or Foundations of Intelligent Physical Agents (FIPA)-ACL for inter-agent communication, along with the robotic control system of Georgia Tech to build cooperating smart agents.

Autonomic Systems

NASA requires many of its future missions (spacecraft, rovers, constellations/ swarms of spacecraft, etc.) to possess greater capabilities to operate on their own with minimal human intervention or guidance [180-182]. Autonomy essentially describes independent activity toward goal achievement, but spacesystem autonomy alone is not sufficient to satisfy the requirement. Autonomicity, the quality that enables a system to handle effects upon its own internal subsystems and their interactions when those effects correspond to risks of damage or impaired function, is the further ingredient of space assets that will become more essential in future advanced space-science and exploration missions. Absent autonomicity, a spacecraft or other asset in a harsh environment, will be vulnerable to many environmental effects: without autonomic responses, the spacecraft's performance will degrade, or the spacecraft will be unable to recover from faults. Ensuring that exploration spacecraft have autonomic properties will increase the survivability, and therefore, their likelihood of success. In short, as missions increasingly incorporate autonomy (selfgoverning of their own goals), there is a strong case to be made that this needs to be extended to include autonomicity (mission self-management [160]). This chapter describes the emerging autonomic paradigm, related research, and programmatic initiatives, and highlights technology transfer issues.

8.1 Overview of Autonomic Systems

Autonomic Systems, as the name suggests, relates to a metaphor based on biology. The autonomic nervous system (ANS) within the body is central to a substantial amount of nonconscious activity. The ANS allows us as individuals to proceed with higher level activities in our daily lives [63] without having to concentrate on such things as heartbeat rate, breathing rate, reflex reactions upon touching a sharp or hot object, and so on [42, 146, 161]. The aim of using this metaphor is to express the vision of something similar to be achieved in computing. This vision is for the creation of the self-management of a substantial amount of computing functions to relieve users of low level management activities, allowing them to emphasize on the higher level concerns of the pursuit of happiness, in general, or the activity of the moment, such as playing in a soccer match, cooking a meal, or engaging in a spirited scientific debate.

The need and justification for Autonomic Systems arise from the ever increasing complexity of modern systems. A not uncommon complaint about the information technology (IT) industry identifies its inordinate emphasis on improving hardware performance with insufficient attention to the burgeoning of software features that always seem to require every possible bit of additional hardware power, neglecting other vital criteria. This has created a trillion dollar industry with consumers at the mercy of the hardware-software upgrade cycle. The consequence is a mass of complexity within "systems of systems," resulting in an increasing financial burden per computer (often measured as the TCO: total cost of ownership).

In addition to the TCO implications, complexity poses a hinderance to achieving dependability [156]. Dependability, a desirable property of all computer-based systems, includes such attributes as reliability, availability, safety, security, survivability, and maintainability [8]. Dependability was identified by both US and UK Computer Science Grand Research Challenges: "Build systems you can count on," "Conquer system complexity," and "Dependable systems (build and evolution)" [60]. The autonomic initiatives offer a means to achieve dependability while coping with complexity [156].

8.1.1 What are Autonomic Systems?

An initial reaction to the Autonomic Initiative was "is there anything new?," and to some extent this question can be justified as artificial intelligence (AI) and fault tolerant computing (FTC), among other research disciplines, have been researching many of the envisaged issues within the field of autonomic computing for many years. For instance, the desire for automation and effective, robust systems is not new. In fact, this may be considered an aspect of best-practice systems and software engineering. Similarly, the desires for systems self-awareness, awareness of the external environment, and the ability to adapt are also not new, being major goals of several fields within AI research.

What is new is AC's holistic aim of bringing all the relevant areas together to create a change in the industry's direction: selfware, instead of the hardware and software feature-upgrade cycle of the past, which created the complexity and TCO quagmire. IBM, upon launching the call to the industry, voiced the state of the industry's concerns as complexity and TCO. They presented the solution to be autonomic computing, expressed as comprising the following eight elements [63]:

- Possess system identity: detailed knowledge of components
- Self configure and re-configure: adaptive algorithms

- Optimize operations: adaptive algorithms
- Recover: no impact on data or delay on processing
- Self protection
- Aware of environment and adapt
- Function in a heterogeneous world
- Hide complexity

These eight elements can be expressed in terms of properties that a system should posses in order to constitute autonomicity [156]. These are described in Sect. 8.1.2 and elaborated upon in Sect. 8.1.3, which discusses the very constructs that constitute these properties.

8.1.2 Autonomic Properties

System autonomicity corresponds to the presence of the properties depicted in Fig. 8.1 [156]. The general properties of an autonomic (self-managing) system can be summarized by four objectives as follows:

- Self-configuring
- Self-healing
- Self-optimizing
- Self-protecting



Fig. 8.1. Autonomic computing properties tree

176 8 Autonomic Systems

which are referred to as self-chop, and four attributes as follows:

- Self-awareness
- Environment-awareness
- Self-monitoring
- Self-adjusting

Essentially, the objectives represent broad system requirements, while the attributes identify basic implementation mechanisms. Since the 2001 launch of autonomic computing, the self-* list of properties has grown substantially [169], yet this initial set still represents the general goal.

The self-configuring objective represents a system's ability to readjust itself automatically, either in support of changing circumstances or to assist in self-healing, self-optimization, or self-protection. Self-healing, in reactive mode, is a mechanism concerned with ensuring effective recovery when a fault occurs – identifying the fault and, where possible, recovering from it. In proactive mode, it monitors vital signs and attempts to predict and avoid health problems. Self-optimization means that a system is aware of its ideal performance, can measure its current performance against that ideal, and has policies for attempting improvements. It may also react to policy changes within the system as indicated by the users. A self-protecting system will defend itself from accidental or malicious external attack. This means being aware of potential threats and having ways of handling those threats [156].

In achieving self-managing objectives, a system must be aware of its internal state (self-aware) and current external operating conditions (environmentaware). Changing circumstances are detected through self-monitoring, and adaptations are made accordingly (self-adjusting) [156]. Thus, a system must have knowledge of its available resources, its components, their desired performance characteristics, their current status, and the status of inter-connections with other systems, along with rules and policies of how these may be adjusted. In the broad view, the ability to operate in a heterogeneous environment will require the use of open standards to enable global understanding and communication with other systems [63].

These mechanisms are not independent entities. For instance, recovery from a successful attack will include self-healing actions and a mix of selfconfiguration and self-optimization: self-healing to ensure dependability and continued operation of the system, and self-configuration and self-optimization to increase self-protection against similar future attacks. Finally, these selfmechanisms should ensure that there are minimal disruptions to the pursuit of system goals.

There are two main perceived approaches (Fig. 8.1) considered to be the means for autonomic computing to become a reality [146]:

- Engineer Autonomicity
- Learn Autonomicity

"Engineer Autonomicity" has an implied Systems and/or Software Engineering view, under which autonomic function would be engineered into the individual systems. "Learn Autonomicity" has an implied AI, evolutionary computing, and adaptive learning view, where the approach would be to utilize algorithms and processes to achieve autonomic behavior. However, both approaches rely on each other in achieving the objectives set out in Automatic Computing. Autonomic Computing may prove to require a greater collaboration between the intelligence-systems research and systemand software-engineering fields to achieve the envisaged level of adaptation and self-management within the autonomic computing initiative.

8.1.3 Necessary Constructs

Considering these autonomic properties, the key constructs and principles that constitute an Autonomic Environment are:

- Selfware; Self-*
- AE = MC + AM
- Control Loop; Sensors+Effectors
- $AE \leftrightarrow AE$
- Selfware; Self-*: The principle of selfware (self-managing software and firmware) and the need for self-* properties were discussed in the previous sections.
- AE=MC+AM: Figure 8.2 represents a view of an architecture for an autonomic element, which consists of the component to be managed and the autonomic manager [69, 154]. It is assumed that an autonomic manager (AM) is responsible for a managed component (MC) within a self-contained autonomic element (AE). This AM may be designed as part of the component or may be provided externally to the component, as an agent, for instance. Interaction will occur with remote AMs (e.g., through an autonomic communications channel) through virtual, peer-to-peer, client-server [11], or grid [33] configurations.
- Control Loop, Sensors+Effectors: At the heart of any autonomic system architecture are sensors and effectors [42]. A control loop is created by monitoring behavior through sensors, comparing this with expectations (historical and current data, rules, and beliefs), planning what action is necessary (if any), and then executing that action through effectors [68]. The control loop, a success of manufacturing science for many years, provides the basic backbone structure for each system component [41]. IBM represents this self-monitor-self-adjuster control loop as the monitor, analyze, plan, and execute (MAPE) control loop. The monitor and analyze parts of the structure process information from the sensors to provide both self-awareness and an awareness of the external environment. The plan and execute parts decide on the necessary self-management behavior that will be executed through the effectors. The MAPE components use the



Fig. 8.2. Autonomic element (AE) consisting of autonomic manager (AM) and managed component (MC)

correlations, rules, beliefs, expectations, histories, and other information known to the autonomic element, or available to it through the knowledge repository within the AM.

 $AE \leftrightarrow AE$: The Autonomic Environment requires that autonomic elements, and in particular, AMs, communicate with one another concerning self-* activities to ensure the robustness of the environment. Figure 8.2 views an AE with the additional concept of a pulse monitor (PBM). This is an extension of the embedded systems heart-beat monitor (HBM), which safeguards vital processes through a regular emitting of an "I am alive" signal to another process, with the capability to encode health and urgency signals as a pulse [148]. Together with the standard event messages on the autonomic communications channel, this provides not only dynamics within autonomic responses, but also multiple loops of control, such as reflex reactions, among the AMs [159].

8.1.4 Evolution vs. Revolution

In recognition of, first, the need for differing levels of human involvement, and second, the reality that the overarching vision of autonomic computing will not be achieved overnight, autonomic computing maturity and sophistication have

been categorized into five "stages of adoption" [10, 20, 69]: Basic, Managed, Predictive, Adaptive, and Autonomic.

Assessing where a system resides within these autonomic maturity levels is not necessarily an easy task. Efforts are underway to define the required characteristics and metrics [88]. The overall AC maturity is established from a combination of dimensions forming a natural continuum of autonomic evolution [80], such as increasing functionality (manual, instrument-andmonitor, analysis, closed-loop, to closed-loop-with-business-priorities) and increasing scope (subcomponents, single-instances, multiple-instances-same type, multiple-instances-different types, to business-systems) [80]. Since assessment is becoming even more complex, efforts are currently underway to automate the assessment process itself [41,130]. These efforts imply that the autonomic computing initiative is following an evolutionary path.

8.1.5 Further Reading

The best starting point for further reading is IBM's "call to arms" launch of the initiative [63], the autonomic "vision" paper [79], and the "dawning" paper [42], as well as news about the autonomic initiative [107].

Since the launch of AC, IBM has released various white papers. The general concepts within these have essentially been brought together into a book published by IBM Press [98]. This book covers IBM's view of Autonomicity and how it strategically fits within their other initiatives (such as On-Demand).

Origins of some of the IBM thinking on autonomic computing can be attributed to the active middleware services (AMS) community, where their fifth workshop in Seattle in 2003 became the Autonomic Computing Workshop [104] and evolved, with IBM's backing, into the Autonomic Conference (New York 2004) [72]. The early focus at this stage was very much on its roots, i.e., middleware, infrastructures, and architectures. Other Autonomic workshops include the Workshop on AI for Autonomic Computing, Workshop on Autonomic Computing Principles and Architectures, Workshop on the Engineering of Autonomic Systems, Almaden Institute Symposium: Autonomic Computing, Workshop on Autonomic Computing Systems, and the Autonomic Applications Workshop; and related workshops such as the ACM Workshop on Self-Healing, Adaptive and self-MANaged Systems (SHAMAN), and the ACM Workshop on Self-healing Systems (WOSS).

Special issue journals are also beginning to appear [53, 170]. The papers in [53] generally cover engineering topics such as mirroring and replication of servers, software hot swapping, and database query optimization. Those in [170] strongly represent autonomic efforts for the grid, web, and networks. Appreciating the wider context of autonomic computing, the boiling pot that influenced AC can be found in other research initiatives such as Recovery Oriented Computing [19].

8.2 State of the Art Research

It has been highlighted that meeting the grand challenge of Autonomic Systems will involve researchers in a diverse array of fields, including systems management, distributed computing, networking, operations research, software development, storage, AI, and control theory, as well as others [42]. There is no space here to cover all the excellent research underway, so this section will discuss a selection of the early reports in the literature of state-of-the-art efforts in AC [147].

8.2.1 Machine Design

A paper in [70] discusses affect and machine design [101]. Essentially, it supports those psychologists and AI researchers who hold the view that affect (and emotion) is essential for intelligent behavior [139,140]. It proposes three levels for the design of systems:

- 1. Reaction: The lowest level, where no learning occurs, but where there is an immediate response from the system to state information coming from sensory systems.
- 2. Routine: The middle level, where largely routine evaluation and planning behaviors take place. The system receives inputs from sensors as well as from the reaction level and reflection level. At this level of assessment, there are results in three dimensions of affect and emotion values: positive affect, negative affect, and (energetic) arousal.
- 3. Reflection: The top level of the system receives no sensory input and has no motor output; it receives all inputs from below. Reflection is a metaprocess where the mind deliberates about the system itself. Operations at this level look at the system's representations of its experiences, its current behavior, its current environment, etc.

Essentially, the reaction level sits within the engineering domain, monitoring the current state of both the machine and its environment, and produces rapid reaction to changing circumstances. The reflection level may reside within an AI domain, utilizing its techniques to consider the behavior of the system and learn new strategies. The routine level may be a cooperative mixture of both the reactive and reflection levels.

8.2.2 Prediction and Optimization

A method known as Clockwork provides predictive autonomicity by regulating behavior in anticipation of need. It involves statistical modeling, tracking, and forecasting methods [127] to predict need and is now being expanded to include real-time model selection techniques to fulfill the self-configuration element of autonomic computing [128]. This work includes probabilistic reasoning, and prospectively, should be able to benefit from invoking genetic algorithms for model selection.

Probabilistic techniques such as Bayesian networks (BNs) discussed in [50] are also central in research into autonomic algorithm selection, along with self-training and self-optimizing [50]. Re-optimization of enterprise business objectives [4] can be encompassed by the breadth and scope of the autonomic vision through such far-reaching work combined with AI techniques (machine learning, Tabu search, statistical reasoning, and clustering analysis).

As an example, the application "Smart Doorplates" assists visitors to a building by locating individuals who are not in their offices. A module in the architecture utilizes probabilistic reasoning to predict the next location of an individual, which is reported along with his/her current location [173,174].

8.2.3 Knowledge Capture and Representation

Vital to the success of Autonomic Systems is the ability to transfer expert human knowledge about system management and configuration to the software managing the system. Fundamentally, this is a knowledge-acquisition problem [85]. One current research approach is to capture the expert's actions automatically (keystrokes and mouse movements, etc.) when performing on a live system, and dynamically build a procedure model that can execute on a new system and repeat the same task [85]. Establishing a collection of traces over time should allow the approach to develop a generic and adaptive model.

The Tivoli management environment approaches this problem by capturing in its resource model the key characteristics of a managed resource [77]. This approach is being extended to capture the best practices information into the common information model (CIM), through descriptive logics at both the design phase and the deployment phase of the development lifecycle [83]. In effect, the approach captures system knowledge from the creators, ultimately to perform automated reasoning when managing the system.

8.2.4 Monitoring and Root-Cause Analysis

Event correlation, rule development, and root-cause analysis are important functions for an autonomic system [155]. Early versions of tools and autonomic functionality updates to existing tools and software suites in this area have recently been released by IBM [41] through their AlphaWorks Autonomic Zone website. Examples include the Log and Trace Tool, the Tivoli Autonomic Monitoring Engine, and the ABLE rules engine.

The generic Log and Trace Tool correlates event logs from legacy systems to identify patterns. These patterns can then be used to facilitate automation or support debugging efforts [41]. The Tivoli Autonomic Monitoring Engine essentially provides server-level correlation of multiple IT systems to assist with root-cause analysis and automated corrective action [41]. The ABLE rules engine can be used for more complex analysis. In effect, it is an agent-building learning environment that includes time series analysis and Bayes classification among others. It correlates events and invokes the necessary action policy [41].

It has been noted that correlation, rule discovery, and root-cause analysis activities can benefit from incorporating Bayesian Networks [153], either in the rule discovery process or in the actual model learning to assist with self-healing [150]. Large-scale server management and control has also received similar treatment. Event logs from a 250-node large-scale server were analyzed by applying a number of machine-learning algorithms and AI techniques to establish time-series methods, rule-based classification, and BN algorithms for a self-management and control system [129].

Another aspect of monitoring and root-cause analysis is the calculation of costs, in conjunction with the self-healing equation in an autonomic system. One approach utilizes naive Bayes for cost-sensitive classification and a feedback approach based on a Markov decision process for failure remediation [89]. The argument is easily made that the autonomic system involves decisions and decisions involve costs [25]. This naturally leads to work with agents, incentives, costs, and competition for resource allocation and extensions thereof [25, 105].

8.2.5 Legacy Systems and Autonomic Environments

Autonomic Systems is arguably widely believed to be a promising approach to develop new systems. Yet organizations continue to have to deal with the reality of either legacy systems or building "systems of systems" composed of new and legacy components that involve disparate technologies from numerous vendors [76]. Work is currently underway to add autonomic capabilities to legacy systems in areas such as instant messaging, spam detection, load balancing, and middleware software [76].

Generally, the engineering of autonomic capabilities into legacy systems involves providing an environment for monitoring the system's sensors and providing adjustments through effectors to create a control loop. One such infrastructure is Kinesthetics eXtreme (KX). It runs a lightweight, decentralized, easily integratable collection of active middleware components tied together via a publish-subscribe (content-based messaging) event system [76]. Another tool, called Astrolabe, may be used to automate self-configuration and monitoring and control adaptation [14]. The AutoMate project, incorporating ACCORD (an autonomic component framework), utilizes the distributed interactive object substrate (DIOS) environment to provide mechanisms to directly enhance traditional computational objects/components with sensors, actuators, rules, a control network, management of distributed sensors and actuators, interrogation, monitoring, and manipulation of components at runtime through a distributed rule-engine [3,90,97].

8.2.6 Space Systems

As discussed earlier, with the increasing constraints on resources and the greater focus on the cost of operations, NASA and others have started to utilize adaptive operations and move toward almost total onboard autonomy in certain classes of mission operations [176, 195]. Autonomy provides selfgovernance, giving responsibility to the agents within the system to meet their defined goals. Autonomicity provides self-management in addition to self-governance as essential to a system's ability to meet its own functional goals. There is also a shared responsibility to ensure the effective management (through self-* properties) of the system, which may include responsibilities beyond the normal task-oriented goals of an individual agent. For instance, monitoring another agent's health signs to ensure self-protection, self-healing and self-configuration, and/or self-optimization activities take place as needed. Autonomic computing, then, can be identified as a key technology [27, 66, 146, 151] for future NASA missions, and research is paying the way for incorporation of both autonomicity and autonomy [182]. These will be discussed in more detail later in Part III of this book.

8.2.7 Agents for Autonomic Systems

Agents, as autonomous entities, have the potential to play a large role in Autonomic Systems [49, 63, 102, 105, 168, 169], though, at this stage, there are no assumptions that agents must necessarily be used in an autonomic architecture. However, as in complex systems, there are substantial arguments for designing a system with agents [75]. Agents can help provide inbuilt redundancy and greater robustness [67], as well as help retrofit legacy systems with autonomic capabilities [76]. With reference to work previously mentioned, a potential contribution of agents may come from environments that require either learning, rules, and norms, or agent monitoring systems.

8.2.8 Policy-Based Management

Policy-based management becomes particularly important with the future vision of autonomic computing, where a manager may simply specify the business objectives and the system will make it so – in terms of the needed information and communications technology (ICT) [94]. A policy-based management tool may reduce the complexity of product and system management by providing a uniform cross-product policy definition and management infrastructure [41].

8.2.9 Related Initiatives

Other self-managing initiatives include:

- Cisco (adaptive network care) [71]
- HP (adaptive infrastructure) [54]
- Intel (proactive computing) [187]
- Microsoft (dynamic systems initiative) [95]
- Sun (N1) [164]

All of these initiatives are concluding that the only viable long-term solution is to create computer systems that manage themselves.

The latest related research initiative is autonomic communications [149, 150]. An European Union brainstorming workshop in July 2003 to discuss novel communication paradigms for 2020 identified "Autonomic Communications" as an important area for future research and development [142]. This can be interpreted as further work on self-organizing networks, but is undoubtedly a reflection of the growing influence of the autonomic computing initiative.

Autonomic communications has the same motivators as the autonomic computing concept, except it has a focus on the communications research and development community. Research in autonomic communications pursues an understanding of how an autonomic network element's behaviors are learned, influenced, or changed, and how this effects other elements, groups, and networks. The ability to adapt the behavior of the elements was considered particularly important in relation to drastic changes in the environment, such as technical developments or new economic models [142]. This initiative has now evolved into a major European research program, known as "situated and autonomic communications" (SAC) [141].

8.2.10 Related Paradigms

Related initiatives, as in perceived future computer paradigms, include grid computing, utility computing, pervasive computing, ubiquitous computing, invisible computing, world computing, ambient intelligence, ambient networks, and so on. The driving force behind these future paradigms of computing is the increasing convergence between the following technologies:

- Proliferation of devices
- Wireless networking
- Mobile software

Weiser first described what has come to be known as ubiquitous computing [188] as the move away from the "dramatic" machine (where hardware and software was to be so exciting that users would not want to be without it) toward making the machine "invisible" (so embedded in users' lives that it would be used without thinking or would be unrecognized as computing).

Behind these different terms and research areas lie three key properties:

- Nomadic
- Embedded
- Invisible

In effect, this may lead to the creation of a single system with (potentially) billions of networked information devices. All of these next generation paradigms, in one form or another, will require an autonomic-self-managinginfrastructure to provide the successful reality of this envisaged level of invisibility and mobility.

Currently, and for the foreseeable future, the majority of users access computing through personal devices. Personal Computing offers unique challenges for self-management in itself due to its multidevice, multisituation, and multiuser nature. Personal autonomic computing is much less about achieving optimum performance or exploiting redundancy (as in AC) and more about simplifying use of the equipment and the associated services [10,11]. Thus, it is particularly relevant to deal with the move toward a nomadic, embedded, and invisible computing future [152, 157].

8.3 Research and Technology Transfer Issues

The challenge of autonomic computing requires more than the re-engineering of today's systems. Autonomic computing also requires new ideas, new insights, and new approaches.

Some of the key issues that will need to be addressed are:

- Trust: Even if the autonomic community manages to "get the technology right," the trust of the user will be an issue in terms of the user's handing over control to the system. AI and autonomous agent domains have suffered from this problem. For instance, neural networks (due to concerns over the "black-box" approach) and a number of AI techniques (due to their inherent uncertainty) are often not adopted. Rule-based systems, even with all their disadvantages, often win adoption, since the user can trace and understand (and thus implicitly trust) them [153]. Note that even within autonomic computing and autonomic communications, the bulk of the literature assumes rules will be used instead of other, less brittle and more adaptable stochastic AI approaches.
- Economics: New models of reward will need to be designed. Autonomy and Autonomicity may derive another self-* property: selfishness. For instance, why would an AE perform an operation, e.g., relay information, for another AE that was outside its organization and did not affect or benefit from it? In particular, if it was operating within a mobile (battery powered) environment and to do so, incurred personal cost, performing the operation for the outside unit could shorten its useful life or make it necessary to recharge earlier.

Standards: The overarching vision of autonomic computing will only be achievable through standards, particularly for communicating between AEs. Like the agent community standardizing on a communications protocol, AEs also need a protocol standard so that they can be added to a system and be able to communicate immediately. Besides, agile ways to define these communications are needed, for which a key enabler would be the self-defining property.

It has been expressed that in AC's initial deployment take-off, many researchers and developers have zeroed in on self-optimization because it is perceived as easier to translate into technology transfer [41]. Essentially, this focus on optimization from the four self-chop attributes may be considered to be going against the grain of technology trends (toward ever faster machines), as such fine-grained optimization is not necessarily a major concern [41]. For autonomic computing to succeed in the longer term, the other self-* attributes must be addressed equally and in an integrated fashion.

As well as addressing complexity, autonomic computing also offers the promise of a lower TCO and a reduced maintenance burden as systems become self-managing. Achieving this vision will likely make substantial demands on legacy maintenance budgets in the short-term as autonomic function and behavior are progressively designed into systems.

Achieving the overarching vision of Autonomic Systems will require innovations in systems and software engineering, as well as collaboration involving many other diverse fields. Early R&D presented in this chapter highlights the momentum that is developing on a broad front to meet the vision. The NASA community, with its increasing utilization of autonomy in missions, will only benefit from the evident paradigm shift within computing that brings Autonomicity into the mainstream.

Applications

Autonomy in Spacecraft Constellations

In this chapter, we discuss the application of the autonomy technologies considered in previous chapters to spacecraft constellations. The needs of constellations that can be supported by onboard autonomy are described along with the enhancements attainable by constellation missions through the application of onboard autonomy. A list of hypothetical constellation mission types is also posed and a list of governance concepts is then presented in relation to the degree of central control being exercised on the constellation. Finally, the chapter discusses mobile agent concepts to support autonomic constellations.

9.1 Introduction

As described in Chap. 7, spacecraft constellations are organized into virtual platforms that appear as a single entity to the ground. They are often flown in formation so that different spacecraft can view science phenomena from a different perspective, or view contiguous areas at the same time. Simple constellations are groups of identical spacecraft that coordinate their data collection and merge the collected data to create a more complete view of the science. Complex constellations are heterogeneous mixes of unlike spacecraft. They share the characteristics of simple constellations, but may comprise different types of spacecraft and/or have different instruments. These differences in spacecraft and instruments make the resulting data fusion more difficult, but allows richer sets of science data to be collected. This configuration also allows older, preexisting spacecraft to be used in new ways not thought of by the original designers of the spacecraft [26].

An example of a NASA constellation is the ST5 mission. The ST5 mission [171], launched in March 2006, has three identical spacecraft that fly in a "string of pearls" formation (Fig. 1.1), utilizing a single uplink/downlink to the ground station.

As in other mission types discussed earlier, the motivations for improved autonomy in constellations arise from (among other things) resource constraints pertaining to onboard processor speeds, memory, spacecraft power, etc. Even though onboard computing power will increase in the coming years, the resource constraints associated with space-based processes will continue to be a major factor that needs to be considered when dealing with, for example, agent-based spacecraft autonomy. The realization of "economical intelligence," i.e., constrained computational intelligence that can reside within a process under severe resource constraints (time, power, space, etc.), is a major goal for future missions such as *nano-sat constellations*, where resources are even more constrained due to their small size.

Like other missions, satellite constellation missions can have a wide range of characteristics. Future missions may vary in their data rate and total data volume. Orbits may range from low earth orbits to very elliptical orbits with multiday periods. Air-to-ground protocols may vary, and the satellites themselves may be low-cost with low autonomy or may be sophisticated with a high level of onboard self-management. Traditional ground-support systems designed for single satellite support may not efficiently scale up to handle large constellations. The interested reader may refer to [15, 64, 143, 166, 190] for additional information on the challenges of spacecraft constellations.

Table 9.1 summarizes current and future types of constellation, their applications, some of the critical distinctions between the applicable mission models, and various relevant issues. To begin to address the implicit challenges, new approaches to autonomy need to be developed for constellations. As discussed in Chap. 4 relative to the agent concept testbed (ACT) proto-type, a possible two-step approach for achieving constellation autonomy is as follows:

- 1. Develop a community of surrogate ground-based agents representing the satellites in the constellation. This will enable the mission to establish, in a prototype environment, the centralized and distributed agent behaviors that eventually will be used in space.
- 2. Migrate the surrogate agents to the space-based satellites on a gradual or as-needed basis. This step is referred to as *progressive autonomy*.

This chapter will focus on Step 1, and at the end, present some ideas relating to Step 2. First we present a brief overview of constellations, reasons for using constellations, and the associated challenges in developing them. These challenges will motivate the agent-based technology discussion in relation to the goal of achieving autonomy in constellations (Fig. 9.1).

9.2 Constellations Overview

Constellations have the potential to provide the data that are needed to yield greater scientific insight and understanding into the cause-and-effect processes that occur in a region. As discussed in Chap. 6, constellations can

Туре	Application	Typical design/ manufacture	Data acquisition	Operations	
Simple	University	Very low cost.	Not a major	Extremely low	
(varied	sponsored.	Minimally	issue. Low rate.	cost. University	
number of	Corporate	space-rated	May operate	level	
satellites)	R&D	components	at amateur		
			radio frequencies		
Cluster:	Coordinated	Complex.	Not a major	Similar to single	
Cluster II (4),	science.	Satellite	issue. Typically	large satellite.	
Magnetospheric	Virtual	crosslinks.	high rate due to	Multiple	
multiscale (5)	telescopes.	Extensive	science mission,	satellites	
	Stereo	testing	but number of	performing	
	imaging	required. High	satellites is	a coordinated	
		redundancy	limited or	function. Added	
		within	downlink access	effort for	
		satellites	can be controlled	mission	
Coverage-	Commercial	Satellites	Large number	May involve	
Constellation:	phone/paging/	operate	of satellites	hundreds of	
Globalstar (48),	Internet	independently,	using many	passes per day.	
Orbcomm (36),	systems. Earth	designed for	ground sites	Ideal for	
TIROS (5) ,	(or planetary)	mass production	concurrently.	automation, as	
NASA	observation	with limited	Dedicated	there are many	
NanoSat (100)	(multi-point	redundancy,	antenna sites	nearly identical	
	data	high duty cycle	may be needed	passes. Space	
	collection,		due to identical	comm architecture	
	broad survey		satellites working	may be needed to	
	or coverage)		continuously	be fully networked	
Military/	Inspection,	New concepts are	Only a few	Mostly orbit/	
Tactical:	imagery	for very small,	satellites active	maneuver and	
XSS-10,		low-cost, mass-	at a time. May	data-acquisition	
ESCORT,		produced spacecraft	use portable	activity.	
Orbital Express		with no redundancy	data acquisition	Data are	
		and minimal	sites. May have	for immediate	
		mission durations	a video downlink	use only. No	
			plus minimal	long-term	
			status info.	trending, etc.	

 ${\bf Table \ 9.1.} \ {\rm Current \ and \ future \ types \ of \ constellation \ missions \ and \ possible \ issues}$

possess significant, and perhaps obvious, advantages over using just one or two spacecraft. For example, NASA's proposed magnetotail constellation (MC) mission, which will use a fleet of 30+ nanospacecraft, would offer space-physics scientists the ability to perform 100 concurrent observations over the magnetosphere, allowing conditions and events recorded to be correlated spatially and temporally.

Constructing and launching constellations will introduce many new and significant challenges. For example, building, launching, and then properly deploying as many as 100 spacecraft housed on one launch vehicle into their required orbits will require the development and demonstration of new spacecraft control solutions so that mission operations costs associated with supporting a constellation comprising a large number of spacecraft do not spiral out of control.



Fig. 9.1. The agent concept testbed (ACT) consists of a community of cooperating agents each of which is component-based (from Chap. 4)

There are a number of implementation issues that are unique to spacecraft constellations. Four examples presented below provide some insight into the challenges that will undoubtedly confront aerospace hardware and software engineers in launching, deploying, and then routinely operating constellations:

- Monitoring engineering telemetry data from one spacecraft is a routine task for mission operations personnel and the ground system computer hardware and software systems. However, responding to time-critical events and identifying, evaluating, and quickly resolving spacecraft subsystem anomalies can frequently be challenging for humans and computers alike. Effectively monitoring and reacting to conditions reported by the telemetry data from 100 identical spacecraft without also incurring a concomitant and potentially significant increase in staff and ground equipment will be a major challenge.
- Spacecraft that compose a constellation still will need to communicate with the ground system. Commands must be uplinked to the spacecraft, engineering health and safety telemetry data must be transmitted to the missions operations center, and payload data must be returned to the science community for ground-based processing and product distribution. Available (and limited) ground resources (e.g., spacecraft tracking stations, communications networks, and computing resources) will need to be scheduled and managed so that realistic contact plans can be created to support forward and return link telemetry processing for constellations with large numbers of spacecraft. Potentially, advanced space networking technologies will lead to more efficient networked communications capabilities,

partially offsetting the need for many direct-to-spacecraft communications paths from ground antennas.

- Some constellation missions may require that the spacecraft communicate with one another for science or formation purposes. One spacecraft may need to broadcast information to many other spacecraft in its vicinity. Alternatively, one spacecraft may need to communicate with another spacecraft in the constellation, for example, to cue it so that the second spacecraft can record an event that the first could not. But these spacecraft may be located in orbital planes where they are rarely if ever in direct line of sight of each other. Multihop networked inter-satellite communications may be necessary to synchronize operations of the entire constellation, or just a subset of it.
- Trend analysis is an important element for any spacecraft mission. It helps the mission operations staff determine whether a failure may be imminent so that switchover to backup or redundant subsystems can be performed, or if necessary, to have the spacecraft enter safemode until the problem can be identified and a corrective course of action implemented. Greater automation in ground data processing will be required to support this. Perhaps data mining techniques that are presently implemented for terrestrial databases and e-commerce applications may provide solutions for consideration and adaptation to this new problem.

9.3 Advantages of Constellations

A wide variety of missions could be best implemented with constellations of satellites working together to meet a single objective. Reasons cited for using constellations include lower mission costs, the need for coordinated science, special coverage or survey requirements, and the need for quick-reaction tactical placement of multiple satellites. The following discusses these in more detail.

9.3.1 Cost Savings

The cost of producing spacecraft for a constellation and getting them to orbit may actually be lower than traditional "one of a kind" satellites that use a dedicated launch vehicle. With a traditional satellite, system reliability requirements force a high level of component protections and redundancy, which leads to higher overall weight and launch costs. Due to their size and/or weight, a dedicated launch is often required for these missions. With a constellation, system reliability can be met by having spare satellites. The use of per-satellite redundancy can be significantly reduced. In some cases, it may be practical to use lower-rated components at a much lower cost combined with an on-orbit sparing plan. Additional savings could be obtained through the use of assembly-line production techniques and coordinated test plans, so that the satellites could basically be mass-produced. With a reduced size and weight, new options would be available for launches: lower cost launch vehicles, multiple satellites of the constellation launched on one launch vehicle, and piggy-back launch slots where launch costs are shared with another mission.

9.3.2 Coordinated Science

A constellation of as few as two satellites could be used to perform coordinated science. For example:

- Storms and other phenomena observed from multiple angles could be used to generate 3-D views
- Satellites with a wide spatial separation could be used for parallax studies of distant objects
- A cluster of satellites flying in formation and working together could form a virtual lens (or mirror) hundreds of miles across to achieve unprecedented resolution for observations of astronomical objects

The currently predominant application of satellite constellations aims to extend area coverage. Low earth orbiting constellations such as Globalstar use dozens of satellites to provide continuous global or near-global coverage. The global positioning system (GPS) uses a constellation to provide global coverage and spatial diversity of the multiple satellites in view. Earth imaging missions can use multiple satellites to shorten the time between successive observations of the same area, and can coordinate observations so that dynamic phenomena (hurricanes, earthquakes, volcanic eruptions, etc.) receive augmented attention by additional members of the constellation.

Military applications for constellations include earth observation, weather, and equipment resource monitoring. In the future, it may be possible to launch very small satellites with a very specific purpose and a very short mission duration. The satellites could be produced by the hundreds and launched as needed. The "constellation," at any point in time, would include those satellites currently performing their intended function.

9.4 Applying Autonomy and Autonomicity to Constellations

With the above discussion as motivation, the following section describes how autonomy could be applied in constellation ground control systems and in constellation spacecraft themselves to overcome the issues mentioned above. Finally, the goal of achieving autonomicity in constellations is discussed.

9.4.1 Ground-Based Constellation Autonomy

Figure 4.7 (Sect. 4.3.4, p. 89) shows a high-level representation of a constellation simulation of ground-based autonomy for a constellation of four satellites. In this simulation, a number of agents are connected to an environment in which the ground control systems and satellites [177, 183] are simulated. In the simulation, the satellites are in orbit collecting magnetosphere data. The simulation environment propagates the orbits based on ideal conditions. Faults can also be inserted into the telemetry stream to simulate an anomaly.

The group of surrogate spacecraft agents, as a major component of the ground-based community, maintains an awareness of the actual physical constellation. The surrogates act on behalf of their respective spacecraft in status monitoring, fault detection and correction, distributed planning and scheduling, and spacecraft cooperative behaviors (as needed).

A next phase in the evolution of the above ACT scenario will be to have communities of agents, each associated with a particular spacecraft in the constellation. Each of these communities would have specialist subsystem agents that would monitor the various subsystems of the spacecraft and cooperate with one another in the handling of anomalous situations. An overall coordinator, or spacecraft agent, would lead the community and represent the spacecraft to ground controllers. It would also represent the spacecraft to other spacecraft agents in the constellation community for activities such as distributed planning and scheduling, and other forms of collaboration.

In the context of spacecraft constellations, the ground-based group of surrogate agents illustrates two major themes in our discussion: (1) surrogate agents can indeed support the concept of constellation autonomy in a meaningful way, and (2) having a ground-based community of surrogate agents allows developers and users (controllers) to gain confidence and trust in the approach.

9.4.2 Space-Based Autonomy for Constellations

Constellation autonomy (as opposed to a single spacecraft autonomy) corresponds to an intrinsic property of the group. Constellation autonomy would not apply to a fleet of spacecraft where each is operated without reference to the others (e.g., a tracking and data relay satellites (TDRS) servicing the communications needs of many user spacecraft). However, if each satellite manages itself in the common external environment so as to maintain a group functionality relative to the environment, even when the members of the group do not communicate directly with one another, and the group accomplishes the end purpose of the constellation, then the constellation can be considered as a self-managing one, i.e., autonomous. This is considered justifiable, from the mission perspective, because the group accomplishes the end result as a system even though its members do not intentionally interchange information. We make this distinction because of the question of autonomy and viability. As long as the constellation produces, i.e., delivers its end purpose on its own without any external support, it is viable and autonomous. Some members of the group may fail, but the remainder will continue to produce, thereby maintaining the essential nature of the constellation.

The second and final step (progressive autonomy) in the proposed overall plan to realize space-based autonomy is to migrate the spacecraft surrogate community of agents to the actual spacecraft. As discussed in previous chapters, this is a nontrivial step. A major step in the direction of actual onboard spacecraft autonomy is to have the agent community demonstrate its correctness in actual ground-based spacecraft control centers [178, 184]. This is discussed in Sect. 9.6.

There are many issues that need to be addressed before this becomes a reality. Some of the major issues are as follows:

- Adaptation to resource constraints. As an example, a spacecraft subsystem agent must be able to exist and operate within the microprocessor associated with the subsystem. This is where the concept, which we call "economical intelligence," comes into play. Reasoning code and knowledge and information structures and management need to be "optimized" in order to function properly in the resource-constrained environment of a spacecraft subsystem microprocessor.
- Integration with existing subsystem autonomy. As discussed earlier, most spacecraft subsystems already have a degree of autonomy built into their operations. This is realized usually through the use of expert systems or state-based technologies. A subsystem agent should be able to take advantage of the existing capability and build upon it. The existing capability would become an external resource to the subsystem agent that would be used to realize a higher level of autonomy for the subsystem. The agent would need to know about the external resource and how to use it, i.e., factor its information into its reasoning process.
- *Real-time activity.* Most situations experienced by a spacecraft require real-time attention. If the situation is not readily handled by built-in sub-system autonomy, the associated subsystem agent will need to respond in real-time. This will require the agent to have a working reflexive behavior.

9.4.3 Autonomicity in Constellations

A step beyond an autonomous constellation is an autonomic constellation – a collective of autonomous agents that are self-governing and learn individual sequences of actions so that the resultant sequence of joint actions achieves a predetermined global objective. As discussed earlier, this approach is particularly useful when centralized human control is either impossible or impractical, such as whenever timely and adequate communication with humans is impractical. Constellations controlled by in situ "intelligent" spacecraft, in

comparison with the more common externally controlled constellation, can have the following advantages:

- 1. In locations not reachable in a timely manner through human contact:
 - (a) Initiating and/or changing orders automatically.
 - (b) Evaluating and summarizing global health status, and therefore, being able to assess the constellation's ability to conduct a specific exercise, and further, being able to engage in self-protection and selfreconfiguration, among other self-* functions to ensure the survival and viability of the entire constellation.
 - (c) Providing summary status of the entire constellation.
- 2. On-the-job training or programming:
 - (a) Provoking a new mode of behavior on the part of constituent satellites by observing operators and the environment [185].

The role of autonomicity in constellations depends on the needs of the satellites individually and collectively, and on the needs of mission control. For example, how robust an autonomic function on any single satellite must be would depend on such things as the proximity to standard tracking and telecommunications facilities, the urgency of data and command access by the ground, and for survey missions, the area of simultaneous coverage needed by satellites in the constellation. Availability of communication channels is another factor, whether it be the timing/accessibility and individual channel bandwidth capacity or, in close formation flying, frequency separation of channels. Of course, there are various ways of getting around mutual interference constraints, such as limiting the individual contact events to a single communications link at a time when in close formation, or by compressing the bandwidth requirements, as was mentioned previously; a fully networked inter-spacecraft communications architecture involving multihop packet routing also represents an alternative approach for some types of mission. But providing an autonomic function that can address the constraints of a system given the current situation in a mission provides a much more flexible and reusable solution than can be solved by single-point solutions.

The autonomicity of constellation governance influences the spacecraft inter-connectivity design and ground-connectivity design (i.e., human control) in much the same ways as does the individual satellite's subsystem control structure. However, the geometry, scale, and desired performance of the constellation as an integral entity come into play adding additional design complexity. As a basic consideration, the level of interdependency among constellation members is a factor influenced strongly by mission class, e.g., what type of payload is being carried.

In certain types of mission (for surveillance, analysis, monitoring, etc.), the requirements for accuracy and speed of data or event notification are becoming more demanding. In addition, the resolution requirements for imaging systems require ever larger apertures, and hence, larger instrument sizes. Formation flying concepts have been identified as the canonical means for achieving very large apertures in the space environment. (It should be noted that, as with other types of constellation mission, formation flying presents an opportunity to create a synergistic system (e.g., an imaging system) where the members of the group operate cooperatively to give rise to group capabilities that no single platform could provide by itself. Further, as with other types of constellation mission, the mission can be designed, in many cases, so that the loss of a member leaves the remainder of the group functioning. In the alternative mission design, based on using a single large spacecraft, the loss of the spacecraft means losing the entire mission.)

In a representative formation-flying concept for an astronomical observatory, the formation itself creates the effect of a large "instrument" whose aperture can be changed along with range to target by maneuvering the formation – which presents several issues in control:

- Timing
- Timing knowledge accuracy and synchronization
- Positioning
- Positioning and timing knowledge confidence

and these, in turn, raise issues of:

- Performing inter-satellite communications
- Relaying data
- Commanding via a master control

Such control issues generally have no possible solution apart from an autonomous mechanism (e.g., laser cross-links between the members of the formation to permit minute, near-instantaneous relative position adjustments on a scale measured by the diameter of an atomic nucleus), and for similar reasons (inadequate ability of humans to deal with distant or rapidly occurring phenomena), some level of autonomicity will be required for the system's survival and viability when, for example, the system experiences the effects of a threat that was not predicted.

Autonomic satellite designs will make major contributions to resolving these issues. Each design should be approached first from the viewpoint of constellation architecture, considering the control options available.

9.5 Intelligent Agents in Space Constellations

For single-agent systems in domains similar to space, intelligent machine learning methods (e.g., reinforcement learners) have been successfully used and could be used for single-spacecraft missions. However, applying such solutions directly to multiagent systems often proves to be problematic. Agents may work at cross-purposes, or have difficulty in evaluating their contribution to achievement of the global objective, or both. Constellations based on intelligent multiagent systems would have similar challenges. Concepts from collective intelligence [144] could be applied to the design of the goals for the agents so that they are "aligned" with the global goals, and are "learnable" in that agents could see how their behavior affects their utility and could overcome unforeseen issues.

Satellite intelligence was considered initially by Schetter et al. [133]. They performed comparisons on several high-level agent organizations, along with varying degrees of satellite intelligence, to assess analytically their impact on communication, computation, performance, and reliability. The results indicate that an autonomous, agent-based design provides increased reliability and performance over traditional satellite operations for the control of constellations.

The following sections discuss some of the approaches to using intelligent multiagents for constellation control.

9.5.1 Levels of Intelligence in Spacecraft Agents

Based on the sum of spacecraft functions, four levels of spacecraft intelligence have been identified [133], where I1 denotes the highest level of intelligence and I4 the lowest level (Fig. 9.2):

• The spacecraft-level agent I4 represents the most "unintelligent" agent. It can only receive commands and tasks from other spacecraft-level agents



Fig. 9.2. Identification of spacecraft-level agents based on levels of capable intelligence 2

²Reprinted from Artificial Intelligence, 145(1-2), Thomas Schetter, Mark Campbell and Derek Surka, Multiple agent-based autonomy for satellite constellations, page 164, Copyright (2003), with permission from Elsevier.

in the organization, or from the ground, and execute them. An example includes receiving and executing a control command sequence to move to a new position within the cluster. This type of intelligence is similar to that being flown on most spacecraft today.

- The next higher spacecraft-level agent is I3, with local planning functionalities onboard. "Local" means the spacecraft-level agent is capable of generating and executing only plans related to its own tasks. An example would be trajectory planning for orbital maneuvers.
- Agent I2 adds a capability to interact with other spacecraft-level agents in the organization. This usually requires the agent to have at least partial knowledge of the full agent-based organization, i.e., of other spacecraft-level agents. It must, therefore, continuously keep and update (or receive) an internal representation of the agent-based organization. An example includes coordinating/negotiating with other spacecraft-level agents in case of conflicting requirements.
- The spacecraft-level agent I1 represents the most "intelligent" agent. The primary difference between I1 and the other spacecraft-level agents outlined is that it is capable of monitoring all spacecraft-level agents in the organization and planning for the organization as a whole. This requires planning capabilities on the cluster level (a cluster being a subset of a constellation), as well as a capability by which an agent has full knowledge of all other spacecraft-level agents in the constellation. An example includes calculation of a new cluster configuration and assigning new satellite positions within the cluster.

Selecting the level of intelligence of a multiagent organization is a complex design process. The organization must be:

- Adaptive, able to avoid bottlenecks, and able to reconfigure
- Efficient in terms of time, resources, information exchange, and processing
- Distributed in terms of intelligence, capabilities, and resources

A design selection process starts from an initial spacecraft-level intelligence hierarchy. An example used for TechSat21 [133] is shown in Fig. 9.3. Here, high-level mission tasks were decomposed into lower-level tasks. The spacecraft functions required to support these tasks are listed down the left hand column with subfunctions grouped by category. Across the top are high level spacecraft tasks with subtasks underneath. Tasks are then arranged in a matrix form to provide a visualization of the agent capabilities associated with each level of spacecraft intelligence as presented in Fig. 9.2. The boxes contain the ID's for function and subfunction categories.

9.5.2 Multiagent-Based Organizations for Satellites

Figure 9.4 shows a summary of options as a function of individual spacecraftlevel agent intelligence. Note that lower level functional agents are implied in each of the architectures. As can be seen, the number and composition

HIGH-LEVEL TASKS	HT 1: Science Imaging	Format	HT 2: Science tion Maint nd Contro	aining 1	HT Scie Clu Reconfig	3: ence ster guration	HT Scie Cluster	4: nce Upgrade
SUB-LEVEL TASKS	ST 11: Science	ST 21: Reject. Disturb.	ST 22: Collis. Avoid.	ST 23: Orbit Maneuv.	ST 31: Fault Detect.	ST 32: Form. Change	ST 41: Accept. new S/C	ST 32: Form. Change
INTERACTION 1. Sensing s/c info 2. Transmitting s/c info	F11 F12	F11 F12	F11 F12	F11 F12	F11 F12		F11 F12	
 DECISION-MAKING Imaging Disturbance Collision Avoidance Failure/Loss Upgrade/Gain 	F20	F21	F22		F23		F24	
3. ORGANIZATIONAL 0. Scheduler 1. Planner (Cornwell) 2. Planner (Assign positic 3. Task allocator 4. FF planner (Trajectory	F30 ons)	F30	F30	F30		F31 F32 F33		F31 F32 F33
4. REPRESENTATION 0. Storing Cluster Information	AL ation					F40		F40
 5. OPERATIVE DAR-Imaging Orbit maneuvering (LQR, bang-bang, FF-con 	F50	F51	F51	F51				

Fig. 9.3. Functional breakdown of the task structure specifically for Techsat21³



Group-level: Levels of Coordination/Organization

Fig. 9.4. Coordination architectures for coordination of multiple spacecraft-level ${\rm agents}^4$

³Reprinted from Artificial Intelligence, 145(1-2), Thomas Schetter, Mark Campbell and Derek Surka, Multiple agent-based autonomy for satellite constellations, page 154, Copyright (2003), with permission from Elsevier.

⁴Reprinted from Artificial Intelligence, 145(1-2), Thomas Schetter, Mark Campbell and Derek Surka, Multiple agent-based autonomy for satellite constellations, page 166, Copyright (2003), with permission from Elsevier.

of the different spacecraft-level agents I1–I4 determine the organizational architecture. The top-down coordination architecture includes only one single (highly intelligent) I1 spacecraft-level agent, and the other spacecraft are (unintelligent) I4 agents. The centralized coordination architecture requires at least local planning and possibly interaction capabilities between spacecraft, requiring I3 or I2 agents. The distributed coordination architecture consists of several parallel hierarchical decision-making structures, each of which is "commanded" by an I1 intelligent spacecraft-level agent. In the case of a fully distributed coordination architecture, each spacecraft in the organization represents an I1 spacecraft-level agent, resulting in a totally "flat organization."

9.6 Grand View

The next level of space-based autonomy is to develop and verify agents and agent-communities concepts to the point that they can migrate to actual ground-based operations and, when fully verified and validated in an operational context, migrate to the spacecraft to provide onboard autonomy. Figure 9.5 is a view, a grand view (not the only one), of what such a system might look like [178,179,184], and is one possible representation of *progressive autonomy*. It paints a picture in which we can see many threads of agent-based activity, both ground-based and space-based. The major theme of the figure is that of agent migration from one level to another. The figure depicts the various migration paths that could be taken by agents and communities of agents enroute to a spacecraft. This is the essential theme of our proposed approach to realizing complete autonomy for constellations as well as other mission types.

Progressive autonomy refers to the levels of autonomy that can be incrementally achieved by a dynamic community of agents. Achieving a higher level of autonomy in a community means either increasing an already existing agent's capabilities through reprogramming it, introducing a new agent into the community with the desired capabilities, or allowing an agent to develop a new or modified capability via learning.

Progressive autonomy is advantageous for at least two reasons:

- 1. It allows a new capability to appear from a community of agents supporting an operational mission after that capability has been verified and is trusted outside the testing environment.
- 2. A qualified agent can be dispatched to a community in need on a temporary basis. Once the need has been fulfilled, the agent can be removed. This keeps the operational resource requirements for the community to a minimum.

Figure 9.5 illustrates some of the concepts that are associated with progressive autonomy in agent-based communities on the ground and in space.



Fig. 9.5. Progressive autonomy of a spacecraft

There are three levels represented in the figure:

- 1. An agent development component
- 2. A ground-based autonomy component
- 3. A space-based autonomy component

Agents can migrate from one level to the next depending on their degree of development and validation. Communication between agents on the different levels facilitates the development and validation of agents, since the agents can receive real data from the other levels. The following subsections discuss each of these parts in more detail.

9.6.1 Agent Development

The lower part of Fig. 9.5 represents the agent-development level, where developers write the code, modify existing agents, or use an automated agent-development system that searches for previously developed agents that perform a needed task and updates them based on new requirements input from the developer.

Domain-specialist agents are available to assist in the development of new agents by interacting with the agents being developed, giving a new agent other agents to interact with and then testing for proper functionality. In addition, data may be received from operational agents in the ground control system and spacecraft for additional testing purposes. Agents in the operational mode would know that messages from agents under development are not operational by virtue of a marking of the messages by the messaging service that passes messages to the operational agents. As agents are developed, they are added to the community of domain experts and provide the developers with additional example agents to modify and test against.

The agent-development level also represents an agent incubator. After an agent is developed, there is an incubation period during which agents are tested in a background or shadow mode. When confidence in the agent's behavior is attained, it is moved into an online community doing real work in its domain. It is at this level that the credentials of the agent come into play. These credentials attest to the development methodology and the verification and validation procedures that would directly ensure the agent's correct behaviors.

9.6.2 Ground-Based Autonomy

The middle section of Fig. 9.5 comprises two parts. The right side represents the ground control system, to which some of the agents may migrate after development. In this part, the agent can run in a shadow or background mode where its activities can be observed before it is put into full operation. This allows the users to gain confidence in the agent's autonomy before committing to or deploying it. If a problem is found with the operation of the agent or its operation is not as envisioned or required by the end user, or if enhancements are desired, the agent (or copy of the agent) can be migrated back down to the development area for further modifications or testing. Once modifications are made, it can then be sent back up to the ground operations level, and the process is repeated.

The right side of the middle section also contains an area where agents can be cloned to support parallel processing or fault tolerance: identical agents can be run on multiple (even geographically distributed) platforms. The left side of the middle section represents the agents that are ready to be sent up to a spacecraft, or are operating in a proxy mode. Those agents that are waiting to be sent to a spacecraft may be ones waiting to be uploaded for emergency-resolution purposes (e.g., for anomaly situations), or may carry functionality updates for other agents and are waiting to be uploaded when needed or when resources become available. The agents that are operating in a proxy mode are operating as if they were on the spacecraft, but due (for example) to resource restrictions are temporarily or permanently operating within the ground control system. The proxy agents communicate with other agents and components on the spacecraft as if they were running onboard (subject, as always, to communications constraints).

In a situation requiring an agent to be uploaded, a managing agent in mission control would make a request to the agent-development area (or the repository of validated agents) for an agent with the needed capability, and would (subject to communications constraints) notify the original agent that requested the capability as to the availability of the requested agent. The original requesting agent would then factor this information into its planning, which would be particularly important if the situation were time-critical and alternate actions needed to be planned if the new agent could not be put into service in the constellation within a needed timeframe (e.g., as a result of communications vagaries).

9.6.3 Space-Based Autonomy

The upper part of Fig. 9.5 depicts other communities that are purely operational on a spacecraft or other robotic system, the members of which would be mature agents that would have been approved through an appropriate process. These agent communities may be based around spacecraft subsystem (e.g., instrument, recorder, spacecraft state and management, etc.) or represent a functionality (e.g., anomaly detection, health and safety, science opportunity, etc.).

An agent would be able to migrate from its initial community to other nodes in "agent space" (for lack of a better name). These communities may be logically or physically distinct from the agent's initial community. A single agent may either migrate to a new community when it is no longer needed, or clone itself when needed in multiple communities simultaneously.

The idea of realizing constellation autonomy first through ground-based communities of spacecraft surrogate agents and then migrating the agent community to the actual spacecraft is a flexible, dynamic approach to providing ongoing updates to spacecraft functionality. The progressive autonomy that could be realized through this approach would enable mission control to upload only those agents in the community that have been thoroughly verified and in which there is the appropriate degree of trust.

Swarms in Space Missions

New NASA mission concepts now being studied involve many small spacecraft operating collaboratively, analogous to swarms in nature. The swarm concept offers several advantages over traditional large spacecraft missions: the ability to send spacecraft to explore regions of space where traditional craft simply would be impractical, greater redundancy (and, consequently, greater protection of assets), and reduced costs and risk, among others [176, 181]. Examples are as follows:

- Several unmanned aerial vehicles (UAVs) flying approximately 1 m above the surface of Mars, which will cover as much of the surface of Mars in minutes as the now famous Mars rovers did in their entire time on the planet
- Armies of tetrahedral walkers to explore the Martian and Lunar surface
- Miniaturized pico-class spacecraft to explore the asteroid belt

Under these concepts for future space exploration missions, swarms of spacecraft and rovers will act much like insects such as ants and bees. Swarms will operate as a large group of autonomous individuals, each having simple, cooperative capabilities, and no global knowledge of the group's objective. Such systems entail a wide range of potential new capabilities, but pose unprecedented challenges to system developers. Swarm-based missions, with a new level and kind of complexity that makes untenable the idea of individual control by human operators, suggest the need for a new level and kind of autonomy and autonomicity. Instead of employing human operators to individually control the members of the swarm, a completely different model of operations would be required, where the swarm will operate completely autonomously or with some control at the swarm level.

This chapter will describe swarm-based systems, the possible use of swarms in future space missions, technologies needed to implement them, and some of the challenges in developing them. We will outline the motivation for using swarms in future exploration missions. We will describe one concept mission in relation to the characteristics that such a mission (and similar systems) would need to exhibit in order to become a reality.

W. Truszkowski et al., Autonomous and Autonomic Systems, NASA Monographs 207 in Systems and Software Engineering, DOI 10.1007/978-1-84628-233-1_10,
10.1 Introduction to Swarms

In nature, swarms are large groupings of insects such as bees or locusts where each insect has a simple role, but where the swarm as a whole produces complex behaviors. Strictly speaking, such emergence of complex behavior is not limited to swarms, and there are similar complex social structures occurring with higher order animals and insects that do not swarm per se such as colonies of ants, flocks of birds, packs of wolves, etc. The idea that swarms can be used to solve complex problems has been taken up in several areas of computer science. The term "swarm" in this book refers to a large grouping of simple components working together to achieve some goal and produce significant results [12]. The result of combining simple behaviors (the microscopic behavior) is the emergence of complex behavior (the macroscopic behavior) and the ability to achieve significant results as a "team" [16]. The term should not be taken to imply that these components fly (or are airborne); they may just as well operate on the surface of the earth, under the surface, under water, or in space (including other planets).

Intelligent swarm technology is based on swarm technology where the individual members of the swarm also exhibit independent intelligence [13], and thus, act as agents. Intelligent swarms may be heterogeneous or homogeneous. Even if the swarm starts out as homogeneous, the individual members, with differing environments, may learn different things and develop different goals, and in this way, the swarm becomes heterogeneous. Intelligent swarms may also be made up of heterogeneous elements from the outset, reflecting different capabilities as well as a possible social structure.

Agent swarms are being used in computer modeling and have been used as a tool to study complex systems [55]. Examples of simulations that have been undertaken include swarms of birds [21,115], problems in business and economics [93], and ecological systems [131]. In *swarm simulations*, each of the agents is given certain parameters that it tries to maximize. In terms of bird swarms, each bird tries to find another bird to fly with, and then flies off to one side and slightly higher to reduce its drag, and eventually the birds form flocks. Other types of swarm simulations have been developed that exhibit unlikely emergent behavior. These emergent behaviors are the sums of often simple individual behaviors, but, when aggregated, form complex and often unexpected behaviors. Swarm behavior is also being investigated for use in such applications as telephone switching, network routing, data categorization, command and control systems, and shortest path optimizations.

Swarm intelligence techniques (note the slight difference in terminology from "intelligent swarms") are population-based stochastic methods used in combinatorial optimization problems. In these models, the collective behavior of relatively simple individuals arises from local interactions between each individual and its environment and between each individual and other members of the swarm, which finally results in the emergence of global functional actions by the swarm. Swarm intelligence represents a metaheuristic approach to solving a wide range of problems.

Swarm robotics is the application of swarm intelligence techniques to robotic devices. These systems are programmed to act much like insect swarms where each robot senses and reacts to near-by robots as well as the environment with a given behavior. For example, each robot of an underwater swarm may be watching and following a neighbor, but is also sensing its environment. When something of interest is found by one, it will communicate the information to its neighbors and swim toward it. The others will follow the new leader until they get to the object of interest and then swarm around and examine it. It may be that only a portion of the swarm breaks off, forming a subteam, while the others continue the swarm's search. When the subteam is finished examining the object, they rejoin the team, so there is a constant breaking off and rejoining by members of the swarm.

Swarms may also operate in a tight or loose group and move between extremes depending on the current state of the mission. The group may be tight not only physically, but also operationally. For example, during exploration, the swarm may be scattered over a large area and communicate very little while they each perform their searches. Then, when one finds something of interest, it may broadcast information to inform the rest of the swarm. Others may respond regarding something similar that has already been examined and a group of swarm members may work computationally close to each other (even if they are physically separated) to determine whether it should be further investigated. If so, a subteam would be dispatched to the location and they would work cooperatively to obtain further information (physically close).

10.2 Swarm Technologies at NASA

The Autonomous Nano Technology Swarm (ANTS) project was a joint NASA Goddard Space Flight Center (GSFC) and NASA Langley Research Center (LARC) collaboration whose purpose was to develop revolutionary mission architectures and exploit artificial intelligence techniques and paradigms in future space exploration [29, 32]. This project researched mission concepts that could make use of swarm technologies for both spacecraft and surface-based rovers.

ANTS consists of a number of mission concepts that include:

- SMART: Super Miniaturized Addressable Reconfigurable Technology uses miniaturized robots based on tetrahedrons to form swarms of configurable robots.
- *PAM*: *Prospecting Asteroid Mission* would also launch 1,000 pico-class spacecraft with the aim of exploring the asteroid belt and collecting data on particular asteroids of interest. PAM is described below in more detail.

- 210 10 Swarms in Space Missions
- SARA: The Saturn Autonomous Ring Array would use 1,000 pico-class spacecraft, organized as ten subswarms, each with specialized instruments to perform in situ exploration of Saturn's rings, so as to understand their make up and how they were formed. The concept mission would require self-configuring structures for nuclear propulsion and control. Additionally, autonomous operation would be necessary for both maneuvering around Saturn's rings and collision avoidance between spacecraft.
- ANTS: Application Lunar Base Activities would exploit new NASAdeveloped technologies in the field of miniaturized robotics, which would form the basis of remote landers to be launched to the moon from remote sites and would exploit innovative techniques (described below in Sect. 10.2.1) to allow rovers to move in an amoeboid-like fashion over the moon's uneven terrain.

The following sections describe the SMART and PAM mission concepts. The description of SMART covers similar technologies that would also be needed for the Lander Amorphous Rover Antenna (LARA) (and other) concept missions. Since SARA and PAM have many attributes in common (as regards to autonomous operation), we will concentrate on a description of PAM in the following.

10.2.1 SMART

The ANTS SMART architectures were initiated to develop new kinds of structures capable of:

- Goal-oriented robotic motion
- Changing form to optimize function (morphological capabilities)
- Adapting to new environmental demands (learning and adaptation capabilities)
- Repairing-protecting itself (autonomic capabilities)

The basic unit of the structures is a tetrahedron (Fig. 10.1) consisting of four addressable nodes interconnected with six struts that can be reversibly deployed or stowed. More complex structures are formed from interconnecting these reconfigurable tetrahedra, making structures that are scalable and leading to massively parallel systems. These highly-integrated, 3D meshes of actuators/nodes and structural elements hold the promise of providing a new approach to robust and effective robotic motion. The current working hypothesis is that the full functionality of such a complex system requires fully autonomous intelligent operations at each node.

The tetrahedron (tet) "walks" by extending certain struts, changing its center of mass, and "falling" in the desired direction. As the tetrahedral structure "grows" by interfacing more and more tets, the falling motion evolves to a smoother walking capability, i.e., the smoother walking-climbing-avoiding



Fig. 10.1. Basic unit of tetrahedral structures



Fig. 10.2. Prototype of a tetrahedron rover (image credit: NASA)

capabilities emerge from the orchestration of the capabilities of the tetrahedra involved in the complex structure. Figure 10.2 shows a picture of a prototype tetrahedron.

The basic tetrahedron structure was modeled as a communicating and cooperating/collaborating four-agent system with an agent associated with each node of the tetrahedron. An agent, in this context, is an intelligent autonomous process capable of bi-level deliberative and reactive behaviors with an intervening neural interconnection (the structure of the neural basis function [30]). The node agents also possess social and introspective behaviors. The problem



Fig. 10.3. A picture of a 12-tet rover (image credit: NASA)

to be solved is to scale this model up to one capable of supporting autonomous operation for a 12-tet rover, a structure realized by the integration of 12 tets in a polyhedral structure. The overall objective is to achieve autonomous robotic motion of this structure. Figure 10.3 shows a drawing of a 12-tet rover.

10.2.2 NASA Prospecting Asteroid Mission

The ANTS PAM concept mission [31, 32, 181] would involve the launch of a swarm of autonomous pico-class (approximately 1 kg) spacecraft that would explore the asteroid belt for asteroids with characteristics of scientific interest. Figure 10.4 gives an overview of the PAM mission concept [176]. In this mission, a transport ship launched from earth would travel to a Lagrangian point. From this point, 1,000 spacecraft, which would have been assembled en route from earth, would be launched into the asteroid belt. Each spacecraft would have a solar sail as the primary means of propulsion (using photon pressure from the Sun's light), supplemented with tiny thrusters to maneuver independently. Each spacecraft would carry just one specialized instrument for collecting a specific type of data from asteroids in the belt. With onboard computational resources, each spacecraft would also have artificial intelligence and heuristics systems for control at the individual and team levels. For communications, spacecraft would use low bandwidth to communicate within the swarm and high bandwidth for sending data back to earth. It is expected that 60-70% of the spacecraft would be lost during the mission, primarily because of collisions with each other or with asteroids during exploration operations, since their ability to maneuver will be severely limited.

As Figs. 10.4 and 10.5 show, teams would consist of members from three classes of spacecraft within the swarm, with members in each class combining



Fig. 10.4. NASA's autonomous nano technology swarm (ANTS) mission overview



Fig. 10.5. ANTS encounter with an asteroid (image credit: NASA)

to form teams that explore particular asteroids. Approximately 80% of the spacecraft would be workers that carry the specialized instruments to obtain specific types of data. Examples of instruments include magnetometers

and x-ray, gamma-ray, visible/infrared, or neutral mass spectrometers. Each worker would gather only its assigned data types. Some of the spacecraft would be coordinators (called leaders or rulers) that would coordinate the efforts of the workers. They would apply rules to determine the types of asteroids and data of interest to the mission. The third type of spacecraft are called messengers. Messengers would coordinate communications among the workers, rulers, and mission control on earth.

Figure 10.5 depicts the flow of activity as teams explore an asteroid, exchange data, and return data to earth. A single ANTS spacecraft could also survey an asteroid in a flyby, sending quick-look data to the ruler, which would then decide whether the asteroid warranted further investigation using a team. The ruler would choose team members according to the instruments they carry.

Many operational scenarios are possible within the overall concept of missions that act like a natural swarm culture. In one scenario, the swarm would form subswarms under the control of a ruler, which would contain models of the types of science that it wants to perform. The ruler would coordinate workers, each of which would use its individual instrument to collect data on specific asteroids and feed this information back to the ruler, which would determine which asteroids are worth examining further. If after consulting its selection criteria and using heuristic reasoning it determines that the asteroid merits further investigation, an imaging spacecraft would be sent to the asteroid to ascertain its shape and size and to create a rough model to be used by other spacecraft for maneuvering around the asteroid. The ruler would also arrange for additional workers to transit to the asteroid with an expanded repertoire of instruments to gather more complete information. In effect, the spacecraft would form a team. The leader would be the spacecraft that contains models of the types of experiments or measurements the team needs to perform. The leader would relay parts of this model to the team workers, which then would take measurements of asteroids using whatever type of instrument they have until something matched the goal the leader sent. The workers would gather the required information and send it to the team leader, which would integrate it and return it to the ruler that formed the team. The ruler might then integrate this new information with information from previous asteroid explorations and use a messenger to relay findings back to earth.

10.2.3 Other Space Swarm-Based Concepts

An autonomous space exploration system studied at Virginia Tech, funded by the NASA Institute for Advanced Concepts (NIAC), consists of a swarm of low altitude, buoyancy-driven gliders for terrain exploration and sampling, a buoyant oscillating wing that absorbs wind energy, and a docking station that could be used to anchor the energy absorber, charge the gliders, and serve as a communications relay [96]. The work was built on success with underwater gliders used for oceanographic research. The intent was to develop low-cost planetary exploration systems that could run autonomously for years in harsh environments, such as in the sulfuric acid atmosphere of Venus or on Titan (the largest of Saturn's moons).

A second NASA swarm-related project titled "Extremely Large Swarm Array of Picosats for Microwave/RF Earth Sensing, Radiometry, and Mapping" [73] was also funded by NIAC. The proposed telescope would be used to do such things as characterize soil moisture content, atmospheric water content, snow accumulation levels, flooding, emergency management after hurricanes, weather and climate prediction, geological feature identification, and others. To accomplish this would require an antenna size on the order of 100 km at a GEO orbit. To implement such an large antenna, a highly sparse spacefed array antenna architecture was proposed that would consist of 300,000 picosats, each being a self-contained onechip spacecraft weighing 20 g.

10.3 Other Applications of Swarms

The behavior of swarms of bees has also been studied as part of the BioTracking project at Georgia Tech [9]. To expedite the understanding of the behavior of bees, the project videotaped the behavior of bees over a period of time, using a computer vision system to analyze data on sequential movements that bees use to encode the location of supplies of food, etc. It is anticipated that such models of bee behavior can be used to improve the organization of cooperating teams of simple robots capable of complex operations. A key point is that the robots need not have a priori knowledge of the environment, and that direct communication between robots in the teams is not necessary.

Eberhart and Kennedy have developed an optimization technique based on particle swarms [78] that produces fast optimizations for a wide number of areas including UAV route planning, movement of containers on container ships, and detecting drowsiness of drivers. Research at Penn State University has focused on the use of particle swarms for the development of quantitative structure activity relationships (QSAR) models used in the area of drug design [23]. The research created models using artificial neural networks and k-nearest neighbor and kernel regression. Binary and niching particle swarms were used to solve feature-selection and feature-weighting problems.

Particle swarms have influenced the field of computer animation also. Rather than scripting the path of each individual bird in a flock, the Boids project [115] elaborates a particle swarm of simulated birds. The aggregate motion of the simulated flock is much like that in nature. The result is from the dense interaction of the relatively simple behaviors of each of the (simulated) birds, where each bird chooses its own path.

Much success has been reported from the use of ant colony optimization (ACO), a technique that studies the social behaviors of colonies of ants and

uses these behavior patterns as models for solving difficult combinational optimization problems [35]. The study of ants and their ability to find shortest paths has lead to ACO solutions to the traveling salesman problem, as well as network and internet optimizations [34, 35].

Work at University of California Berkeley is focusing on the use of networks of unmanned underwater vehicles (UUVs). Each UUV has the same template information, containing plans, subplans, etc., and relies upon this and its own local situation map to make independent decisions, which will result in cooperation between all of the UUVs in the network. Experiments involving strategies for group pursuit were also done.

10.4 Autonomicity in Swarm Missions

Swarms are being used in devising solutions to various problems principally because they present an appropriate model for those problems. Sections 10.2 and 10.3 described several application areas of swarm technology where the approach seems to be particularly successful.

But swarms (in nature or otherwise) inherently need to exhibit autonomic properties. To begin with, swarms should be self-directed and self-governed. Recall that this is achieved through the complex behavior that emerges from the combination of several simple behaviors and their interaction with the environment. It can be said that in nature, organisms and groups/colonies of individuals, with the one fundamental goal of survival, would succumb as individuals and even as species without autonomicity. A natural conclusion is that artificial swarms with planned mission objectives must also possess autonomicity.

The described ANTS PAM concept mission would need to exhibit almost total autonomy to succeed. The mission would also exhibit many of the properties required to qualify it as an autonomic system [161, 181, 182]:

- Self-configuring: The ANTS's resources must be fully configurable to support concurrent exploration and examination of hundreds of asteroids. Resources must be configured at both the swarm and team (subswarm) levels in order to coordinate science operations while simultaneously maximizing resource utilization.
- Self-optimizing: Rulers self-optimize primarily through learning and improving their ability to identify asteroids that will be of interest to scientists. Messengers self-optimize through positioning themselves appropriately for optimum communications. Workers self-optimize through learning and experience. Self-optimization at the swarm level propagates up from the self-optimization of individuals.
- Self-healing: ANTS must self-heal to recover from damage due to solar storms or collisions with an asteroid or between ANTS spacecraft. Loss of a ruler or messenger may involve a worker being "upgraded" to fulfill that role.

Additionally, loss of power may require a worker to be listed as dead ("killed off" via an apoptosis mechanism [145, 161]).

Self-protecting: In addition to protecting themselves from collision with asteroids and other spacecraft, ANTS teams must protect themselves from solar storms, where charged particles can degrade sensors and electronic components and destroy solar sails (the ANTS spacecraft's sole source of power and primary means to perform maneuvering). ANTS teams must re-plan their trajectories or, in worst-case scenarios, must go into "sleep" mode to protect their sails and instruments and other subsystems.

The concept of autonomicity can be further elaborated beyond the selfchop properties listed above. Three additional self-properties – self-awareness, self-monitoring, and self-adjusting – will facilitate the basic self-properties. Swarm (ANTS) individuals must be aware (have knowledge) of their own capabilities and limitations, and the workers, messengers, and rulers will all need to be involved in constant self-monitoring and (if necessary) selfadjusting, thereby forming a feedback control loop. Finally, the concept of autonomicity would require environmental awareness. The swarm (ANTS) individuals will need to be constantly environmentally aware to enable effective self-adaptation and ensure mission success.

10.5 Software Development of Swarms

Developing the software for the ANTS missions would be monumentally complicated. The total autonomy requirement would mean that the software would likely be based on a heuristic approach that accommodates the swarm's social structure. Artificial-intelligence technologies, such as genetic algorithms, neural nets, fuzzy logic, and on-board planners are candidate solutions. But the autonomic properties, which alone make the system extremely complex, are only part of the challenge. Add intelligence for each of the thousand interacting spacecraft, and it becomes clear that the mission depends on several breakthroughs in software development.

10.5.1 Programming Techniques and Tools

A primary requirement would be a new level or new class of programming techniques and tools that either replace or build on object-oriented development. The idea is to reduce complexity through novel abstraction paradigms that would essentially "abstract away" complexity. Developers would use predefined libraries or components that have been solidly tested and verified. The level of programming languages would need to be high enough that developers could use constructs that are natural extensions to the software type under development. Another requirement would be tools and techniques that would have builtin autonomic, intelligent, and interacting constructs to reduce development time and increase developer productivity. Tools would need to allow rapid simulation so that developers might identify errors in requirements or code at the earliest stage possible. For now, ideas about creating standard intelligent, autonomic components are still evolving: there is yet no consensus as to what constitutes a system of such components. Hopefully, more research and development in these areas will yield effective and timely results.

10.5.2 Verification

These new approaches to exploration missions simultaneously pose many challenges. Swarm missions will be highly autonomous and will have autonomic properties. Many of these missions will be sent to parts of the solar system where manned missions are regarded as infeasible, and where, in some instances, the round-trip delay for communications between earth and the spacecraft exceeds 40 min., meaning that the decisions on responses to problems and undesirable situations must be made in situ rather than from ground control on earth. The degree of autonomy that such missions will require would mean an extreme burden of testing in order to accomplish system verification. Furthermore, learning and adaptation toward continual improvements in performance during mission operations will mean that emergent behavior patterns simply cannot be fully predicted through the use of traditional system development methods. Consequently, formal specification techniques and formal verification will play vital roles in the future development of these types of missions.

Full testing of software of the complexity of the ANTS mission may be recognized as a heavy burden and may have questionable feasibility, but verification of the on-board software, especially the mechanism that endows the spacecraft with autonomy and the ability to learn, is crucial because the oneway communications delay makes real-time control by human operators on earth infeasible. Large communications delays mean human operators could not, in many scenarios, learn of problems or errors or anomalies in the mission until the mission had substantially degraded or failed. For example, in a complex system with many concurrently communicating processes on board or among the members of the swarm, race conditions are highly likely, but such conditions rarely come to light during the testing or mission development phase by inputting sample data and checking results. These types of errors are time-based, occurring only when processes send or receive data at particular times or in a particular sequence, or after learning takes place. To find these errors, testers must execute the software in all the possible combinations of the states of the communicating processes. Because the state space is extremely large (and probably extremely difficult to project in sufficient detail for actual testing), these systems become untestable with a relatively small number of elements in the swarm. Traditionally, to get around the state-space explosion problem, testers have artificially reduced the number of states of these types of systems and approximated the underlying software using models. This approach, in general, sacrifices fidelity and can result in missed errors. Consequently, even with relatively few spacecraft, the state space can be too large to realistically test.

One of the most challenging aspects of using swarm technology is determining how to verify that emergent system behavior will be proper and that no undesirable behaviors will occur. Verifying intelligent swarms is even more difficult because the swarms no longer consist of homogeneous members with limited intelligence and communications. Verification will be difficult not only because each individual is tremendously complex, but also because of the many interacting intelligent elements. To address the verification challenge, ongoing research is investigating formal methods and techniques for verification and validation of swarm-based missions.

Formal methods are proven approaches for ensuring the correct operation of complex interacting systems. Formal methods are particularly useful in specifying complex parallel and distributed systems – where a single person finds it difficult to fully understand the entire system and where there are typically multiple developers. Testers can use a formal specification to prove that system properties are correct – for example, that the underlying system will go from one state to another or not into a specific state. They can also check for particular types of errors such as race conditions, and use the formal specification as a basis for model checking.

Most formal methods do not address the problem of verifying emergent behavior. Clearly in the ANTS PAM concept, the combined behavior of individual spacecraft is far more complex than the behavior of each spacecraft in isolation. The formal approaches to swarm technologies (FAST) project surveyed formal methods techniques to determine whether any would be suitable for verifying swarm-based systems and their emergent behavior [125, 126]. The project found that there are a number of formal methods that support the specification of either concurrency or algorithms, but not both. Though there are a few formal methods that have been used to specify swarm-based systems, the project found only two formal approaches that were used to analyze the emergent behavior of swarms.

Weighted synchronous calculus of communicating systems (WSCCS), a process algebra, was used by Sumpter et al. to analyze the nonlinear aspects of social insects [163]. X-Machines have been used to model cell biology [61,62], and with modifications, the X-Machines model has the potential for specifying swarms. Simulation approaches are also being investigated to determine emergent behavior [55]. However, these approaches do not predict emergent behavior from the model, but rather model the emergent behavior after the fact.

The FAST project defined an integrated formal method, which is appropriate for the development of swarm-based systems [121]. Future work will concentrate on the application of the method to demonstrate its usefulness, and on the development of appropriate support tools. NASA is pursuing further development of formal methods techniques and tools that can be applied in the development of swarm-based systems to help achieve confidence in their correctness.

10.6 Future Swarm Concepts

A brief overview of swarm technologies was presented with emphasis on their relevance for potential future NASA missions. Swarm technologies hold promise for complex exploration and scientific observational missions that require capabilities that would be unavailable in missions designed around single spacecraft.

While swarm autonomy is clearly essential for missions where human control is not feasible (e.g., when communications delays are too great or communications data rates are inadequate for effective remote control), autonomicity is essential for survival of individual spacecraft as well as the entire swarm as a consequence of hostile space environments.

Although ANTS was a concept mission, the underlying techniques and technologies that were developed are also motivating other technologies and applications. ANTS technology has many potential applications in military and commercial environments, as well as in other space missions. In military surveillance, smaller craft, perhaps carrying only a basic camera or other instrument, could coordinate to provide 3D views of a target. The US Navy has been studying the use of vehicle swarms for several years. In mining and underwater exploration, autonomous craft could go into areas that are too dangerous or small for humans. For navigation, ANTS technology could make GPS cheaper and more accurate because using many smaller satellites for triangulation would make positioning more accurate.

Finally, in other types of space exploration, a swarm flying over a planetary surface could yield significant information in a short time. The ANTS technology could also benefit commercial satellite operations, making them both cheaper and more reliable. With its autonomic properties, a swarm could easily replace an individual pico-satellite, preserving operations that are now often lost when satellites become damaged. Mission control could also increase functionality simply by having the swarm add members (perhaps from a collection of pico-satellites already in orbit as standby spares) with the needed functionality, rather than launching a new, large, complex satellite.

The obvious need for advances in miniaturization and nanotechnology is prompting groundbreaking advances at NASA and elsewhere. The need for more efficient on-board power generation and storage motivates research in solar energy and battery technology, and the need for energy-efficient propulsion motivates research on solar sails and other technologies such as electric-field propulsion. The ANTS concepts also push the envelope in terms of software technologies for requirements engineering, nontrivial learning, planning, agent technology, self-modifying systems, and verification and validation technologies. The paradigms, techniques, and approaches stimulated by concept missions like ANTS open the way for new types of future space exploration missions: namely, large numbers of small, cooperating spacecraft conducting flexible, reliable, autonomous, and cost-effective science and exploration operations beyond the capabilities of the more familiar large spacecraft.

Concluding Remarks

In this book, we have examined technologies for system autonomy and autonomicity. We have considered what it means for a system to be autonomous and autonomic, and have projected how the concepts might be applied to spacecraft and other aspects of NASA missions. We discussed current spacecraft ground and flight operations, described how autonomy and autonomic technology is currently applied to NASA missions, and identified the areas where additional autonomy could be beneficially effective. We also considered artificial intelligence techniques that could be applied to current and future missions to provide additional autonomy and autonomicity.

We now proceed to identify factors that drive the use of new technology and discuss the necessity of software reliability for space missions. We will discuss certain future missions and their needs for autonomy and autonomic technology, and finally will consider the NASA strategic plan and the manner in which autonomy and autonomic systems may be involved in supporting that plan.

11.1 Factors Driving the Use of Autonomy and Autonomicity

As discussed in other parts of this book, there are a number of factors that drive the application of autonomy and autonomicity. New science using spacebased platforms gives rise progressively to new methods and means, and calls for new and increasingly sophisticated instruments and complex new instrument configurations in space, as will be seen from a number of examples to be given later. Future explorations in more remote environments bring new challenges for mission control when real-time information for human operators becomes impossible due to the fundamental reality of signal delays under speed-of-light limitations. Reductions of mission costs is a continuing and evidently unavoidable necessity, and as we have seen, has been increasingly realized through reducing human involvement in routine mission operations – a likely recurring theme in future missions. These are among the principal factors that make autonomy and autonomicity an increasing necessity in many future space missions.

New scientific discovery is enabled by observing deeper into space using instruments that are more sensitive and complex, by making multiple observations simultaneously with coordinating and cooperating spacecraft, and by reacting to the environment or science of opportunity more quickly. All of these capabilities can be realized either through the use of autonomous systems and autonomic properties or by having a human onboard the spacecraft or, except in the case of very remote assets, by having a sufficiently large operations staff at the mission control center. In some missions, a human onboard the spacecraft would not be able to react fast enough to a phenomenon, or maintain required separations between spacecraft. In other missions, such as missions to another planet or the asteroid belt, a crewed spacecraft would be infeasible, too dangerous, or too costly.

We also saw in Chap. 3 that adding autonomy to missions is not new. Autonomy or automation has been an increasing aspect of flight and ground software with the gradually increasing complexity of missions and instruments and with ongoing budgetary pressures. This trend is continuing into future NASA missions, more particularly robotic (un-crewed) missions.

11.2 Reliability of Autonomous and Autonomic Systems

In NASA missions, software reliability is extremely important. A software failure can mean the loss of an entire mission. Ground-based systems can be tended by humans to correct any errors. In unmanned space systems, with only a few exceptions (e.g., the Hubble Space Telescope, which was designed to be serviced on orbit by space-walking humans), any corrections must be performed strictly via radio signals, with no possibility of human presence. Therefore, software and hardware for space missions must be developed, verified, and tested to a high level of assurance, with a corresponding cost in time and money.

Because of the need for high assurance, one of the challenges in adding autonomy and autonomicity to spacecraft systems is to implement these concepts so they work reliably and are verifiable. The software must be robust enough to run on a spacecraft and perhaps as part of a community of spacecraft. Further, the software must be implementable in a reasonable timeframe and for a reasonable cost (relative to the type and importance of the mission).

Autonomous systems often require flexible communication systems, mobile code, and complex functionality, not all of which is always fully understood at the outset. A particular problem with these types of systems is that such systems can never really be tested to any degree of sufficiency, as an intelligent system may adapt its behavior on every execution. New ways of testing and monitoring this type of software are needed to give mission principal investigators the assurance that the software is correct [117]. This was addressed to some extent in a previous book in this series [119].

In addition to being space-based, many of the proposed missions will operate very remotely and without frequent contact with a ground-based operations control center or with a large communications lag time. Such conditions make detecting and correcting software errors before launch even more important, because patching the software during the mission's operational phase will be difficult, impractical, or impossible.

Autonomous missions are still at a relatively early stage of evolution in NASA, and the software development community is still learning approaches to their development. These are highly parallel systems that can have very complex interactions. Even simple interacting systems can be difficult to develop, as well as debug, test, and validate. In addition to being autonomous and highly parallel, these missions may also have intelligence built into them, and they can be distributed and can engage in asynchronous communications. Consequently, these systems are difficult to verify and validate.

New verification and validation techniques are required [100, 113, 118, 120, 124]. Current techniques, based on large monolithic systems, have worked well and reliably, but do not translate to these new autonomous systems, which are highly parallel and nondeterministic.

11.3 Future Missions

Some future missions were discussed in other parts of this book. The following material from indicated sources describes additional mission concepts that are undeniably ambitious in nature – the success of which will require autonomous and autonomic properties (Table 11.1).¹

Mission	Launch date	Number of spacecraft
Big Bang Observer	2020 +	Approx. 12
Black Hole Imager	2025 +	Approx. 33
Constellation X	2020 +	1
Stellar Imager	After 2020	Approx. 17
LISA	2020	3
MIDEX SIRA	2015 +	12 - 16
Enceladus	2018	2
Titan	2018	2
JWST	$2013 \pm$	1

 Table 11.1.
 Some NASA future missions

LISALaser Interferometer Space Antenna; MIDEXMedium Explorer; SIRASolar Imaging Radio Array; JWSTJames Webb Space Telescope

¹The following descriptions are summarized from the NASA future missions web site.

226 11 Concluding Remarks

The Laser Interferometer Space Antenna (LISA) mission concept specifies spacecraft that measure passing gravitational waves. LISA will be a constellation of three spacecraft that uses laser interferometry for precise measurement of distance changes between widely separated freely falling test masses housed in each spacecraft (Fig. 3.3). The spacecraft are at the corners of an approximately equilateral triangle about 5 million kilometers on a side in heliocentric orbit. The science instrument is created via laser links connecting the three spacecraft. It is formed by measuring to high levels of precision the distances separating the three spacecraft (i.e., the test masses) via the exchange of the laser light. From the standpoint of Bus FSW providing spacecraft attitude and position control, the number of sensors and actuators that must be interrogated and commanded is at least twice the number associated with a more traditional mission. Similarly, the number of control modes is double that of a typical astrophysics mission, as are the number of parameters solved-for by the state estimator.

The Big Bang Observer and the Black Hole Imager are part of the Beyond Einstein program that will further test Einstein's general theory of relativity. The Big Bang Observer will explore the beginning of time and will build on the LISA mission to directly measure gravitons from the early Universe still present today. The Black Hole Imager mission will calculate the aspects of matter that fall into a black hole by conducting a census of hidden black holes, revealing where, when, and how they form.

The Stellar Imager mission will help increase understanding of solar/ stellar magnetic activity and its impact on the origin and continued existence of life in the Universe, structure and evolution of stars, and habitability of planets. It will also study magnetic processes and their roles in the origin and evolution of structure and the transport of matter throughout the Universe. The current baseline architecture for the full Stellar Imager mission is a spacebased, UV-Optical Fizeau Interferometer with 20–30 1-m primary mirrors, mounted on formation-flying "mirrorsats" distributed over a parabolic virtual surface whose diameter can be varied from 100 m upto as much as 1,000 m, depending on the angular size of the target to be observed (Fig. 11.1). The hub and all of the mirrorsats are free-flyers in a tightly-controlled formation in a Lissajous orbit around the Sun-Earth Lagrange L2 point. The mission will also use autonomous analysis of wavefronts and will require real-time correction and control of tight formation flying.

The Solar Imaging Radio Array (SIRA) mission will be a Medium explorer (MIDEX) mission and will perform interferometric observations of low frequency solar and magnetospheric radio bursts. The primary science targets are coronal mass ejections (CMEs), which drive radio-emission-producing shock waves. A space-based interferometer is required because the frequencies of observation (<15 MHz) do not penetrate the ionosphere. SIRA will require 12–16 microsatellites to establish a sufficient number of baselines with separations on the order of kilometers. The microsat constellation consists of microsate located quasi-randomly on a spherical shell, initially of radius



Fig. 11.1. The Stellar Imager mission (image credit: NASA)

5 km. The baseline microsat is 3-axis stabilized with an earth-pointing bodymounted high gain antenna and an articulated solar array. The microsats will have limited inter-mirosat communications to help maintain prescribed distances each to the others.

A mission to Enceladus is one of several options NASA is considering for an outer planets Flagship mission that would be launched no earlier than 2015. Several types of missions are being studied for Enceladus. One is where a lander is sent down to collect samples. The lander would use autonomous hazard avoidance to land safely. The autonomous hazard avoidance would use a descent camera to identify and avoid rocks and blocks of ice in order to land at a relatively smooth and flat location. The lander would perform its observations, collect the samples, and analyze them, while the orbiter continues to orbit Saturn, beyond communication range. When the orbiter returns 8.22 days later, the lander would uplink its data to the orbiter and conclude its mission.

The Titan Explorer with Orbiter mission will map Titan with a highresolution radar and study the atmosphere, prebiological chemistry, and potential life. The mission will include an orbiter to relay communications from a module that will land on the surface. The orbiter will also perform aerocapture to test the atmosphere.

The James Webb Space Telescope (JWST) mission is the replacement for the Hubble Space Telescope (Fig. 11.2). The JWST flight and ground system will be developed as an integrated system that will provide seamless operations from science proposal to data distribution, with a minimum of interfaces. The spacecraft will incorporate greater onboard autonomy, executing a high



Fig. 11.2. The James Webb Space Telescope (JWST) (image credit: NASA)

level list of observations rather than a detailed timeline. This will simplify the ground system scheduling activities and permit the spacecraft to perform over many days with little or no direction from the ground. JWST will be able to operate autonomously for periods of several days between these uplinks, and therefore, continuous staffing of the control center will not be required in the latter years of operations. The observatory will operate nearly autonomously throughout its science operations phase. Control of most housekeeping and science collecting functions will be provided onboard the observatory by an event-driven command system. During the science operations phase of the mission, one communications contact per day and one command load per week or two will be sufficient to support operations. Figure 11.3 shows an early design of autonomous attitude control transitions for JWST [74].

11.4 Autonomous and Autonomic Systems in Future NASA Missions

The NASA 2006 Strategic Plan states:

NASA also will develop and test technologies for power and autonomous systems that can enable more affordable and sustainable space exploration by reducing both consumables launched from earth and the risks for mission operations. Advanced power systems, including solar, fuel cell, and potential nuclear power, will provide abundant



Fig. 11.3. Early design of autonomous attitude control mode transitions

power to a lunar outpost so that exploration will not be limited by the available energy. Intelligent robotics will assist the crew in exploring, setting up, operating, and maintaining the outpost. Autonomous systems will reduce mission risk by alerting the crew to impending failures, automatically reconfiguring in response to changing conditions, and performing hazardous and complex operations.

The plan continues:

Therefore, NASA's long-term Earth science plan is to use sentinel orbits (e.g., Lagrange points, geostationary, and medium Earth orbit) and constellations of smart satellites as parts of an integrated, interactive "sensorweb" observing system that complements satellites in low Earth orbit, airborne sensors, and surface-based sensors. NASA will mature active remote sensing technologies (radars and lasers) to take global measurements of Earth system processes from low and geostationary Earth orbits.

As new types of Earth observations become available, information systems, modeling, and partnerships to enable full use of the data for scientific research and timely decision support will become increasingly important. The sensorweb observing systems of the future will perform satellite constellation management, automated detection of environmental phenomena, tasking of other elements of the observing network, onboard data processing, data transmittal, and data archival and distribution to users of earth observations. The sensorweb will be linked to "modelwebs" of prediction systems enabled by NASA and formed by Agency partners to improve the forecast services they provide. NASA's investment in these areas (through such means as the Advanced Information Systems Technology program) will help the Nation take full advantage of enhanced information availability. In particular, the role of models in converting the satellite-produced information into useful products for environmental characterization and prediction will become more crucial.

The above statements in the NASA strategic plan project the increasing importance of mission autonomy and autonomicity. In the first quote, the strategic plan notes that autonomous systems are needed to make future missions possible and make them affordable. It further indicates the need for adding intelligence to robotics and reducing mission risks through autonomy, by, for example, alerting the crew to failures, and automatically reconfiguring (autonomicity) as mission conditions change.

The second quoted passage describes the planned sensor web in terms of constellations of smart satellites, with the prospect of automatic science of opportunity, and with heterogenous systems working together to make up the sensor web. Again, all of these will require autonomous and autonomic systems.

With the new mission concepts that are taking shape, an excellent opportunity now exists to insert autonomy and autonomicity, as well as agent technologies, into these missions. Since these technologies make many of the missions feasible, the science community is now looking to the artificial intelligence and agent software community to implement these ideas in future flight software.

This book has attempted to provide background on NASA ground and space systems and exploration thrusts, and has presented autonomy and autonomicity as a technological means to enhance space mission usefulness, cost effectiveness, and functionality. The motivation behind this book has been to help others direct their research and development into this area, as well as to stimulate future missions to adopt more autonomy and autonomicity, thereby enhancing new exploration and making scientific discovery more productive.

Attitude and Orbit Determination and Control

In order to perform its science mission, a spacecraft must, in general, know its orientation in space and its position relative to the targets it plans to observe. The term *orbit* refers to the spacecraft's position and velocity with respect to an inertial reference frame. The term *attitude* refers to the spacecraft's orientation with respect to some inertial reference frame (for Earth-orbiting spacecraft, this ordinarily is the geocentric inertial (GCI) frame).

Traditionally, the determination of a spacecraft's orbit was exclusively a ground system function. While in communication with the spacecraft, the ground system would collect tracking data, and then the Flight Dynamics computational facility (in the context of NASA/Goddard missions) processed the data to calculate the definitive (i.e., actual) spacecraft orbit associated with the time(s) corresponding to the collected tracking data. Flight Dynamics then would calculate a *predicted* spacecraft orbit by combining the definitive orbit data with mathematical models describing the gravitational interactions and orbital perturbations experienced by the spacecraft. If the spacecraft's flight software (FSW) did not include these models, the ground system would uplink a set of "fit parameters" (or a position/velocity-vector file) tailored to the FSW's orbit propagator (or orbit interpolator), which the spacecraft would then use to compute its position and velocity. If the spacecraft FSW did include at least a subset (or simplified version) of these models, the ground system would uplink "seed" vectors that the FSW would utilize as starting input when integrating the spacecraft equations of motion to calculate its position and velocity at some arbitrary time. In either case, the FSW output was a predicted orbital position and velocity as opposed to a measured orbital position and velocity. More recently, for low earth orbit (LEO) spacecraft, orbital position can be directly measured to high accuracy using the onboard global positioning system (GPS), which also can be used to synchronize the spacecraft clock to GPS time.

By contrast, measuring spacecraft attitude has been a standard onboard function practically since flight computers were introduced. What has evolved over time has been the accuracy and sophistication with which those measurements have been made. In early spacecraft, onboard storage and CPU capabilities limited onboard attitude calculations to coarse attitude determination using sensors such as Sun sensors and magnetometers, or limited fineattitude determination to processing star-tracker output using very small onboard star catalogs, or using reference stars carefully preselected by the ground system. In the 1990s, more powerful onboard processing capabilities and cheap onboard storage facilitated the enhancement of this capability. For example, for the Rossi X-ray Timing Explorer (RXTE) mission, the FSW included a star catalog and star identification capabilities permitting the FSW to calculate the spacecraft's own fine attitude without ground support. Later still, the Wilkinson Microwave Anisotropy Probe (WMAP) mission's star tracker contained its own star catalog and star identification algorithms, and so could output the spacecraft attitude directly (although it did not compensate for the star tracker's misalignment relative to the spacecraft), providing a "Lost in Space" capability.

Knowing the spacecraft position and orientation is only half of the guidance, navigation, and control (GN&C) problem. The other half is controlling the spacecraft position and attitude such that the desired science can be performed. As with orbit determination, planning orbit maneuvers historically has been solely a ground function because of the CPU intensive and lookahead nature of the problem. According to the traditional paradigm (again in the NASA/Goddard context), after determining the current spacecraft orbit, Flight Dynamics would evaluate whether that orbit satisfied the mission requirements, and if it did, would calculate when orbit perturbations were likely to drive the orbit outside those requirements. In collaboration with the flight operations team (FOT), Flight Dynamics would then plan and schedule (as necessary) a small stationkeeping orbit maneuver designed to restore the orbit to its operational geometry. The stationkeeping plan created by this process would include highly detailed instructions (i.e., which thrusters to fire, how long they should fire, the thruster configuration when firing, etc.), which would be uplinked to the spacecraft as "Delta-V" (change-of-velocity) commands, and would be executed open-loop by the FSW, which would have no way to evaluate the success/failure of the orbit maneuver. Lastly, Flight Dynamics would evaluate the postburn orbit to determine whether further corrections were necessary. This same basic procedure also applied to the major orbit maneuvers required to acquire mission orbit, except the planning and scheduling would be more elaborate, typically requiring several burns to complete. Recently, however, serious consideration has been given to planning and scheduling routine stationkeeping orbit maneuvers autonomously onboard. Migrating this function to today's more capable flight computers not only would reduce the cost associated with using ground staff to perform a routine operation, but also would reduce operational risk by eliminating unnecessary command uplinks.

Onboard control of attitude, by contrast, has been a necessary part of all space missions other than a handful of generally low-cost missions utilizing passive control methods such as gravity gradient stabilization. Using attitude sensor measurements to determine the current attitude, the FSW compares that attitude to the commanded (i.e., desired) attitude and determines an attitude error. For spacecraft employing gyros (an attitude sensor that measures the change in spacecraft attitude during a set time period, as opposed to measuring the spacecraft's absolute orientation with respect to inertial space), a Kalman filter usually is utilized both to calculate the current attitude and to calibrate the gyro's drift bias (which ramps with time) relative to an absolute attitude sensor, such as a star tracker. The attitude error is estimated and fed into a control law that calculates on each control cycle what attitude actuator commands (e.g., reaction-wheel control torques) must be generated in order to null the error. On the next control cycle, feedback from the attitude sensors provides the information needed to determine how good a job of reducing attitude error the previous cycle's actuator commands did, as well as how much new attitude error has been introduced this cycle by external perturbative torques. Although this description of onboard attitude control has implicitly addressed maintenance of a constant commanded attitude in the presence of perturbative torques, it can equally well be applied to the execution of large desired attitude changes, called *slews*. Slews can be dealt with two ways. First, the FSW can calculate the amount of attitude change to be performed during a given control cycle, and modify the previous commanded attitude to reflect that change, which would then be used directly as part of that control cycle's attitude error. A second approach is simply to make the commanded attitude the slew target attitude. Although the control law would not be able to null that very large error (say, 90°) in one control cycle, by limiting the size of the commanded control torques in a given control cycle the FSW could gradually work off the error over a series of control cycles, eventually reaching the slew's target attitude.

Operational Scenarios and Agent Interactions

To show more fully how the Remote Agent implementation introduced in Chap. 6 would work in an actual on-orbit situation in the uncrewed sciencemission context, this chapter provides a series of operational scenarios that illustrate the interaction of agents among themselves onboard, the interaction of flight-based agents with ground-based agents, and the interaction of members of a spacecraft constellation with each other.

B.1 Onboard Remote Agent Interaction Scenario

To illustrate the behavior of flight system Remote Agents (incorporating the full FSW subsystems and functionality discussed previously) cooperating to achieve a mission objective, consider the operational scenario defined by the following somewhat simplified assumptions:

- 1. The mission type is a Lagrangian-L2 celestial pointer.
- 2. The primary mission goal is to observe all ground-specified targets while minimizing fuel expenditure so as to maximize mission lifetime.
 - (a) The ground will group observations in clusters
 - (b) The ground defines the nominal order in which the observations within a cluster are performed
 - (c) The FSW defines the cluster order, subject to the following restrictions:
 - (i) If not prohibited by another restriction, on completion of all observations within a cluster, transition to the cluster whose first observation is closest to the final spacecraft attitude on completion of observations within the current cluster.
 - (ii) If the new cluster's observations cannot be completed before an angular-momentum dump is required, select the nearest cluster that can be completed, subject to the momentum restriction.

236 B Operational Scenarios and Agent Interactions

- (iii) If no cluster satisfies restriction (2), perform an angular momentum dump and slew to the nearest remaining unobserved cluster.
- (iv) If an angular-momentum dump is needed while still observing within a cluster, pause science and execute the dump on current exposure completion.
- 3. The secondary mission goal is to survey opportunistically ground-specified areas of the celestial sphere for new targets. These areas have been subdivided by the ground into small regions of equal size and shape.
 - (a) The FSW may schedule surveys of a region if no part of the region is further than a database-specified angle from the current pointing direction.
 - (b) No more than 2 h out of any 24-h period may be spent surveying.
 - (c) If an angular-momentum dump must be performed beforehand to ensure that the survey can be performed without interruption, scheduling the survey is forbidden. If a dump is needed while observing the region, pause science and execute the dump on current exposure completion.
- 4. Ground-issued realtime commands with an attached time window will be scheduled by the FSW within that window, so as, if possible, not to interfere with ongoing activities. On completion of the ongoing activity, the realtime command is executed. If the ongoing activity will not complete by the end of the window, the realtime command is executed by the FSW no later than a database-specified time period prior to the expiration of the window. If the duration of the window is zero seconds, the realtime command is executed immediately upon receipt. If a realtime command needs to be executed prior to the completion of the ongoing activity, the procedure for interruption of the activity and its subsequent treatment is specified by the onboard smart fault detection, diagnosis, isolation, and correction (SFDDIC) Agent.
- 5. Sun angle constraints may not be violated, either when slewing to a target or when observing a target. If any target within a cluster is in violation of this constraint, all targets within the cluster are considered to be in violation. If any part of a region is in violation of this constraint, the entire region is considered to be in violation.
- 6. A constraint on executing science observations is that SI calibrations must meet observation-specific accuracy requirements prior to initiation of the observation.

Relative to the previous assumptions, the following is a characteristic example of how FSW processing would function in performance of typical daily activities. Having completed the last science observation in the current cluster, the scheduling agent requests that the SFDDIC Agent validate the remaining clusters relative to observing constraints. SFDDIC reports back that all remaining clusters are valid. Scheduling then asks the data monitoring and trending agent to identify those clusters that can be observed successfully given the current state of SI calibration. Monitoring and trending reports back that current calibration accuracy is insufficient to support successful observations at the nearest cluster, but is satisfactory at the remaining clusters. Planning and scheduling determines that the ground-specified priority attached to the nearest cluster is not high enough to justify scheduling an SI calibration update at this time, and instead directs that the attitude control agent generate appropriate commanding to produce a slew to the first target in the next nearest cluster.

To effect a slew to the next target requires reaction wheel commanding, so the attitude control agent's commands must pass through an executive agent that interfaces with the FSW backbone. The backbone accepts the slew directive from the agent community and interfaces with the reaction wheels to effect the slew. Following successful completion of the slew (as determined by the SFDDIC Agent), the ACS software in the backbone facilitates entry into fine-pointing mode by activating the quaternion star trackers (QSTs) and fine error sensor (FES). Once the necessary QST and FES data are available, the fine attitude determination agent begins computing high accuracy attitude products. The data are simply stored in a file associated with the observation (managed by the SI data-storage agent) and may be accessed by all users requiring the information. At the same time, and regularly before arrival at the target, the orbit determination agent has produced a steady stream (once a second) of spacecraft, Solar position and velocity vectors, and trends position and velocity vectors, again in support of applications needing the information.

In particular, the attitude control agent uses both the high accuracy attitude and orbit data to generate high precision attitude control commands that, again, are passed through the executive agent to the FSW backbone, which in turn interfaces with the reaction wheels (after the backbone quality assures (QAs) the commands) to produce the desired pointing performance. Having established a stable platform at the target attitude, the planning and scheduling agent directs the acquisition of the science target by the SI and initiation of the science observation specified by the ground. To this end, it notifies the SI commanding-and-configuration agent to effect the necessary SI adjustments required to perform the desired science activity. The SI commanding-and-configuration agent generates the associated commanding and forwards it to the executive agent, which communicates the hardware changes to the FSW backbone so it can directly command the SI. All data output from the SI, whether from the target acquisition or execution of the science observation itself, is stored in the science observation file.

Specialized processing of the first SI data is performed by the SI data processing agent to support target acquisition. Once the target has been acquired successfully, the activity proceeds to the science observation itself, which is processed onboard to support compact packaging prior to downlink. For example, lossless compression will be performed, and possibly processing to reject errors caused by cosmic rays. As the observation continues, the SI data-storage agent will progressively build the observation file so that, by the end of the activity, a complete, coherent record of the observation will have been compiled for downlink as a unified file under the authority of the SI data-storage and communications agent.

When "data-take" at this target has been completed, the planning and scheduling agent determines that a survey observation can be performed conveniently and schedules the necessary slew to the target. Processing flow then proceeds as described above for the earlier target, except following completion of data collection from the survey activity, the SI data processing agent processes the survey data and identifies several point-targets of interest.

These targets are reported to planning and scheduling, which then adds them to the target list for immediate revisits, with observations to be performed according to canned, ground-specified scripts. The targets are then visited in an order defined by the planning and scheduling agent. At the end of this activity, the SFDDIC (using data from the look-ahead modeling agent) determines that a momentum dump is required. Planning and scheduling is notified, which decides that the dump should be performed now, and issues the necessary directive to the orbit maneuvering agent, which also handles the thruster commanding for momentum reduction. As with the attitude control agent, the orbit maneuvering agent must forward its thruster commands to the executive agent, which in turn relays them to the FSW backbone, which then performs its QA and communicates directly to the thrusters to cause the thrusters to fire in the manner specified. Finally, by this time the ground station antenna has become visible to the spacecraft (or vice-versa, depending on one's point of reference), and an electronic handshake between the ground station's and spacecraft's communications agents is established. The handshake is initiated by the ground station agent, but downlink of the science data onboard, including the recently built ground-specified and opportunistic survey files, is managed by the onboard communications agent. The lights-out ground system autonomously validates each file as it is downlinked. As transmission of a file is deemed to be successful, the ground system notifies the onboard communications agent and SI data-storage agent that the onboard addresses associated with that data are free to be overwritten.

This completes the narrative illustrating the mutual communication and interaction of FSW subsystem agents (along with some interaction with ground system agents) in nominal performance of inflight activities of a typical science mission. In reality, the description provided is highly oversimplified: the communication flow in the example is quite sequential, whereas in reality, there will be many parallel conversations in progress at any given time. Also, the assumptions specify a model of a far lower level of complexity than would be characteristic of a real mission. So the example provided should be viewed as simply a token of what would obtain in an actual application.

B.2 Space-to-Ground Dialog Scenario

The dialog in this scenario is initiated by the spacecraft and driven by the following assumptions:

- 1. The mission type is a 1-AU (i.e., 1 *astronomical unit*) drift-orbit survey mission.
- 2. The mission goal is to map out the entire celestial sphere to chart the microwave structure of space. Each mapping will take 6 months, so four mappings will be performed during the 2 year mission lifetime.
- 3. As the onboard antenna size is quite small and transmitting power is highly limited, the deep space network (DSN) must be used for data capture. To reduce transmission costs by reducing downlink volume, the spacecraft will process all raw science data onboard and only downlink science endproducts. The spacecraft will utilize beacon mode and will burst-transmit its processed science data on a low priority basis.
- 4. In the event of major anomalies that the autonomous SFDDIC Agent cannot handle, the spacecraft will notify the ground, downlinking a diagnostic file whose contents characterize the problems encountered.

Relative to these assumptions, consider this scenario for space-ground communications. As the survey work continues, the SI data processing agent processes SI output and forwards the end-product to the SI data-storage agent. In addition, raw SI data are stored in buffers (a precaution against the event of an SI anomaly). When the SFDDIC Agent (in conjunction with the data monitoring and trending agent) validates a given subset of data and declares it to be acceptable, and also verifies nominal SI performance during that time period, the storage locations associated with the raw SI data are designated as available to be over-written.

When sufficient processed survey data have been accumulated to warrant scheduling a downlink, the SI data-storage and communications agent forwards to the executive agent a request to turn on the transmitter to contact the DSN and request a downlink opportunity. This request is then relayed to the FSW backbone, which activates the transmitter, establishing contact with DSN. The DSN informs the spacecraft of its telemetry window. After the start of the window, the onboard agent then downlinks all available, validated SI end-products. The lights-out ground system automatically verifies that all data received from the spacecraft in this pass are intact (i.e., have not been corrupted in transmission) and notifies the onboard SI data-storage agent that the memory areas used for storage of the telemetered science data may now be overwritten.

Sometime later, after this conversation has terminated, the spacecraft loses (at least temporarily) one of its four reaction wheels. The FSW backbone responds by transitioning both the platform and payload to safemode. While in safemode, the SFDDIC (in conjunction with data monitoring and trending, as well as look-ahead modeling) evaluates the situation, both with respect to the failed component and the overall spacecraft state. SFDDIC concludes that the spacecraft must remain in safemode pending consultation with ground resources, and establishes an emergency communications link with the ground via DSN using the procedure already discussed above.

Once a link is established, the SI data-storage and communications agent dumps all recent data stored onboard to the ground system. The agent also provides the results of the SFDDIC Agent's analysis as a starting point for the ground system's more definitive trouble shooting, to be conducted by an integrated team consisting of senior system engineers supported by the ground system's intelligent software agents. As the ground system's work proceeds, requests for additional data from the spacecraft may be made via more regular and frequent DSN contacts. As new ideas are developed and need to be experimented with onboard, the onboard agents may well join the ground system team and participate in a more active fashion until the problem is resolved and nominal function is restored.

This completes the narrative illustrating space-to-ground dialogs initiated by the flight system in nominal performance of typical inflight activities. In reality, the description provided is somewhat oversimplified, as the communication flow in the example is sparse and sequential, whereas in reality, communications will probably be more frequent and there may be parallel conversations in progress during a single contact. Also, the assumptions specify a model of far less complexity than that characteristic of a real mission. So the example provided should be viewed as simply a token of what would be obtained in an actual application.

B.3 Ground-to-Space Dialog Scenario

The interaction in this scenario is initiated by the ground station. Consider a ground-space agent dialog driven by the following assumptions:

- 1. The mission type is LEO Earth-pointer.
- 2. The spacecraft determines its own orbit via global positioning system (GPS). Orbit stationkeeping maneuvers are performed autonomously onboard.
- 3. The mission goal is to observe all ground-specified targets while minimizing fuel expenditure so as to maximize mission lifetime.
 - (a) The spacecraft is provided with Earth coordinates of targets desired to be observed. Repeated observations are performed every 16 days. The spacecraft autonomously determines when the targets may be viewed during a 16-day cycle.
 - (b) The ground maintains onboard a set of observing scenario templates. Each target will be observed using one of those templates. The templates are populated by ground-alterable parameters controlling the observing process.

- 4. Targets are acquired autonomously by the spacecraft.
 - (a) The spacecraft utilizes "quick-look" realtime image data to verify target acquisition and to determine whether targets are obscured by cloud cover such that data collection is useless during this pass.
 - (b) The spacecraft uses the results of its analysis of quick-look data to transition SI configuration autonomously to nominal high data rate mode if target conditions are suitable for the science observation to commence.
 - (c) Pattern recognition is performed onboard, as necessary, to support the observation.
- 5. The spacecraft autonomously generates H&S commanding where necessary (e.g., SI re-configuration at SAA entrance/exit).

Each day for each ground station pass, the ground initiates contact with the spacecraft for the purpose of receiving science and engineering/diagnostic data. During ground-selected passes, the ground system uplinks to the spacecraft an updated target list, as well as changes to parameters controlling the science observations.

At each orbit, the planning and scheduling agent uses data supplied by the orbit determination agent to identify which ground targets can be observed and when. At a database-specified lead-time prior to encountering the target, planning and scheduling notifies the SI commanding-and-configuration agent of the need to configure the SI for use according to the state specified by the template. SI commanding and configuration then generates the appropriate SI commanding and forwards its requests to the executive agent, which relays the package to the FSW backbone, which (following its own command validation) ships the commands to the SI.

When a target is encountered, the SI data processing agent examines the initial quick-look data to verify that the observation can be performed. The agent passes its assessment to SI commanding and configuration and, if the conditions are suitable, commanding and configuration issues the necessary directives to initiate generation of high volume data. The directives are then relayed as before to the backbone so that the required adjustments can be made. Similarly, if special fine SI adjustments are needed to home-in on a specific landmark or feature, SI data processing performs the necessary pattern recognition function and informs SI commanding and configuration of its results. As science data are output from the SI, an observation file is constructed by the SI data-storage and communication agent for downlink when requested by the ground system.

While these observations are performed, parallel onboard processing (controlled by the data-monitoring and trending agent in conjunction with SFD-DIC) determines when the orbit requires correction and informs the planning and scheduling agent, which in turn schedules an orbit stationkeeping maneuver (commanding for which is generated by the orbit-maneuvering agent at the request of planning and scheduling) so as not to conflict with upcoming science data takes. The same agents involved in the orbit correction evaluation also determine when SI reconfigurations need to be performed in response to SAA events. In both cases, agent-generated thruster and SI commands must be forwarded to the executive agent and relayed to the backbone, which QAs the commands and issues the commands directly to the appropriate hardware.

Finally, any windowed realtime commands issued by the ground system are processed by the planning and scheduling agent and inserted into the timeline of onboard commands/activities as necessary. Once scheduled, these realtime commands are treated onboard the same as any commands internally generated.

This completes the narrative illustrating ground-to-space dialogs initiated by the ground system in nominal performance of typical inflight activities. In reality, the description provided is somewhat oversimplified, as the communication flow in the example is sparse and sequential, whereas in reality, communications will be more frequent and there will be parallel conversations in progress during a single contact. Also, the assumptions specify a model of far lower complexity than that characteristic of a real mission. So the example provided should be viewed as simply a token of what would obtain in an actual application.

B.4 Spacecraft Constellation Interactions Scenario

While the scenario discussions above were restricted to cooperative efforts between Remote Agents on a single spacecraft or with counterpart agents in the ground system, this subsection examines the far more elaborate topic of integrating the efforts of multiple teams of agents on several spacecraft (as well as the ground). This higher level of complexity introduces a whole new set of issues unique to constellation work, including the following:

- 1. Is a moderating agent/entity required to facilitate and referee dialogs among the members of a constellation?
- 2. Although all members of a constellation need to be aware of the results from a constellation dialog, how do you decide which members should be direct participants in a given dialog, and "who" makes that decision?
- 3. How are dialogs created, i.e., how are the topics for a dialog selected? For example, does a constellation member with a "problem" simply call a "town meeting," does it submit its problem to a moderator for consideration, do problems un-resolvable by a single member get referred to the ground system to be dealt with, etc.?
- 4. As a dialog progresses, can constellation members "casually" drop in and drop out? If so, is an ongoing record maintained as the dialog proceeds so newly arriving or returning members can quickly get back up to speed? If so, how is the record maintained and by whom?
- 5. Are all dialogs short-term things with well-defined starts and endings, or can a dialog extend over a long time duration, with gaps in activity

interspersed among periods of very high activity? In other words, is there the equivalent of regularly scheduled meetings?

- 6. Must all dialogs be originated from a well-defined issue or concern, or can some arise from general topics of interest, for example, as a mechanism for developing and/or sharing knowledge among the various constellation members? Note the relationship of this question to question #5.
- 7. How does a new constellation member join the community of agent groups? A new spacecraft may have improved methodologies or algorithms useful to the entire constellation. Would an agent dialog be initiated so it can share its new knowledge with the entire constellation?
- 8. Since the existing constellation members may "learn" new things in the course of their day-to-day activities, how do they share this knowledge with the other members so it becomes generally available? Prior to its absorption by the rest of the constellation, would other members be responsible for validating or QAing that new knowledge?
- 9. Once the knowledge base of the constellation grows beyond its original aslaunched confines, how does this new "school-of-hard-knocks" knowledge get passed on to new members?
- 10. How does the constellation resolve differences between the new knowledge carried by new members vs. the empirical experience gained inflight by the old members?
- 11. Could different members of the constellation take on special roles to achieve general constellation goals? For example, could a subset of constellation members (equipped with higher capacity flight computers) conduct simulations or long-term studies of interest to the constellation as a whole?
- 12. For a communications satellite constellation, one can assume all members of the constellation will be at least compatible, if not identical. So interface incompatibility among members should not be a problem. However, suppose in the future it is desirable to form at least a temporary constellation of science satellites to engage in an observation campaign to study a celestial object or phenomenon of great scientific import. By what means could this goal be enabled or facilitated through the use of Remote Agents, and what special interface/architecture issues obtain?

The full impact of these issues for constellations is beyond the scope of this book. Here we will simply consider a high-level (and perhaps overly simplified) scenario for constellation member interactions purely for the purpose of illustrating how constellation behavior might be supported by the proposed FSW design. To this end, assume the following is the case:

- 1. The constellation consists of a set of 16 LEO Sun-synchronous and four GEO weather satellites (the fourth is a spare).
- 2. The OBCs of the larger GEOs have oversized processing power to enable transfer of LEO overflow processing, conduct background long-term studies and simulations, and manage overall direction of constellation behavior.

- 3. The GEOs rarely are replaced. The small LEOs are replaced fairly frequently due to orbit deterioration (to reduce LEO costs, it is assumed that onboard propulsion capacity is weak and fuel limits lifetime to 2 years).
- 4. The ground stations only communicate with GEOs. The GEOs communicate among themselves, the ground stations, and the LEOs. The LEOs only communicate with GEOs.
- 5. The LEOs are fairly primitive, and can be viewed somewhat simplistically as only possessing a FSW backbone. The GEOs possess the full range of Remote Agent functionality discussed previously. The ground system also is equipped with autonomous agents for lights-out operation.
- 6. In support of the higher-level constellation goals, the LEOs' jobs are to pass their SI data to the GEOs and accept commanding and changes in mission objectives from the GEOs. Otherwise, the LEOs operate in a purely local manner, little different from the behavior of a ground-controlled single LEO. They collect their science data, perform necessary housekeeping functions (largely in response to external directives from the GEOs), and conduct very simple FDC functions.
- 7. The GEOs' job is to collect SI data from their equatorial orbits, communicate with the ground via the ground stations with which they are in permanent contact, and run the constellation. Running the constellation includes managing the collection of science data (including all LEO commanding), monitoring and trending all H&S data from the LEOs and GEOs, performing relatively short-term continuing process improvement (CPI) type analytical studies to increase operational efficiency, and employing SFD-DIC Agents to deal with some anomalies that cannot be handled optimally by FDC in the FSW backbones of the LEOs and GEOs.

The ground's job is to receive and archive all science data generated by the constellation, perform long-term analytical studies to increase operational efficiency, and support the GEOs in dealing with major inflight anomalies or failures.

A typical operational scenario is now presented, where some of the agent interactions within a single spacecraft are glossed over in favor of a cleaner description of spacecraft-spacecraft dialogs. At the start of this scenario, the ground has just generated calibration updates for several of the SIs and updated SI observation templates for the LEO and GEO spacecraft. These data are uplinked to the GEOs with instructions to implement the updates as soon as possible without interfering with ongoing observations by the individual spacecraft, while maintaining consistency (as much as possible) among observations by those spacecraft. The planning and scheduling agents of the three GEO spacecraft caucus (by way of a "teleconference" established by their communications agents) and examine their anticipated processor loading (in consultation with their monitoring and trending and look-ahead modeling agents) over the next few hours. In this case, they determine that none of the three GEOs on their own can easily accommodate this planning and scheduling assignment within the high priority time demands specified by the ground without impact to their own science responsibilities. Rather than attempting to segment the planning and scheduling work, parceling out pieces to the individual GEOs, and assigning one of the three GEOs to coordinate the effort, they decide instead to utilize the idle processing power of the fourth (spare) GEO and assign the job to GEO-4, leaving to GEO-1 responsibility for interfacing with GEO-4 when its task is completed.

While GEO-4 is performing the new intermediate-term planning and scheduling task, the other GEOs concentrate on their immediate routine jobs, namely receiving science data from LEOs for relaying to the ground and performing their own science assignments (and communicating the results to the ground) according to their current operating instructions.

For the first duty, the GEOs function in a manner quite similar to the traditional ground station, which has knowledge of the time and angle at which to view any given LEO as it comes into view over the horizon, thereby starting a "view period" during which communications can take place. For the GEOs to perform their similar communications function with the LEOs, switching curves (in latitude and longitude) are maintained onboard the GEOs for use by the planning and scheduling agents in conjunction with the look-ahead modeling and data-monitoring and trending agents. The switching curves divide the "sky" into three 120° slices (where, it may be noted, none of the GEOs will be able to "see" the LEOs near the north and south poles). The curves defining the segment boundaries are padded so that when a LEO enters a padded region, preparations to initiate communications with the new GEO are begun and completed before the LEO leaves the padded region. To effect change in control, the GEO currently directing the LEO's actions instructs (via its communications agent) the LEO to terminate its current telemetry link and reorient its main antenna toward the new GEO. The new GEO then initiates a link with the LEO and requests that flow of completed science products be renewed. Data received from LEOs are then formatted by the individual GEOs in observation files by the SI data-storage and communications agents and are relayed to a designated single GEO (say, GEO-3) for integration into an overall global picture/assessment. As the GEOs themselves conduct their own science observations (as discussed in a previous section), the data from GEO-1 and GEO-2 are relayed to GEO-3 for merger with the LEO data.

When GEO-3 has completed the integration process, it also performs the data reduction processing required to convert the raw measurements into science end products. These results are then transmitted to the lights-out ground station for archiving and dissemination. Optionally, the raw measurements themselves may be transmitted to the ground for archiving. Note that GEO-3 does not integrate and process all the data all the time. Once GEO-3 begins its integration job, another of the three GEOs (say GEO-2) will be the collection point for new data as it is generated by the constellation. So at any given time, one GEO will be collecting science data from the constellation,
one GEO will be integrating and processing data, and a third GEO will be available for planning and scheduling work, which in this case was assigned to GEO-4, leaving GEO-1 available to respond to other ground requests as well as communications from GEO-4.

Once GEO-4 has completed its planning and scheduling work, it communicates its results to GEO-1. GEO-1 then passes the plan for installing the new calibrations and procedures to the other GEOs. Each GEO then instructs those LEOs under its control as to the changes that should be made. The time at which those changes should be made is also specified to ensure that when the next mass of data from the whole constellation is integrated, all data in the mix would have been generated, ideally, using the same calibrations and procedures. If that ideal condition is not possible, then any data obtained using old calibrations/procedures will, when contributed to integration, be so identified in the data/results transmitted to the ground.

B.5 Agent-Based Satellite Constellation Control Scenario

Consider the scenario when one has many satellites with different viewing capabilities (IR, visible, or UV) orbiting a planet and one desires a full spectrum sweep of a certain portion of the planet. Traditionally, science team members and human controllers would need to identify the satellites with each different capability that will be making a pass over the section of the planet indicated. Human controllers would need to form a detailed, possibly quite intricate plan for the needed observations and to organize a series of requests addressed to the satellites to perform the sweep, with all relevant details down to the transmission of the data back to earth. This type of activity entails inefficiencies and represents a questionable or wasteful use of manpower.

Through the use of an agent community that hierarchically and intelligently parses instructions, this could be done much more efficiently. Ideally, the human operator needs only to transmit a command similar to "Take a full spectrum picture of the area bounded by given latitude and longitude data and transmit the picture back in 2 days." An executive level satellite could receive this information, decompose it, and then negotiate with the agent community (where each satellite in orbit is part of the community) to attempt to schedule a plan. From there, each satellite could respond with information such as "will be passing over the site in 36 h, I can take the picture in IR" or "will not be passing over the site for another 96 h, I cannot take the picture." Certain constraints may come into play also; for example, UV and visible light sensors are only useful when it is not night time at the given site. Responses of this nature may be similar to "will be passing over the site in 4 h, but the site is currently on the dark side of the planet" or "will be passing over in 4 h when site is on dark side, but will pass over again in 20 h when it is local noon." Of course, there are certain requests that just cannot be fulfilled. It is the executive's job to notice these, come up with the "closest fit" to the request issued from the human controller, and report back with the closest fit to ask for a go-ahead on that schedule.

When all of the planning has been performed through the negotiation, the executive satellite could issue the plan to all image gathering satellites. The satellites will receive their plans, and internally they will schedule their own control (perhaps via an internal agent network for subsystem control) for setting up their imaging systems, recording the image, and transmitting it. The satellite will pass over the section of the planet in due course, record the images, and transmit them back to the executive satellite. The executive satellite will assemble the images when the whole spectrum has been covered, and transmit them back to the human controller at the next appropriate opportunity.

This scenario is a perfect illustration of a negotiating agent network. Although it is hypothetical, it unifies the concepts of hierarchical parsing networks between spacecraft, and within spacecraft.

B.6 Scenario Issues

One aspect of the presented design concepts that could create problems in flight is the layering of communications engendered by the FSW backbone. Commands, status messages, and data often must go through several checkpoints before they reach their intended hardware destination in order to guarantee platform and/or pavload H&S, no matter what anomaly or failure conditions may obtain with the Remote Agents. There are two potential problems arising from this security paradigm. First, some commanding has associated with it very severe timing requirements. For example, in the case of HST, in order to satisfy its pointing accuracy and stability requirements, reaction wheel commands must be executed within 7 ms of receipt of the gyro data from which the commands are derived. For this case, a data latency problem engendered from the time delays in relaying commands to the reaction wheels potentially could jeopardize meeting a fundamental mission requirement, unless the OBC and bus infrastructure are adequate to support the design. Second, the multiplication of messages and commands, especially if receipt of one always triggers an acknowledgement, could create a blizzard of traffic on the bus (or busses), leading to loss of information or even processing lock-up, again, unless the OBC and bus infrastructure are adequate to support the design. In other words, the presented design concepts, involving communications-and-computation-intensive negotiation processes between multiple agents, imply the need for research into new spacecraft architectures and the crucial need for certain minimum levels of performance of future onboard communications and computing resources.

Acronyms

AC	Autonomic computing
ACE	Attitude control electronics
ACL	Agent communication language
ACS	Attitude control subsystem
ACT	Agent concepts testbed
AE	Autonomic element
AFLOAT	An agent-based flight operations associate
AI	Artificial intelligence
AIFA	Archive interface agent
AM	Autonomic manager
AMS	Active middleware service
ANS	Autonomic nervous system
ANTS	Autonomous nano technology swarm
AOS	Acquisition of signal
ARPA	Advanced Research Projects Agency
ASM	All sky monitor
BAT	Burst alert telescope
BN	Bayesian networks
C&DH	Command and data handling
CBR	Case based reasoning
CCC	Constellation Control Center
CCD	Charge-coupled device
CGRO	Compton Gamma Ray Observatory
CIM	Common information model
CLIPS	C-language integrated production system
CMA	Contact manager agent
COTS	Commercial off the shelf
CPU	Central processing unit
CSS	Coarse Sun sensor
DARPA	Defense Advanced Research Projects Agency
DBIFA	Database interface agent

DIG	Distributed interactive simulation
	Distributed interactive simulation
DS DS1	Deep Space
DSI	Deep Space I
DSN	Digital gun gangan
EO 1	Farth Observing 1
EO-I FOS	Earth observing system
FD	Explorer platform
FIIVE	Explorer platform
FAST	Formal approaches to swarm technologies
FDC	Formal approaches to swarm technologies
FES	Fine error sensor
FGS	Fine guidance sensor
FIPA	Foundations of Intelligent Physical Agents
FIRE	Fault isolation and resolution expert
FOT	Flight operations team
FOV	Field of view
FSS	Fine Sun sensor
FSW	Flight software
FTC	Fault tolerant computing
GCI	Geocentric inertial
GEO	Geosynchronous earth orbit
GIFA	GenSAA/genie interface agent
GN&C	Guidance, navigation, and control
GPM	Global precipitation measurement mission
GPS	Global positioning system
GRB	Gamma ray burst
GSFC	Goddard Space Flight Center
GSS	Ground station simulator or generalized support software or
	ground support system or ground support software
GTDS	Goddard trajectory determination system
H/W	Hardware
H&S	Health and safety
HBM	Heart-beat monitor
HEAO	High Energy Astronomical Observatory
HGA	High gain antenna
HLA	High level architecture
HST	Hubble Space Telescope
I/O	Input/output
IBM	International business machines
ICT	Information and communications technology
IMU	Inertial measurement unit
IR	Infrared
IRU	Inertial reference unit

ISA Interface services agent

250 C Acronyms

IT	Information technology
IUE	International ultraviolet explorer
$_{\rm JPL}$	Jet propulsion lab
JWST	James Webb Space Telescope
KIF	Knowledge interchange format
KQML	Knowledge query and manipulation language
KSE	Knowledge sharing effort
KX	Kinesthetics eXtreme
LARA	Lander amorphous rover antenna
LEO	Low earth orbit
LISA	Laser interferometer space antenna
LOG	Log agent
LOS	Loss of signal
MA	Multiple access
MAP	Microwave anisotropy probe
MAPE	Monitor, analyze, plan and execute
MAS	Multi-agent system
MC	Managed component
MIDEX	Medium-class explorer
MIFA	MOPSS Interface Agent
MMA	Mission manager agent
MMS	Magnetospheric multiscale
MOCC	Mission Operations Control Center
MOPSS	Mission Operations Planning and Scheduling System
MTB	Magnetic torquer bar
NASA	National Aeronautics and Space Administration
NFI	Narrow field instruments
NIR	Near-infrared
NSSC	NASA Standard Spacecraft Computer
OAO	Orbiting Astronomical Observatory
OBC	Onboard computer
OPE	Observation plan execution
OSO	Orbiting Solar Observatory
OTA	Optical telescope assembly
PAGER	Pager interface agent
PAM	Prospecting Asteroid Mission
PBM	Pulse-beat monitor
PCA	Proportional counter array
PCS	Pointing control subsystem
PID	Proportional-integral-derivative
PSA	Planner/scheduler agent
QA	Quality assurance
QST	Quaternion star tracker
R&D	Research and development
RF	Radio frequency

RSDO	Rapid Spacecraft Development Office
RTS	Relative time sequence
RXTE	Rossi X-ray timing explorer
SA	Solar array
SAA	South Atlantic Anomaly
SAC	Situated and autonomic communications
SAMPEX	Solar Anomalous and Magnetospheric Particle Explorer
SARA	Saturn autonomous ring array
SC	Spacecraft
SDO	Solar dynamics observatory
SFDDIC	Smart fault detection, diagnosis, isolation, & correction
SI	Science instrument
SMEX	Small explorer
SMM	Solar Maximum Mission
SMP	Statistics Monitor Program
SSA	S-band single access
SSA	System services agent
SSR	Solid state recorder
SysMMA	System monitoring and management agent
TAM	Three axis magnetometer
TBS	To be specified
TCO	Total cost of ownership
TDRS	Tracking and data and relay satellite
TDRSS	TDRS system
TMON	Telemetry monitor
TOO	Target of opportunity
TRMM	Tropical Rainfall Measuring Mission
TSM	Telemetry and statistics monitor
TWINS	Two wide-angle imaging neutral-atom spectrometers
UARS	Upper Atmosphere Research Satellite
UI	User interface
UIA	User interface agent
UIFA	User interface agent
UV	Ultraviolet
UVOT	UV/optical telescopoe
VIFA	VisAGE interface agent
WMAP	Willsinson microwave anisotropy probe
XRT	X-ray telescope

Glossary

Angular momentum dump See momentum dump.

- Attitude The orientation of the spacecraft in inertial space. Usually, attitude defines the orientation of all three spacecraft axes with respect to an inertial reference frame, though for spin-stabilized spacecraft, it often is the case that only the orientation of the spin axis is specified.
- Attitude actuator A control hardware component that generates control torques required to maintain spacecraft stability, null attitude errors, reorient the spacecraft to a new orientation with respect to an inertial reference frame, etc.
- Attitude control The mechanism for establishing and maintaining a desired spacecraft orientation with respect to an inertial reference frame.
- Attitude control accuracy A quantitative measurement of the error in maintaining the spacecraft attitude at its desired orientation with respect to an inertial reference frame.
- Attitude control torque Torques intentionally applied to the spacecraft to maintain or establish a desired spacecraft orientation with respect to an inertial reference frame.
- Attitude determination The computation of the spacecraft orientation relative to a specified reference frame. For celestial-pointing spacecraft, the reference frame is usually either the geocentric inertial (GCI) frame or the heliocentric inertial frame.
- Attitude determination accuracy A quantitative measurement of the error in the computed spacecraft attitude.
- **Attitude dynamics** The study of a spacecraft's motion about its center of mass.
- **Attitude maneuver** A commanded change in the spacecraft's desired attitude, as opposed to a torque applied to null errors in the actual attitude relative to the desired attitude. A large attitude maneuver is called a *slew*.
- Attitude matrix A specification of the spacecraft orientation in direction cosine matrix format. The direction cosine matrix is a 3×3 square matrix. The nine elements of the matrix are the cosines of the angles between the

D

three spacecraft-body unit vectors and the three reference axes relative to which the spacecraft orientation is defined.

- **Attitude quaternion** A specification of the spacecraft orientation in quaternion format.
- Attitude sensors Spacecraft hardware and electronics providing measurement data that can be used to determine the spacecraft orientation or changes in the spacecraft orientation.
- **Autonomic** Of or pertaining to the capacity of a system to control its own internal state and operational condition.
- Autonomic communications A research field with the same motivators as the autonomic computing concept with particular focus on the communications research and development community (see SAC).
- Autonomic computing Overarching initiative to create self-managing computer-based systems, with a metaphor inspired by the biological autonomic nervous system.
- Autonomic element An autonomic manager and a managed component considered together.
- Autonomic manager A control loop and components to provide autonomic self-* for the managed component.
- Autonomic systems Often used synonymously with autonomic computing; sometimes used as a synonym for autonomicity from a systems perspective, and sometimes used in the sense that AutonomicSystems = AutonomicComputing + AutonomicCommunications.
- Autonomicity The quality of having an autonomic capability.
- **Autonomy** A system's capacity to act according to its own goals, percepts, internal states, and knowledge, without outside intervention.
- **Celestial-pointer** A spacecraft whose fine-pointing science instruments are oriented toward "fixed" points on the celestial sphere. For example, a spacecraft that slews from one attitude to another to observe a series of X-ray point-sources would be a member of the class of celestial pointers.
- **Center of mass** Average position of a system of objects, weighted in proportion to their masses.
- **Charge coupled device star tracker** A star tracker that detects stars by digitally scanning an array of photosensitive components (pixels). The star tracker integrates the electrical charge in the pixels "struck" by (for example) starlight. The measurement is performed by reading the pixel output line by line.
- **Coarse Sun sensor** A device that measures photocell output as a function of Sun angle. Since the amount of energy hitting the photocell varies as the cosine of the Sun angle (normal incidence for null Sun angle, grazing incidence for 90° Sun angle), CSSs are also referred to as cosine detectors. Note that this scheme provides an analog representation of the Sun angle.
- **Command quaternion** Desired spacecraft orientation, with respect to an inertial reference frame, expressed in quaternion format.

- **Commanded attitude** Desired spacecraft orientation with respect to an inertial reference frame.
- **Constellation** Two or more spacecraft engaged in coordinated operations with the objective of meeting a set of mission requirements.
- Control loop Closed loop of feedback control.
- **Control torque** Torque generated by spacecraft actuators for the purpose of controlling the spacecraft's attitude.
- **Digital Sun sensor** A device that measures output from a series of photocells to determine Sun angle. If a photocell's output is greater than a threshold, it is considered to be "on." The pattern of "on" photocells is directly associated with the Sun angle. Note that this scheme provides a digital representation of the Sun angle.
- **Distributed satellite systems** Multiple networked spacecraft, analogous to distributed client-server, or networked, computing systems of today, as opposed to the traditional "monolithic" centralized computing environment of the past.
- **Distributed science mission** Multiple networked space assets, analogous to distributed client-server, or networked, computing systems of today, as opposed to the traditional "monolithic" centralized computing environment of the past. The assets could include rovers, stationary instrument packages, and spacecraft.
- **Downlink** A point-to-point RF communications channel carrying data from the spacecraft to the ground.
- **Earth-pointer** A spacecraft whose fine-pointing science instruments are oriented along the nadir vector toward the Earth.
- **Effector** In the context of autonomic management, a defined means to bring about a change to a part of the managed component.
- **Environment awareness** A capability of a system through which the system is continually able to perceive its external operating conditions in relation to its knowledge of its abilities. From this perspective, environment awareness may be considered a part of self-awareness the ability to know one's place in the environment. In another view of environment awareness, the environment is aware of the individuals themselves for instance, through their heartbeat or pulse.
- **Ephemeris** In the context of space missions, a table of positions (of either celestial objects or spacecraft) at a given time (or at given times) in a given coordinate system. Also, colloquially referred to as an empirical (as opposed to algorithmic) specification of a spacecraft's orbital position (either historical or predicted).
- **Fine Sun sensor** A high precision digital Sun sensor (DSS).
- **Formation flying** A flight concept in which multiple spacecraft perform their science operations while keeping a fixed position relative to one another.
- Geomagnetic field The magnetic field of the Earth.

- **Geostationary orbit** An orbit around the Earth that maintains a spacecraft directly above the same point on the Earth's surface at all times. This is necessarily an equatorial orbit (with orbit inclination to the equatorial plane equal to zero degrees) where the orbit altitude is precisely chosen so that its associated period is 24 h, matching the Earth's rotational rate.
- **Geosynchronous orbit** An orbit having the same altitude as a geostationary orbit, but not necessarily maintaining the spacecraft above the same point on the Earth at all times. Geosynchronous orbits may have nonzero inclinations.
- **Gimbal** A rotatable or pivotable hardware contrivance whereby an attached element of a spacecraft, for example a dish antenna or solar array, may be reoriented (in two degrees of freedom) relative to the body of the spacecraft.
- **Global positioning system** A constellation of low Earth orbiting satellites continually broadcasting radio signals supporting realtime, onboard determination of position by spacecraft with compatible receivers.
- **Goddard trajectory determination system** The primary ground software system at NASA Goddard Space Flight Center for computing definitive and predictive spacecraft ephemerides.
- **Gyro** Originally, a mechanical sensor containing a spinning mass that exploits conservation of angular momentum to measure changes in attitude. Recently, spacecraft have employed gyros that utilize nonmechanical structures (for example, hemispheric resonating gyros (HRGs) and laser ring gyros).

Gyro drift A ramping error (varying with time) in a gyro's output.

Gyro scale factor and alignment calibration An operational process consisting of a series of large slews executed to generate gyro data that the ground system uses to determine the gyro scale factors and gyro alignments relative to absolute attitude sensors (such as star trackers).

Gyroscope See Gyro.

HBM Heart-beat monitor

Inertial measurement unit See Gyro.

Inertial reference unit See Gyro.

- Kalman filter A sequential estimator with a fading memory commonly used for realtime spacecraft attitude estimation. Implemented in the ACS flight software of most GSFC spacecraft flying gyros. Enables onboard estimation of both attitude error and gyro drift bias.
- Lagrange point Positions of stable or pseudo-stable gravitational equilibrium within a three-body system consisting of two major bodies and one body of negligible mass relative to the other two. In a gravitational three-body system, there are, relative to one of the two major bodies, always exactly five such points, referred to as L1, L2, L3, L4, and L5, all in the plane of the orbit of the one major body about the other. The three pseudo-stable points (L1, L2, and L3) lie along an axis between the one major body and the other. The two stable points (L4 and L5) are off-axis,

in positions leading and following the one major body on its orbital path around the other major body. An example of clumping of small natural objects at off-axis Lagrange points is the Trojan Asteroids (here, the two large gravitational bodies are the Sun and Jupiter). See Fig. 3.2 for Lagrange Points illustrations.

- Lagrange point orbit The complex, nonplanar motion of a spacecraft when in orbit, the so-called *halo orbit*, near one of the unstable Lagrange Points, L1, L2, or L3. The halo orbit in general cannot be maintained without station-keeping.
- Libration point See Lagrange point.
- **Limit checking** A validation procedure in which a data point value is checked against a threshold (upper, lower, or both).
- Low Earth orbit An orbit whose altitude above the Earth's surface is between about 160 km and about 2,000 km. Atmospheric drag increases dramatically at lower altitudes. A popular LEO altitude range is 500-600 km.
- **Magnetic coil** A wire wrapped about a cylinder in a series of loops; a magnetic field results when the wire carries an electrical current. When the cylindrical space contains a ferromagnetic core, the configuration is called a magnetic torquer bar (MTB).
- **Magnetic disturbance torque** A spacecraft torque generated by the interaction of the spacecraft's residual magnetic dipole as the spacecraft moves along its orbital path through the external magnetic field.
- **Magnetic torque** Torque on the spacecraft arising from the interaction between fields generated by the spacecraft's magnetic coils and the external magnetic field. If the torque arises from an interaction between a spacecraft residual dipole moment and the external magnetic field, the torque is called the *magnetic disturbance torque*.
- Magnetic torquer bar See magnetic coil.
- Magnetometer An attitude sensor measuring the strength and direction of the magnetic field external to the spacecraft. Note that this measurement will include not only the geomagnetic field (for spacecraft in Earth orbit), but also contributions from MTBs and spacecraft residual dipole moment. Processing of the raw magnetometer measurements must remove these other field sources to obtain accurate geomagnetic field measurements.
- Managed component A component that is protected by the autonomic manager.
- **MAPE** Monitor, analyze, plan and execute components within autonomic manager.
- Measured attitude The current best estimate of the spacecraft attitude.
- **Molniya orbit** An orbit designed to support Russian communications satellites. It freezes the orbit's perigee over the Southern hemisphere, tilting the orbit to place the apogee over the Northern latitudes, ensuring maximum communication opportunities between spacecraft and ground stations.

- 258 D Glossary
- Momentum dumping A procedure for spacecraft controlled by reaction or momentum wheels, whereby angular momentum is removed from the wheel system to prevent the wheels from saturating (a circumstance where wheel speeds can be run up to maximum values). Angular momentum dumping utilizes other attitude actuators to shift or remove angular momentum from the wheels (e.g., by applying current to magnetic coils, thereby generating a spacecraft counter-torque by means of coupling to the geomagnetic field).
- **Nadir vector** In the context of the gravitational field of a massive body, the nadir at a given point is the direction along the force of gravity at that point. For Earth-orbiting spacecraft, the vector is directed toward the geocenter. See Zenith.
- **Nadir-pointer** A spacecraft whose primary science instrument is directed along the nadir vector.
- Nanospacecraft A small spacecraft characterized as having a weight of about 10 kg or less, a cylindrical diameter about 30 cm or less, and a low cost.
- Networked science mission See distributed science mission.
- **Occultation** A geometric condition when the view of the target of a spacecraft sensor or science instrument is obstructed by a celestial body.
- Orbit acquisition Achieving mission orbit.
- **Orbit decay** For Earth-orbiting spacecraft, decrease in altitude due to atmospheric drag.
- **Orbit determination** Computation of the orbit of a spacecraft (or celestial body) in inertial space.
- **Orbit determination accuracy** A quantitative measurement of the validity of the computed spacecraft orbit.
- Orbit dynamics The study of the motion of the spacecraft's center of mass.
- **Orbit elements** A set of parameters specifying the size, shape, and orientation of the orbit in inertial space, as well as the location of the spacecraft at a given time (the epoch time). An example is the osculating Keplerian elements.
- **Orbit generator** An algorithm-based computational means of predicting the future position and velocity of body in orbit. Commonly, the algorithm can also be used to compute orbits into the past from a given set of orbital elements.
- **Orbit maneuver** A commanded change in the spacecraft's orbit, accomplished by producing a net thrust force on the spacecraft, generally by firing a thruster. If applied to null errors in the actual orbit relative to the desired orbit, the maneuver is called a *station-keeping maneuver*.
- **Orbit normal** A vector perpendicular to the orbital plane. The normal vector obtained by the *right hand rule*, with the fingers curled in the direction of spacecraft orbital motion, is *positive orbit normal*. The opposite vector is *negative orbit normal*.

- **Orbit propagation** Extrapolation of spacecraft position and velocity vectors from actual measurements at an earlier time.
- **Period** An orbital period is the time it takes a spacecraft (or celestial object) to complete an orbit.
- **Predicted orbit** Extrapolated spacecraft position and velocity (as a function of time) determined by applying mathematical models of physical processes (e.g., the Earth's gravitational potential) to actual measurements of the spacecraft orbit at an earlier time.
- **Propagation** For orbits, extrapolation of a spacecraft position and velocity from a known starting value. For attitudes, use of relative attitude information (from gyros) to extrapolate spacecraft absolute pointing from a starting, measured absolute attitude (for example, from star tracker data).

Propellant Thruster fuel.

- **Pulse-beat monitor** Extension of heart-beat monitor with health urgency tones.
- **Quaternion** A mathematical representational system involving parameterization of the three pieces of information supplied by Euler angles in terms of a four-dimensional unit vector, which facilitates descriptions of spacecraft attitudes. The quaternion lacks the singularity issues present in Euler angle formulations and is more compact than the nine elements of the direction cosine matrix. Quaternions also are easily manipulated to determine (or incorporate) changes in attitude, a useful feature when supporting attitude slews.
- **Reaction wheel** A flywheel rotated with an electric motor, used on a spacecraft to transfer angular momentum to or from the spacecraft body, thereby effecting a change in the spacecraft's attitude without firing a thruster.
- **Reference data** Inertial frame information from a model, catalog, etc. that can be combined with attitude observations in order to determine the spacecraft orientation.
- Self-* Self-managing properties.
- **Self-anticipating** The ability to predict likely outcomes or to simulate self-* actions.
- **Self-awareness** The ability to perceive and compute with its own internal state in relation to its own knowledge and capabilities. Relates to the concept of "Know thy self."
- **Self-chop** The initial four (and generic) self properties (configuration, healing, optimization, and protection).
- **Self-configuring** A system's ability to configure and re-configure itself to meet policies/goals.
- **Self-critical** The ability to consider whether policies are being met or goals are being achieved (see self-reflecting).
- **Self-defining** The ability to reference (and operate in a manner determined by) internal data and the internal definitions of that data (i.e., metadata).

Primarily related to definitions of goals and policies (perhaps as derived from self-reflection).

- **Self-governing** The capability of operating autonomously and being responsible for achieving goals and performing tasks.
- **Self-healing** In the reactive sense, the capability of self-fixing faults; in the proactive sense, the capability of predicting and preventing faults.
- **Self-managing** The capability of operating autonomously and being responsible for wider self-* management.
- **Self-optimizing** A system's capability of dynamically optimizing its own operation.
- **Self-organizing** A system's capability of organizing its own efforts. Often used relative to networks and communications.
- **Self-protecting** A system's capability of protecting itself through perception of potential threats and prediction of outcomes of situations in the environment, and through self-configuring to minimize potential harm.
- **Self-reflecting** The capability of assessing routine and reflex operations of self-* operations and determining whether they are as expected. May involve self-simulation to test scenarios.
- **Self-simulation** The capability of generating and testing possible scenarios without affecting the live system.
- Selfware Self-managing software or firmware.
- **Sensor** In the context of autonomic capabilities, a means to measure a part of the managed component. In an ACS context, a measuring device whose output can be used to determine the spacecraft's attitude, either in absolute or relative terms.
- Situated and autonomic communications Local, self-managed information flows in reacting to environment and context changes. Refers to the communication and networking vision of being task- and knowledgedriven and fully scalable.
- **Slew** A large change in orientation, e.g., a large attitude maneuver by a spacecraft.
- **South Atlantic Anomaly** A region of space near the Earth over the south Atlantic ocean where the van Allen radiation belt makes its closest approach to the Earth's surface, with increased intensity of radiation.
- **Star tracker** A star detecting device used for spacecraft attitude determination and control. Stars registering an intensity above a commanded threshold are detected and tracked until a break-track is ordered. While the star is being tracked, the star tracker measures the location of the star in the field of view (FOV), as well as the star's magnitude. For early star trackers, the output was location and magnitude, but recently quaternion star trackers have been flown that output an attitude quaternion (relative to the tracker frame as opposed to the body frame) directly.
- **Station-keeping maneuver** A spacecraft orbital maneuver performed to null errors in the actual orbit relative to the desired orbit.

- **Stellar-pointer** A spacecraft that conducts its science by pointing a science instrument (or instruments) at targets on the celestial sphere, i.e., stars, galaxies, etc. Stellar-pointers are slewed to a commanded attitude, acquire their target(s), conduct their science, and then slew to the next target(s).
- **Sun sensor** An attitude sensor that detects the presence of the Sun and, in the case of analog and digital Sun sensors, also outputs measurements from which the Sun vector or Sun angle can be computed.
- **Sun-pointer** A spacecraft whose attitude is maintained so as to keep its science instrument(s) oriented toward the Sun.
- Sun-synchronous orbit A spacecraft orbit whose plane rotates at the same rate as the Earth's orbital rate about the Sun. The spacecraft orbital plane motion then matches the apparent motion of the Sun (relative to the Earth). Sun-synchronous spacecraft imaging the Earth will, therefore, (for example) always image points on the Earth's equator at the same local time.
- **Swarm** A large group of autonomous individuals each having simple capabilities, cooperative actions, and no global knowledge of the group's objective.
- **Tachometer** A device that measures the rotation rate of a reaction wheel or momentum wheel.
- **Three-axis attitude** A specification of the orientation of a spacecraft's body axes with respect to a reference frame, typically expressed as a transformation from the reference frame to the body frame.
- **Three-axis magnetometer** An attitude sensor that measures the magnitude and direction of the magnetic field in which the spacecraft is immersed.
- **Three-axis stabilized** A type of spacecraft the orientations of all of whose body axes are controlled; distinct from a spin-stabilized spacecraft, which only controls the orientation of its spin axis.
- Thruster Spacecraft hardware that generates thrust by expelling propellant.
- **Torque** The vector rate of change of the angular momentum. A rigid body experiences a torque if a net force is applied to the body along a line that does not pass through the body's center of mass. Such forces introduce rotations about the object's center of mass.
- **Tracking and data relay satellite** A communications relay satellite in NASA's TDRSS. TDRS satellites are geostationary spacecraft.
- **Tracking and data relay satellite system** The collection of NASA's TDRS satellites, their ground stations, and their control systems.
- **Uplink** An RF communications channel for data flow from the ground to the spacecraft.
- **Virtual mission** A mission that consists of both virtual and real spacecraft assets to meet a mission objective that would otherwise require one or more real spacecraft and their instruments to be designed, constructed, launched, and operated.

- **Virtual platform** A "virtual spacecraft" whose "payload" is composed of the instrument(s) or payload(s) aboard two or more "real" spacecraft. By extension, the virtual payload may also include instruments located at ground-based observatories.
- **Zenith** In the context of the gravitational field of a massive body, the zenith at a given point is the direction opposite to the force of gravity at that point. See Nadir.
- **Zenith-pointer** A spacecraft whose primary science instrument(s) is (are) directed "up" along the local vertical (i.e., along the zenith vector).

References

- A. Aamodt and E. Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7(1):39–59, 1994.
- M. Aarup, K. H. Munch, J. Fuchs, R. Hartmann, and T. Baud. Distributed intelligence for ground space systems. In *Proc. Third International Symposium* on Artificial Intelligence, Robotics, and Automation for Space (I-SAIRAS 94), pages 67–70, Pasadena, CA (USA), 18–20 April 1994. Jet Propulsion Laboratory.
- M. Agarwal, V. Bhat, H. Liu, V. Matossian, V. Putty, C. Schmidt, G. Ahang, L. Zhen, M. Parashar, B. Khargharia, and S. Hariri. AutoMate: Enabling autonomic applications on the grid. In Proc. The Autonomic Computing Workshop, 5th Int. Workshop on Active Middleware Services (AMS'03), Seattle, WA (USA), 25 June 2003.
- 4. S. Aiber, O. Etzion, and S. Wasserkrug. The utilization of AI techniques in the autonomic monitoring and optimization of business objectives. In *IJCAI* Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems, Acapulco (Mexico), 10 August 2003.
- J. Allen. Natural Language Understanding. Benjamin Cummings Publishing Company, 1994.
- E. A. Alluisi. The development of technology for collective training: SIMNET, a case history. In L. D. Voss, editor, A Revolution in Simulation: Distributed Interaction in the '90s and Beyond. Pasha Publications, Arlington, VA (USA), 1991.
- T. Ames and S. Henderson. The workplace distributed processing environment. In C. Hostetter, editor, *Proc. 1993 Goddard Conference on Space Applications of Artificial Intelligence*, pages 181–188, Greenbelt, MD (USA), 10–13 May 1993. NASA Goddard Space Flight Center, NASA Conference Publication 3200.
- A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, 2004.
- 9. T. Balch, A. Feldman, and Z. Khan. Automatic classification of insect behavior using computer vision and behavior recognition. In *Proc. Second International*

Workshop on the Mathematics and Algorithms of Social Insects, Atlanta, Georgia (USA), 16–17 December 2003.

- D. Bantz and D. Frank. Challenges in autonomic personal computing, with some new results in automatic configuration management. In Proc. IEEE Workshop on Autonomic Computing Principles and Architectures (AUCOPA 2003), pages 451–456, Banff, Alberta (Canada), 22–23 August 2003.
- D. F. Bantz, C. Bisdikian, D. Challener, J. P. Karidis, S. Mastrianni, A. Mohindra, D. G. Shea, and M. Vanover. Autonomic personal computing. *IBM Systems Journal*, 42(1):165–176, 2003.
- G. Beni. The concept of cellular robotics. In Proc. 1988 IEEE International Symposium on Intelligent Control, pages 57–62. IEEE Computer Society Press, 1988.
- G. Beni and J. Want. Swarm intelligence. In Proc. Seventh Annual Meeting of the Robotics Society of Japan, pages 425–428, Tokyo (Japan), 1989. RSJ Press.
- K. P. Birman, R. van Renesse, and W. Vogels. Navigating in the storm: Using Astrolabe for distributed self-configuration, monitoring and adaptation. In *The Autonomic Computing Workshop*, 5th Int. Workshop on Active Middleware Services (AMS'03), pages 4–13, Seattle, WA (USA), 25 June 2003.
- D. G. Boden and W. J. Larson. Cost-Effective Space Mission Operations. McGraw-Hill, NY, 1996.
- E. Bonabeau and G. Théraulaz. Swarm smarts. Scientific American, 282: 72–79, 2000.
- A. H. Bond and L. Gasser. *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann Publishers, San Mateo, CA (USA), 1988.
- W. K. Braudaway and S. M. Harkrider. Implementation of the high level architecture into DIS-based legacy simulations. In Proc. Spring 1997 Simulation Interoperability Workshop, Orlando, FL (USA), 1997.
- A. Brown and D. Patterson. Embracing failure: A case for recovery-oriented computing. In *Proc. High Performance Transaction Processing Symposium*, Asilomar, CA, October 2001.
- A. B. Brown, J. Hellerstein, M. Hogstrom, T. Lau, S. Lightstone, P. Shum, and M. P. Yost. Benchmarking autonomic capabilities: Promises and pitfalls. In *Proc. ICAC'04: International Conference on Autonomic Computing*, Hawthorne, NY, pages 266–267, IEEE Computer Society, May 2004.
- S. Carlson. Artificial life: Boids of a feather flock together. Scientific American, 283(5), 2000.
- 22. E. K. Casani, B. Wilson, and R. Ridenoure. The new millennium program: Positioning NASA for ambitious space and earth science missions for the 21st century. *AIP Conference Proceedings*, 361(3):1553–1558, 1996.
- W. Cedeno and D. K. Agrafiotis. Combining particle swarms and k-nearest neighbors for the development of quantitative structure-activity relationships. *International Journal of Computational Research*, 11(4):443–452, 2003.
- H. Chalupsky, T. Finin, R. Fritzson, D. McKay, S. Shapiro, and G. Weiderhold. An overview of KQML: A knowledge query and manipulation language. Technical report, KQML Advisory Group, April 1992.
- 25. G. Cheliotis and C. Kenyon. Autonomic economics: A blueprint for selfmanaged systems. In IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems, Acapulco (Mexico), 10 August 2003.

- 26. S. Chien, R. Sherwood, D. Tran, R. Castano, B. Cichy, A. Davies, G. Rabideau, N. Tang, M. Burl, D. Mandl, S. Frye, J. Hengemihle, J. D'Agostino, R. Bote, B. Trout, S. Shulman, S. Ungar, J. Van Gaasbeck, D. Boyer, M. Griffin, H. hua Burke, R. Greeley, T. Doggett, K. Williams, V. Baker, and J. Dohm. Autonomous science on the EO-1 mission. In *Proc. International Symposium on Artificial Intelligence Robotics and Automation in Space (i-SAIRAS)*, Nara (Japan), May 2003.
- D. J. Clancy. NASA challenges in autonomic computing. In *The Second Almaden Institute*, IBM Almaden Research Center, San Jose, CA (USA), 10–12 April 2002.
- D. Clark, C. Partridge, J. C. Ramming, and J. T. Wroclawski. A knowledge plane for the Internet. In Proc. ACM SIGCOMM 2003: Applications, Technologies, Architectures, and Protocols for Computer Communication, Karlsruhe, Germany, 2003.
- P. E. Clark, S. A. Curtis, and M. L. Rilee. ANTS: Applying a new paradigm to Lunar and planetary exploration. In *Proc. Solar System Remote Sensing* Symposium, Pittsburgh, PA (USA), 20–21 September 2002.
- S. Curtis, M. Rilee, W. Truszkowski, C. Cheung, and P. Clark. Neural basis function control of super micro autonomous reconfigurable technology (SMART) nano-systems. In Proc. First AIAA Intelligent Systems Technical Conference. AIAA, Chicago, IL, 20–22 September 2004.
- 31. S. A. Curtis, J. Mica, J. Nuth, G. Marr, M. L. Rilee, and M. K. Bhat. ANTS (autonomous nano-technology swarm): An artificial intelligence approach to Asteroid Belt resource exploration. In *Proc. Int'l Astronautical Federation*, 51st Congress, Rio de Janeiro (Brazil). AIAA, October 2000.
- S. A. Curtis, W. F. Truszkowski, M. L. Rilee, and P. E. Clark. ANTS for the human exploration and development of space. In *Proc. IEEE Aerospace Conference*, Big Sky, MT (USA), 9–16 March 2003.
- G. Deen, T. Lehman, and J. Kaufman. The almaden optimal grid project. In IEEE Autonomic Computing Workshop (5th AMS), Seattle, WA (USA), 14–21 June 2003.
- M. Dorigo and L. M. Gambardella. Ant colonies for the traveling salesman problem. *BioSystems*, 43:73–81, 1997.
- M. Dorigo and T. Stützle. Ant Colony Optimization. MIT Press, Cambridge, MA (USA), 2004.
- R. J. Doyle. Attention focusing and anomaly detection in systems monitoring. In Proc. Third International Symposium on Artificial Intelligence, Robotics, and Automation for Space (I-SAIRAS 94), pages 57–60, Pasadena, CA (USA), 18–20 April 1994. Jet Propulsion Laboratory.
- 37. A. S. Driesman, B. S. Ballard, D. E. Rodriguez, and S. J. Offenbacher. STEREO observatory trade studies and resulting architecture. In *Proc. IEEE Aerospace Conference*, volume 1, pages 63–80, Big Sky, MT (USA), March 2001.
- O. Etzioni. A Softbot-based interface to the Internet. Communications of the ACM, 37(7):72–76, 1994.
- R. J. Firby. The RAP language manual. Animate Agent Project Working Note AAP-6. Technical report, University of Chicago, Chicago, IL (USA), 1995.
- D. A. Fullford. Distributed interactive simulation: Its past, present, and future. In Proc. WSC '96: 28th Conference on Winter Simulation, pages 179–185, Washington, DC (USA), 1996. IEEE Computer Society.

- A. G. Ganek. Autonomic computing: implementing the vision. Keynote presentation, Autonomic Computing Workshop, AMS 2003, pages 2–3, Seattle, WA (USA), IEEE Computer Society, 25 June 2003.
- A. G. Ganek and T. A. Corbi. The dawning of the autonomic computing era. IBM Systems Journal, 42(1):5–18, 2003.
- E. Gat. ESL: A language for supporting robust plan execution in embedded autonomous agents. In *Proc. IEEE Aerospace Conference*, volume 1, pages 319–324, Aspen, CO (USA), 1–8 February, 1997.
- M. R. Genesereth. Knowledge interchange format. In Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference, pages 599–600, Cambridge, MA (USA), 1991. Morgan Kaufmann.
- M. R. Genesereth and S. P. Ketchpel. Software agents. Communications of the ACM, 37(7):48–53, 1994.
- J. C. Giarratano and G. D. Riley. Expert Systems: Principles and Programming. PWS Publishing Company, Boston, MA (USA), 4th edition, 15 October 2004.
- 47. M. D. Griffin and J. R. French, editors. *Space Vehicle Design*. American Institute of Aeronautics and Astronautics (AIAA), 2004.
- D. Ground and J. Schwab. Concept evaluation program of simulation networking (SIMNET). Technical Report 86-CEP345, Army and Engineering Board, Fort Knox, Kentucky (USA), 1988.
- R. Guitierrez and M. Huhns. Achieving software robustness via multiagentbased redundancy. In IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems, Acapulco (Mexico), 10 August 2003.
- 50. H. Guo. A Bayesian approach for autonomic algorithm selection. In IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems, Acapulco (Mexico), 10 August 2003.
- J. J. Guzman and A. Edery. Mission design for the MMS tetrahedron formation. In *Proc. IEEE Aerospace Conference*, volume 1, pages 540–545, Big Sky, MT (USA), March 2004.
- D. Harel. Comments made during presentation at "Formal Approaches to Complex Software Systems" panel session. In ISoLA-04 First International Conference on Leveraging Applications of Formal Methods, Paphos (Cyprus), 31 October 2004.
- L. Herger, K. Iwano, P. Pattnaik, J. J. Ritsko, and A. G. Davis. Special issue (J. J. Ritsko, Editor-in-Chief): Autonomic computing. *IBM Systems Journal*, 42(1):3–4, 2003.
- Hewlett-Packard Development Company, L.P. Adaptive infrastructure. HP World, 11–15 August 2003.
- D. E. Hiebeler. The swarm simulation system and individual-based modeling. In Proc. Decision Support 2001: Advanced Technology for Natural Resource Management, Toronto (Canada), September 1994.
- 56. M. Hinchey, J. Rash, and C. Rouff. Verification and validation of autonomous systems. In *Proc. SEW-26, 26th Annual NASA/IEEE Software Engineering Workshop*, pages 136–144, Greenbelt, MD (USA), November 2001. NASA Goddard Space Flight Center, Greenbelt, MD (USA), IEEE Computer Society Press.
- 57. M. G. Hinchey, J. L. Rash, and C. A. Rouff. Enabling requirements-based programming for highly dependable complex parallel and distributed systems.

In Proc. 1st International Workshop on Distributed, Parallel and Network Applications (DPNA 2005), Fukuoka (Japan), 20–22 July 2005. IEEE Computer Society Press.

- M. G. Hinchey, J. L. Rash, and C. A. Rouff. A formal approach to requirementsbased programming. In Proc. IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2005). IEEE Computer Society Press, 3–8 April 2005.
- 59. M. G. Hinchey, J. L. Rash, and C. A. Rouff. Towards an automated development methodology for dependable systems with application to sensor networks. In S. F. Andler and A. Cervin, editors, Proc. Real Time in Sweden 2005 (RTiS2005), The 8th Biennial SNART Conference on Real-time Systems (Reprinted from Proc. IEEE Workshop on Information Assurance in Wireless Sensor Networks (WSNIA 2005), Proc. International Performance Computing and Communications Conference (IPCCC-05), 2005), Skövde University Studies in Informatics, pages 73–79, University of Skövde (Sweden), 2005.
- T. Hoare and R. Milner. Grand challenges for computing research. Computer Journal, 48(1):49–52, 2005.
- W. M. L. Holcombe. Mathematical models of cell biochemistry. Technical Report CS-86-4, Sheffield University, UK, 1986.
- W. M. L. Holcombe. Towards a formal description of intracellular biochemical organization. Technical Report CS-86-1, Sheffield University, UK, 1986.
- 63. P. Horn. Autonomic computing: IBM's perspective on the state of information technology. White paper, IBM Research, Armonk, NY (USA), October 2001.
- 64. R. S. Hornstein, J. K. Willoughby, J. A. Gardner, R. Casasanta, J. Donald J. Hei, F. J. Hawkins, J. Eugene S. Burke, J. E. Todd, J. A. Bell, and R. E. Miller. Cost efficient operations: Challenges from NASA administrator and lessons learned from "hunting sacred cows." In *Fourth International Symposium on Space Mission Operations and Ground Data Systems*, Munich (Germany), 16–20 September 1996. American Aeronautical Society.
- P. Hughes, G. Shirah, and E. Luczak. Advancing satellite operations with intelligent graphical monitoring systems. In *Proc. AIAA Computing in Aerospace Conference*, San Diego, CA (USA), October 1993.
- P. M. Hughes. Application of autonomic computing concepts for GSFC's next generation missions. Presentation at the Woodrow Wilson International Center for Scholars, 28 October 2003.
- M. N. Huns, V. T. Holderfield, and R. L. Z. Gutierrez. Robust software via agent-based redundancy. In Proc. Second International Joint Conference on Autonomous Agents & Multiagent Systems (AAMAS 2003), pages 1018–1019, Melbourne, Victoria (Australia), 14–18 July 2003.
- 68. IBM. Autonomic computing concepts. White paper, 2001.
- IBM. An architectural blueprint for autonomic computing. White paper, October, 2003.
- IBM. Special issue on autonomic computing. IBM Systems Journal, 42(1), 2003.
- IBM and Cisco Systems. Adaptive Services Framework. White paper, version 1.0, IBM and Cisco Systems, 14 October 2003.
- IEEE International Conference on Autonomic Computing (ICAC'04), New York, NY (USA), 17–18 May 2004.
- T. Iida, J. N. Pelton, and E. Ashford. Satellite Communications in the 21st Century: Trends and Technologies. AIAA, 2003.

- J. C. Isaacs. Next Generation Space Telescope mission: Operations Concept Document (OCD). Technical Report STScI-NGST-OPS-0001C, Space Telescope Science Institute, Baltimore, MD (USA), 2001.
- N. R. Jennings. On agent-based software engineering. Artificial Intelligence, 117(2):277–296, 2000.
- 76. G. Kaiser, J. Parekh, P. Gross, and G. Valetto. Kinesthetics extreme: An external infrastructure for monitoring distributed legacy systems. In *The Autonomic Computing Workshop, 5th Int. Workshop on Active Middleware Services* (AMS'03), pages 22–30, Seattle, WA (USA), 25 June 2003.
- G. Karjoth. Access control with IBM Tivoli access manager. ACM Transactions on Information and System Security, 6(2):232–257, 2003.
- J. Kennedy and R. Eberhart. Particle swarm optimization. In *IEEE International Conference on Neural Networks*, pages 1942–1948, Perth (Australia), IEEE, 1995.
- J. O. Kephart and D. M. Chess. The vision of autonomic computing. *IEEE Computer*, 36(1):41–50, 2003.
- K. Kistler-Glendon. Beginning the autonomic journey a review and lessons learned from an autonomic computing readiness assessment at a major US telecom. In Proc. CHIACS2: Conference on the Human Impact and Application of Autonomic Computing Systems, Yorktown Heights, NY, 21 April 2004.
- Y. Labrou and T. Finin. A proposal for a new KQML specification. Technical Report TR CS-97-03, Computer Science and Electrical Engineering Department, University of Maryland Baltimore County, Baltimore, MD (USA), 1997.
- D. B. Lange and M. Oshima. Programming and Deploying Java Mobile Agents with Aglets. Addison-Wesley, Reading, MA, 1998.
- G. Langranchi, P. D. Peruta, A. Perrone, and D. Calvanese. Toward a new landscape of systems management in an autonomic computing environment. *IBM Systems Journal*, 42(1):38–44, 2003.
- W. J. Larson and J. R. Wertz, editors. Space Mission Analysis and Design. Microcosm and Kluwer Academic Publishers, Dordrecht, 1992.
- T. Lau, D. Oblinger, L. Bergman, V. Castelli, and C. Anderson. Learning procedures for autonomic computing. In *IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Sys*tems, Acapulco (Mexico), 10 August 2003.
- M. Lauriente, R. Durand, A. Vampola, H. C. Koons, and D. Gorney. An expert system for diagnosing anomalies of spacecraft. In *Proceedings of the Third International Symposium on Artificial Intelligence, Robotics, and Automation for Space (I-SAIRAS 94)*, pages 63–66, Pasadena, CA (USA), 18–20 April 1994. Jet Propulsion Laboratory.
- S. M. Lewandowski, D. J. V. Hook, G. C. O'Leary, J. W. Haines, and L. M. Rossey. SARA: Survivable autonomic response architecture. In *Proc. DARPA Information Survivability Conference and Exposition II*, volume 1, pages 77–88, Anaheim, CA (USA), IEEE Computer Society, June 2001.
- S. Lightstone. Towards benchmarking autonomic computing maturity. In Proc. IEEE Workshop on Autonomic Computing Principles and Architectures (AUCOPA 2003), pages 451–456, Banff, Alberta (Canada), 22–23 August 2003.
- M. Littman, T. Nguyen, and J. Hirsh. A model of cost-sensitive fault mediation. In *IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems*, Acapulco (Mexico), 10 August 2003.

- H. Liu and M. Parashar. Dios++: A framework for rule-based autonomic management of distributed scientific applications. In 9th International Euro-Par Conference (Euro-Par 2003), Lecture Notes in Computer Science, volume 2790, pages 66–73. Heidelberg (Germany), Springer, 2003.
- J. Llinas, C. Bowman, G. Rogova, A. Steinberg, E. Waltz, and F. White. Revisiting the JDL Data Fusion Model II. In *Proc. 7th International Data Fusion Conference*, Stockholm (Sweden), 28 June–1 July 2004.
- A. C. Long, J. O. Cappellari, Jr., C. E. Velez, and A. J. Fuchs, editors. Goddard Trajectory Determination System (GTDS) Mathematical Theory, Revision 1. NASA Goddard Space Flight Center, Greenbelt, MD (USA), 1989.
- F. Luna and B. Stefansson. Economic Simulations in Swarm: Agent-Based Modelling and Object Oriented Programming. Kluwer Academic Publishers, Dordrecht, 2000.
- L. Lymberopoulos, E. Lupu, and M. Sloman. An adaptive policy-based framework for network services management. *Journal of Network and Systems Man*agement, 11(3):277–303, 2003.
- 95. Microsoft Corporation. Dynamic Systems Initiative overview. White paper, 31 March, Revised 15 November, 2004.
- M. T. Morrow, C. A. Woolsey, and G. M. Hagerman, Jr. Exploring Titan with autonomous, buoyancy driven gliders. *Journal of the British Interplanetary Society*, 59(1):27–34, 2006.
- R. Muralidhar and M. Parashar. A distributed object infrastructure for interaction and steering. Concurrency and Computation: Practice and Experience, 15(10):957–977, 2003.
- 98. R. Murch. Autonomic Computing. IBM Press, Prentice-Hall, 2004.
- N. Muscettola, P. P. Nayak, B. Pell, and B. Williams. Remote agent: To boldly go where no AI system has gone before. *Artificial Intelligence*, 103(1/2):5–47, 1998.
- 100. P. P. Nayak, D. E. Bernard, G. Dorais, E. B. Gamble, Jr., B. Kanefsky, J. Kurien, W. Millar, N. Muscettola, K. Rajan, N. Rouquette, B. D. Smith, W. Taylor, and Y. wen Tung. Validating the DS1 remote agent experiment. In Proc. of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS-99), Noordwijk (The Netherlands), 1-3 June, 1999.
- 101. D. A. Norman, A. Ortony, and D. M. Russell. Affect and machine design: Lessons for the development of autonomous machines. *IBM Systems Journal*, 42(1):38–44, 2003.
- 102. J. Padget. The role of norms in autonomic organizations. In IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems, Acapulco (Mexico), 10 August 2003.
- 103. M. Papazoglou, S. Laufmann, and T. K. Sellis. An organizational framework for cooperating intelligent information systems. *International Journal of Intelligent and Collaborative Information Systems*, 1(1):169–202, 1992.
- 104. M. Parashar, editor. The Autonomic Computing Workshop, 5th Int. Workshop on Active Middleware Services (AMS 2003), Seattle, WA (USA), 25 June 2003. IEEE Computer Society.
- 105. D. Parkes. Five AI challenges in strategy proof computing. In IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems, Acapulco (Mexico), 10 August 2003.

- 106. R. Patil, R. Fikes, P. Patel-Schneider, D. McKay, T. Finin, T. Gruber, and R. Neches. The DARPA knowledge sharing effort: Progress report. In *Proc. KR'92, The Annual International Conference on Knowledge Acquisition*, pages 599–600, Cambridge, MA (USA), 1992.
- 107. L. D. Paulson. Computer system, heal thyself. *IEEE Computer*, 35(8):20–22, 2002.
- 108. B. Pell, D. E. Bernard, S. A. Chien, E. Gat, N. Muscettola, P. P. Nayak, M. D. Wagner, and B. C. Williams. An autonomous spacecraft agent prototype. Autonomous Robots, 5(1):29–52, 1998.
- J. P. Pickett, editor. American Heritage Dictionary of the English Language. Houghton Mifflin, Boston, 4th edition, 2005.
- J. Pitt and A. Mamdani. Some remarks on the semantics of FIPA's Agent Communication Language. Autonomous Agents and Multi-Agent Systems, 2(4): 333–356, 1999.
- 111. J. L. Rash, M. G. Hinchey, C. A. Rouff, and D. Gračanin. Formal requirementsbased programming for complex systems. In *Proc. International Conference* on Engineering of Complex Computer Systems, Shanghai (China), 16–20 June 2005. IEEE Computer Society Press.
- 112. J. L. Rash, M. G. Hinchey, C. A. Rouff, D. Gračanin, and J. D. Erickson. Experiences with a requirements-based programming approach to the development of a NASA autonomous ground control system. In Proc. IEEE Workshop on Engineering of Autonomic Systems (EASe 2005) Held at the IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2005). IEEE Computer Society Press, 3–8 April 2005.
- 113. J. L. Rash, C. A. Rouff, W. F. Truszkowski, D. Gordon, and M. G. Hinchey, editors. Formal Approaches to Agent-Based Systems, First International Workshop, FAABS 2000, Greenbelt, MD (USA), April, 2000, Revised Papers, Lecture Notes in Computer Science, Lecture Notes in Artificial Intelligence, volume 1871. Springer, 2001.
- 114. M. D. Rayman, P. Varghese, D. H. Lehman, and L. L. Livesay. Results from the Deep Space 1 technology validation mission. Acta Astronautica, 47(2–9): 475–487, 2000.
- C. W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. Computer Graphics, 21(4):25–34, 1987.
- C. K. Riesbeck and R. C. Schank. *Inside Case-Based Reasoning*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1989.
- 117. C. Rouff. A test agent for testing agents and their communities. In *Proc. IEEE Aerospace Conference*, Big Sky, MT (USA), March 2002.
- 118. C. Rouff and M. Hinchey. Modeling the LOGOS multi-agent system with CSP. In Proc. AAAI Spring Symposium, Technical Report SS-01-04, 2001.
- 119. C. Rouff, J. Rash, M. Hinchey, and W. Truszkowski. Formal methods at NASA Goddard Space Flight Center. In Agent Technology from a Formal Perspective, NASA Monographs in Systems and Software Engineering, pages 287–310. Springer, London (UK), 2005.
- 120. C. Rouff and W. Truszkowski. A process for introducing agent technology into 782 space missions. In *Proc. IEEE Aerospace Conference*, Volume 6, pages 2925–2935, Big Sky, MT (USA), IEEE, 8-15 March 2001. 783.
- 121. C. Rouff, A. Vanderbilt, M. Hinchey, W. Truszkowski, and J. Rash. Properties of a formal method for prediction of emergent behaviors in swarm-based

systems. In Proc. 2nd IEEE International Conference on Software Engineering and Formal Methods, Beijing (China), September 2004.

- 122. C. A. Rouff. Autonomy in future space missions. In Proc. IEEE Aerospace 788 Conference, Big Sky, MT (USA), March 2003. 789.
- 123. C. A. Rouff, M. G. Hinchey, J. L. Rash, W. F. Truszkowski, and D. Gordon-Spears, editors. Agent Technology from a Formal Perspective. NASA Monographs in Systems and Software Engineering. Springer, London (UK), 2006.
- 124. C. A. Rouff, J. L. Rash, and M. G. Hinchey. Experience using formal methods for specifying a multi-agent system. In *Proc. Sixth IEEE International Conference on Engineering of Complex Computer Systems (ICECCS 2000)*, Tokyo (Japan), 2000. IEEE Computer Society Press.
- 125. C. A. Rouff, W. F. Truszkowski, M. G. Hinchey, and J. L. Rash. Verification of NASA emergent systems. In Proc. 9th IEEE International Conference on Engineering of Complex Computer Systems, Florence (Italy), April 2004. IEEE Computer Society Press.
- 126. C. A. Rouff, W. F. Truszkowski, J. L. Rash, and M. G. Hinchey. A survey of formal methods for intelligent swarms. Technical Report TM-2005-212779, NASA Goddard Space Flight Center, Greenbelt, MD (USA), 2005.
- 127. L. Russell, S. Morgan, and E. Chron. Clockwork: A new movement in autonomic systems. *IBM Systems Journal*, 42(1):77–84, 2003.
- 128. L. Russell, S. Morgan, and E. Chron. On-line model selection procedures in Clockwork. In Proc. IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems, Acapulco (Mexico), 10 August 2003.
- 129. R. SAhoo, I. Rish, A. Oliner, M. Gupta, J. Moreira, S. Ma, R. Vilata, and A. Sivasubramaniam. Autonomic computing features for large-scale server management and control. In *IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems*, Acapulco (Mexico), 10 August 2003.
- M. Salehie and L. Tahvildari. Autonomic computing: emerging trends and open problems. SIGSOFT Software Engineering Notes, 30(4):1–7, 2005.
- M. Savage and M. Askenazi. Arborscapes: A swarm-based multi-agent ecological disturbance model. Working paper 98-06-056, Santa Fe Institute, Santa Fe, NM (USA), 1998.
- P. Scerri, D. Pynadath, and M. Tambe. Towards adjustable autonomy for the real-world. *Journal of AI Research*, 17:171–228, 2002.
- T. P. Schetter, M. E. Campbell, and D. M. Surka. Multiple agent-based autonomy for satellite constellations. *Artificial Intelligence*, 145(1–2):147–180, 2003.
- 134. U. M. Schwuttke, J. R. Veregge, and A. G. Quan. Performance results of cooperating expert systems in a distributed real-time monitoring system. In Proc. Third International Symposium on Artificial Intelligence, Robotics, and Automation for Space (I-SAIRAS 94), pages 79–83, Pasadena, CA (USA), 18–20 April 1994. Jet Propulsion Laboratory.
- 135. R. Sherwood, J. Wyatt, H. Hotz, A. Schlutsmeyer, and M. Sue. Lessons learned during implementation and early operations of the DS1 beacon monitor experiment. In Proc. Third International Symposium on Reducing the Cost of Ground Systems and Spacecraft Operations, Tainan (Taiwan), 1999.

- 136. H. A. Simon. *Models of Thought*, "Motivational and emotional controls of cognition," pages 29–38. Yale University Press, New Haven, CT (USA), 1979 (reprinted).
- 137. R. F. Simpson. Explorer platform. In AIAA Aerospace Sciences Meeting, Reno, NV (USA), 11–14 January 1988. AIAA.
- 138. R. L. Simpson. A computer model of case-based reasoning in problem solving: An investigation in the domain of dispute mediation. Technical Report GIT-ICS-85/18, Georgia Institute of Technology, School of Information and Computer Science, Atlanta, GA, 1985.
- A. Sloman. Review of: Rosalind Picard's Affective Computing. AI Magazine, pages 127–137, Springer, 1999.
- 140. A. Sloman and M. Croucher. Why robots will have emotions. In Proc. 7th International Joint Conference on Artificial Intelligence, pages 197–202, Vancouver (Canada), 1981.
- 141. M. Smirnov. Area: Autonomic communications. In Proc. Consultation Meeting on Future Emerging Technologies: "New Communication Paradigms for 2020," Brussels (Belgium), March 2004. European Union.
- 142. M. Smirnov and R. Popescu-Zeletin. Autonomic communication. In Proc. New Communication Paradigms for 2020, "Brainstorming" Meeting on Future Emerging Technologies, Brussels (Belgium), July 2003. European Union.
- 143. D. Smith. Efficient mission control for the 48-satellite Globalstar constellation. In Proc. Third International Symposium on Space Mission Operations and Ground Data Systems, pages 663–670, Greenbelt, MD (USA), 15–18 November 1994.
- 144. J. B. Smith. Collective Intelligence in Computer-Based Collaboration. CRC Press, Boca Raton, 1994.
- 145. R. Sterrit and M. G. Hinchey. SPAACE: Self-properties for an autonomous & autonomic computing environment. In Proc. The 2005 International Conference on Software Engineering Research and Practice (SERP'05), pages 9–15, Las Vegas, NV (USA), 27 June 2005. CSREA Press.
- 146. R. Sterritt. Towards autonomic computing: Effective event management. In Proc. 27th Annual IEEE/NASA Software Engineering Workshop (SEW), pages 40–47, Greenbelt, MD (USA), 3–5 December 2002. IEEE Computer Society Press.
- 147. R. Sterritt. Autonomic computing: The natural fusion of soft computing and hard computing. In *IEEE International Conference on Systems, Man and Cypernetics*, pages 4754–4759, Washington, DC (USA), 5–8 October 2003.
- 148. R. Sterritt. Pulse monitoring: Extending the health-check for the autonomic GRID. In *IEEE Workshop on Autonomic Computing Principles and Architectures (AUCOPA 2003)*, pages 433–440, Banff, Alberta (Canada), 22–23 August 2003.
- 149. R. Sterritt. xACT: Autonomic computing and telecommunications. BT Exact Research Fellowship report, British Telecommunications, 2003.
- R. Sterritt. Autonomic networks: Engineering the self-healing property. Engineering Applications of Artificial Intelligence, 17(7):727–739, 2004.
- 151. R. Sterritt. SelfWares-Episode IV Autonomic computing: A new hope. Presentation, IS&T Colloquium, NASA Goddard Space Flight Center, Greenbelt, MD (USA), 1 December 2004.

- 152. R. Sterritt and D. F. Bantz. PAC-MEN: Personal autonomic computing monitoring environments. In Proc. IEEE DEXA 2004 Workshops – 2nd International Workshop on Self-Adaptive and Autonomic Computing Systems (SAACS04), pages 737–741, Zaragoza (Spain), 30 August–3 September 2004. IEEE.
- 153. R. Sterritt and D. Bustard. Fusing hard and soft computing for fault management in telecommunications systems. *IEEE Transactions on Systems Man and Cybernetics – Part C: Applications and Reviews*, 32(2):92–98, 2002.
- 154. R. Sterritt and D. Bustard. Towards an autonomic computing environment. In 1st International Workshop on Autonomic Computing Systems, pages 694–698, Prague (Czech Republic), 1–5 September 2003.
- 155. R. Sterritt, D. Bustard, and A. McCrea. Autonomic computing correlation for fault management system evolution. In *Proc. IEEE International Conference on Industrial Informatics (INDIN 2003)*, pages 240–247, Banff, Alberta (Canada), 21–24 August 2003.
- 156. R. Sterritt and D. W. Bustard. Autonomic computing a means of achieving dependability? In Proc. IEEE International Conference on the Engineering of Computer Based Systems (ECBS-03), pages 247–251, Huntsville, AL (USA), April 2003. IEEE Computer Society Press.
- 157. R. Sterritt and S. Chung. Personal autonomic computing self-healing tool. In Proc. IEEE Workshop on the Engineering of Autonomic Systems (EASe 2004) at the 11th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2004), pages 513–520, Brno (Czech Republic), May 2004.
- 158. R. Sterritt, G. Garrity, E. Hanna, and P. O'Hagan. Survivable security systems through autonomicity. In *Proc. Workshop on Radical Agent Concepts (WRAC)* 2005, Lecture Notes in Computer Science, volume 3825, Greenbelt, MD (USA), September 2005. Springer.
- 159. R. Sterritt, D. Gunning, A. Meban, and P. Henning. Exploring autonomic options in an unified fault management architecture through reflex reactions via pulse monitoring. In Proc. IEEE Workshop on the Engineering of Autonomic Systems (EASe 2004) at the 11th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2004), pages 449–455, Brno (Czech Republic), May 2004.
- 160. R. Sterritt and M. Hinchey. Engineering ultimate self-protection in autonomic agents for space exploration missions. In Proc. EASe-2005, 2nd IEEE Workshop on Engineering Autonomic Systems, at ECBS 2005, 12th IEEE International Conference on Engineering of Computer Based Systems, pages 506–511, Greenbelt, MD (USA), 4–7 April 2005. IEEE Computer Society Press.
- 161. R. Sterritt and M. G. Hinchey. Apoptosis and self-destruct: A contribution to autonomic agents? In Proc. FAABS-III, 3rd NASA/IEEE Workshop on Formal Approaches to Agent-Based Systems, pages 269–278. Springer, April 2004.
- 162. R. Sterritt and M. G. Hinchey. Birds of a feather session: autonomic computing: Panacea or poppycock? In Proc. EASe 2005: IEEE Workshop on the Engineering of Autonomic Systems at 12th Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS 2005), Greenbelt, MD (USA), 3–8 April 2005.

- 163. D. J. T. Sumpter, G. B. Blanchard, and D. S. Broomhead. Ants and agents: A process algebra approach to modelling ant colony behaviour. *Bulletin of Mathematical Biology*, 63(5):951–980, 2001.
- 164. Sun Microsystems. N1 introducing just-in-time computing. White paper, 2002.
- 165. M. Swartwout. Engineering data summaries for space missions.
- 166. S. J. Talabac. Spacecraft constellations: The technological challenges in the new millennium. Technical report, AETD Information Systems Center, Code 588, NASA Goddard Space Flight Center, Greenbelt, MD (USA), 27 September 1999.
- 167. H. D. Tananbaum, N. E. White, J. A. Bookbinder, F. E. Marshall, and F. Cordova. Constellation X-ray mission: implementation concept and science overview. In O. H. Siegmund and K. A. Flanagan, editors, Proc. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, volume 3765, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pages 62–72, October 1999.
- 168. M. Therani, D. Zeng, and M. Dror. Decentralized resource management in autonomic systems. In *IJCAI Workshop on AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems*, Acapulco (Mexico), 10 August 2003.
- H. Tianfield. Multi-agent based autonomic architecture for network management. In Proc. IEEE International Conference on Industrial Informatics, pages 462–469, Banff, Alberta, Canada, August 2003.
- H. Tianfield and R. Unland (guest eds.). Special issue: Autonomic computing systems. Engineering Applications of Artificial Intelligence, 17(7):689–869, 2004.
- 171. R. L. Ticker and D. McLennan. NASA's New Millennium Space Technology 5 (ST5) project. In *IEEE Aerospace Conference*, volume 7, pages 609–617, Big Sky, MT (USA), March 2000.
- 172. C. N. Toomey, E. Simoudis, R. W. Johnson, and W. S. Mark. Software agents for the dissemination of remote terrestrial sensing data. In Proc. Third International Symposium on Artificial Intelligence, Robotics, and Automation for Space (I-SAIRAS 94), pages 19–22, Pasadena, CA (USA), 18–20 April 1994. Jet Propulsion Laboratory.
- 173. W. Trumler, F. Bagci, J. Petzold, and T. Ungerer. Smart doorplates toward an autonomic computing system. In *The Autonomic Computing Workshop*, 5th Int. Workshop on Active Middleware Services (AMS'03), pages 42–47, Seattle, WA (USA), 25 June 2003. IEEE Computer Society.
- 174. W. Trumler, J. Petzold, F. Bagci, and T. Ungerer. AMUN autonomic middleware for ubiquitous environments applied to the Smart Doorplate project. In *International Conference on Autonomic Computing (ICAC'04)*, pages 274-275, New York, NY (USA), 17–19 May 2004. IEEE Computer Society.
- 175. W. Truszkowski and H. Hallock. Agent technology from a NASA perspective. In Proc. CIA-99, Third International Workshop on Cooperative Information Agents, Uppsala (Sweden), 31 July – 2 August 1999. Springer.
- 176. W. Truszkowski, M. Hinchey, J. Rash, and C. Rouff. NASA's swarm missions: The challenge of building autonomous software. *IEEE IT Professional*, 6(5): 47–52, 2004.

- 177. W. Truszkowski and C. Rouff. An overview of the NASA LOGOS and ACT agent communities. In *Proc. World Multiconference on Systemics, Cybernetics and Informatics*, Orlando, FL (USA), July 2001.
- 178. W. Truszkowski and C. Rouff. Progressive autonomy. In *Proc. 2002 NASA/ IEEE Software Engineering Workshop*, Greenbelt, MD (USA), December 2002. IEEE Press.
- 179. W. Truszkowski, C. Rouff, S. Bailin, and M. Riley. Progressive autonomy: A method for gradually introducing autonomy into space missions. *Innovations in Systems and Software Engineering*, 1(2):89–99, 2005. Springer.
- 180. W. F. Truszkowski, L. Hallock, C. A. Rouff, J. Kerlin, J. L. Rash, M. G. Hinchey, and R. Sterritt. Autonomous and Autonomic Systems with Applications to NASA Intelligent Spacecraft Operations and Exploration Systems. NASA Monographs in Systems and Software Engineering. Springer, London (UK), 2007.
- 181. W. F. Truszkowski, M. G. Hinchey, J. L. Rash, and C. A. Rouff. Autonomous and autonomic systems: A paradigm for future space exploration missions. *IEEE Transactions on Systems Man and Cybernetics – Part C: Applications and Reviews*, 36(3), pages 279–291. May 2006.
- 182. W. F. Truszkowski, J. L. Rash, C. A. Rouff, and M. G. Hinchey. Asteroid exploration with autonomic systems. In Proc. 11th IEEE International Conference and Workshop on the Engineering of Computer-Based Systems (ECBS), Workshop on Engineering of Autonomic Systems (EASe), Brno (Czech Republic), May 2004. IEEE Computer Society.
- 183. W. F. Truszkowski, J. L. Rash, C. A. Rouff, and M. G. Hinchey. Some autonomic properties of two legacy multi-agent systems LOGOS and ACT. In Proc. 11th IEEE International Conference on Engineering Computer-Based Systems (ECBS), Workshop on Engineering Autonomic Systems (EASe), pages 490–498, Brno (Czech Republic), May 2004. IEEE Computer Society Press.
- 184. W. F. Truszkowski, C. A. Rouff, S. Bailin, and M. Rilee. Progressive autonomy – An incremental agent-based approach. In Proc. 2005 International Conference on Software Engineering Research and Practice (SERP'05), pages 9–15, Las Vegas, NV (USA), 27 June 2005. CSREA Press.
- 185. K. Tumer, A. Agogino, and D. Wolpert. Autonomous Agents & Multiagent Systems, Part 1, "Learning Sequences of Actions in Collectives of Autonomous Agents," pages 378–385. ACM Press, 2002.
- 186. C. Vanek. TDRSS, space-based communications for the present and for the future. In Proc. Space Programs and Technologies Conference and Exhibit, Huntsville, AL (USA), 21–23 September 1993.
- 187. R. Want, T. Pering, and D. Tennenhouse. Comparing autonomic and proactive computing. *IBM Systems Journal*, 42(1):129–135, 2003.
- 188. M. Weiser. Creating the invisible interface. In Proc. UIST '94: Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology, Invited talk, page 1, Marina del Rey, CA (USA), 1994. ACM Press.
- J. R. Wertz, editor. Spacecraft Attitude Determination and Control. Kluwer Academic Publishers, Dordrecht, 1978.
- 190. J. R. Wertz, J. L. Cloots, J. T. Collins, S. D. Dawson, G. Gurevich, B. K. Sato, and J. Hansen. Autonomous orbit control: Initial flight results from UoSAT-12. In 23rd Annual AAS Guidance and Control Conference, AAS 00-011, Breckenridge, CO (USA), February 2–6, 2000. American Aeronautical Society.

- J. R. Wertz and W. J. Larson, editors. *Reducing Space Mission Costs.* Microcosm and Kluwer Academic Publishers, Dordrecht, 1996.
- 192. F. White. A model for data fusion. In 1st National Symposium on Sensor Fusion, volume 2, Orlando, FL. 5–8 April 1988.
- B. C. Williams and P. P. Nayak. Immobile robots: Artificial intelligence in the new Millennium. AI Magazine, 17(3):16–35, 1996.
- 194. B. C. Williams and P. P. Nayak. A model-based approach to reactive selfconfiguring systems. In *Proc. AAAI-96*, pages 971–978, Portland, OR (USA), 4–8 August, AAAI Press, 1996.
- 195. J. Wyatt, R. Sherwood, M. Sue, and J. Szijjarto. Flight validation of ondemand operations: The Deep Space One beacon monitor operations experiment. In 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS '99), ESTEC, Noordwijk (The Netherlands), 1–3 June 1999.

Index

Ada, 59 adaptive algorithms, 174 operations, 183 reasoning, 70 scheduler, 135-137 user interface, 69 Advanced Research Projects Agency (ARPA), 170 agent, 17-19, 21, 35, 97, 108, 147, 208 actions, 108, 149, 150 adaptation, 184, 200 aglets, 171 applets, 171 architecture, 69-72, 74, 78, 81-83, 171assurance, 204 attributes, 17, 18 autonomy, 91 based control centers, 69 programming, 171 spacecraft, 115 behavior, 82 beliefs, 149 cloning, 72, 171, 202, 204, 205 collaboration, 76, 103 communication, 70, 74, 76, 83, 85, 108, 109, 172, 203 communication language, 69, 72, 76, 77, 82, 84, 103, 170, 172 communities, 78, 79, 88, 195, 202, 205, 230, 246

constellation, 198, 246 cooperation, 78, 235 development, 138, 203, 204 effector, 11, 22, 84–86, 177, 178, 182 efficiency, 200 executive, 119, 134, 237-239, 241, 242, 247 external interface, 78 fault tolerance, 204 formal verification, 91 Foundations of Intelligent Physical Agents (FIPA), 85, 172 framework, 86 goals, 71, 84, 149, 152, 154, 168 immobots, 17, 18, 23, 108, 145, 146 incubator, 204 informational, 19, 20, 146 learning, 83, 151 migration, 72, 90, 196, 202-205 mobile, 19, 171, 189 model, 70, 71 domain, 83 self, 151 world, 149, 150 monitoring, 200, 225 multi-, 69, 72–74, 90 satellite organization, 200 negotiation, 21, 128, 166, 167, 200, 246, 247 parallelization, 204 perception, 108, 109, 150 perceptor, 83, 85-87, 177 persistence, 72

personal assistant, 20, 21 policy, 75 proxy, 88-91, 190, 195, 196, 204 reactive, 81 reasoning, 83 remote, 35, 61, 115-128, 130-135, 137-144, 168, 235, 242, 243, 247 event-driven, 137 requirements, 110 robots, 18, 21, 22, 108 robust, 81 sensing, 108, 109 spacecraft operations, 72 testing. 204 trust, 195, 205 verification and validation, 204 Agent Concept Testbed (ACT), 69, 81-83, 87-90, 190, 192 Agent-based Flight Operations Associate (AFLOAT), 69-76, 78, 79, 90 archiving, 25, 66 artificial intelligence, 95, 180 assembly language, 55 asteroid belt, 3, 4, 38, 130, 142, 144, 209, 212, 214, 216, 224 automation, 3, 5, 6, 9, 10, 13, 14, 26, 28, 30, 32, 66, 69, 141, 157, 160-163, 166, 168, 174, 181, 191, 193, 224 definition, 9 autonomic, 4, 5, 9-14, 23, 47, 95, 173-183, 185, 186, 189, 194, 196-198, 207, 210, 216-220, 223-225, 228, 230 agents, 183 architecture, 183 Astrolabe, 182 AutoMate, 183 Cisco Adaptive Network Care, 184 communications, 184, 185 computing, 174, 176, 178, 179, 183, 185, 186 control loop, 177 definition, 11 dependability, 174 Distributed Interactive Object Substrate, 183 economics, 185 element, 177, 178, 181, 184, 185

environment, 173-178, 180-184 heart-beat monitor, 178 HP Adaptive Infrastructure, 184 Information Assurance, 12 initative, 179 Intel Proactive Computing, 184 Kinesthetics eXtreme, 182 learning, 176, 182 legacy systems, 182, 183 machine design, 180 maintainability, 174 manager, 177, 178 Microsoft Dynamic Systems Initiative. 184 Monitor, Analyze, Plan and Execute (MAPE), 177 monitoring and control, 181, 182 Nervous System (ANS), 173 Personal Computing, 185 policies, 176, 183 pulse monitor, 178 reliability, 174 safety, 174 security, 174 selfadaptation, 217 adjusting, 11, 176, 177, 217 awareness, 4, 11, 64, 144, 174, 176, 178, 217 chop, 176, 186, 217 configuring, 11, 174–176, 181–183, 209, 216 directed, 11, 12, 216 direction, 10, 11, 64 governing, 10-13, 173, 183, 196, 216 healing, 11, 175, 176, 179, 182, 183, 216managing, 10, 11, 13, 173–178, 182, 183, 185, 186, 190, 196 modifying, 221 monitoring, 11, 32, 176, 177, 217 motivating, 82 optimizing, 11, 175, 176, 179-181, 183, 186, 216 organizing, 184 properties, 217 protecting, 11, 175, 176, 183, 217 regulating, 11 selfishness, 185

situated, 11 training, 181 ware, 11, 13, 174, 177 x, 11, 176-178, 183, 185, 186 Situated and autonomic communications Program, 184 Smart Doorplates, 181 Sun N1, 184 survivability, 173, 174 swarm, 216 Tivoli management environment, 181 ubiquitous computing, 184 autonomy adjustable, 17 All Sky Monitor, 59 constellation, 195 cooperative, 113 debugging, 225 definition, 9 development, 225 flight, 38 mixed, 17 pointing control, 43 semi-autonomous, 196 testing, 218, 224 validation, 225 verification, 218, 225 Bayesian networks, 181, 182 reasoning, 103, 105, 106 statistics, 105 black holes, 189, 226 C++, 59C-Language Integrated Production System (CLIPS), 74, 76, 77

Campbell, Mark, 199 charge-coupled device, 22, 124 chi-squared, 122, 127 Clockwork, 180 collaboration, 69, 96, 128, 146, 147, 156, 157, 169, 170, 186, 195, 207, 211, 232languages, 103 planning, 159, 160, 165 collective intelligence, 199 command and data handling, 34, 48, 118, 123, 126

command sequence generation, 163, 165 planning, 13 commercial off the shelf, 28, 30, 32, 138 common information model, 181 **Communicating Sequential Processes** (CSP), 91 communications, 127, 128, 134, 140, 141, 146 beacon mode, 128, 133 compression, 237 ground to space, 242 management, 33, 34 protocols, 190 space to ground, 239, 240 uplink and downlink, 25 complex systems, 174, 175, 208, 210, 218constellation, 14, 46, 63-65, 81, 88, 132, 139, 144, 149, 164, 173, 189, 190, 192-197, 199, 205, 225, 227, 229, 230, 235, 242-247 advantages, 193 agents, 195, 198, 200, 246 autonomic, 196, 198 autonomy, 195 clusters, 149, 194, 200 communications, 193, 198, 200, 227, 238, 240, 244, 246, 247 community, 195 complex, 189 control, 198, 199 coordination, 201 costs, 191, 193 formation flying, 16, 46, 63, 64, 137, 146, 149, 151, 155, 164, 189, 193, 194, 197, 198, 226 governance, 189, 197 ground stations, 194 learning, 197, 198 low earth orbit, 194 management, 229 microsat, 227 military applications, 194 missions Cluster, 191 Constellation X, 189 ESCORT, 191

Experimental Spacecraft System (XSS), 191 Global Positioning System (GPS), 194 GlobalStar, 14, 15, 191, 194 Iridium, 14 Magnetospheric Multiscale, 191 Magnetotail Constellation, 191 Orbcomm, 191 Orbital Express, 191 ST5, 189 Total Internal Reflection Optical System, 191 organization, 200, 201 overview, 190 simple, 189 simulation, 195 status, 197 types. 191 Continuous Process Improvement, 244 control systems, 104 theory, 180cooperation, 70, 110, 111, 113, 119, 124, 129-132, 139, 144-152, 155, 158-165, 167, 169-172, 180, 192, 195, 224, 235, 242 actions, 155 hierarchical, 153, 154, 158, 160 mission planning, 160 model, 149, 155, 165 perception, 155, 156 planning, 152, 153 problem solving, 169, 170 science planning, 159 spacecraft, 165 technologies, 152, 169 virtual platform, 164, 166 coordination, 119, 151, 155, 160, 170, 191, 193, 200, 224, 245 architecture, 200 cost lifecycle, 53 savings, 15, 134, 135 crew, 3, 26, 37, 38, 115, 224, 229, 230, 235

DARPA Defense Advanced Research Projects Agency, 13 data analysis, 26, 35, 125, 148, 160, 163 archiving, 26, 66, 229 calibration, 26 capture, 25, 26, 28, 29, 66 collection, 165 compression, 44 distribution, 25 errors, 237 fusion, 109, 110, 155, 160, 163, 165 integration, 246 latency, 247 management, 28 mining, 193 monitoring, 26, 33, 34, 66, 119, 120, 122, 123, 127, 131, 132, 134, 140-143, 236 monitoring and trending, 144, 237, 239, 241, 244 processing, 25, 26, 35, 117, 119, 120, 124, 128-130, 133, 134, 140-144, 193, 229, 237-239, 241, 246 onboard, 65 science, 156, 158 reduction processing, 245 storage, 28, 33, 34, 44, 45, 66, 116, 119, 120, 124, 125, 134, 140, 141, 144, 238-240 trending, 119-123, 125, 127, 128, 131, 132, 134, 140-143 validation, 66, 67, 239 verificatoin, 239 deadlock, 134 decision support, 229 Deep Space Network, 239, 240 Defense Advanced Research Projects Agency (DARPA), 12, 113, 169, 170Autonomic Information Assurance, 12Knowledge Sharing Effort, 103, 170 distributed computing, 180 e-commerce, 193 earth

centroid, 62 observation, 194

pointers, 28, 42, 45, 137 science plan, 229 Eberhart, Russell, 215 Einstein, Albert, 66, 226 Embedded Computing, 185 emergence, 202, 208, 209, 211, 216, 218, 219definition, 208 verification, 219 environment aware, 175, 176, 229 event correlation, 181 evolutionary computing, 176 expert system, 16, 35, 196 exploration. 38 robotic, 3 Explorer Platform, 72 failure remediation, 182 fault correction, 26, 131-133 detection, 10, 120, 127, 128, 131detection and correction, 26, 29, 32-35, 47, 50, 56-58, 60, 61, 63, 116, 118, 119, 122, 123, 125, 126, 131, 134, 139-144, 195, 236-241, 244 constraints, 50diagnosis and correction, 26, 66 tolerance, 11, 78, 174 Finin, Tim, 76 flight autonomy capabilities, 54 computer, 28 data storage, 130 dynamics, 231 dynamics team, 232 hardware, 31, 32, 119, 135, 138 operating system, 125 operations team, 9, 10, 25, 29, 37, 39, 51-53, 58, 61, 116, 135, 232, 240 processor, 126 software, 5, 10, 19, 29-33, 35, 37, 40, 41, 43-45, 47-54, 57-63, 67, 115-118, 120-127, 130, 134-139, 141-143, 189, 231-233, 235-239, 244, 247 backbone, 115-119, 121-125, 138, 139, 204, 237-239, 241, 244, 247

bus, 226 design, 243 development, 138, 142 executive, 33, 34, 168 internal data transfer, 44 maintenance, 52, 137, 186 operating systems, 59 patch. 52 reuse, 138 safemode, 10, 21, 29, 32, 33, 35, 39, 44, 45, 47, 50, 53, 56-58, 62, 116, 118-123, 131, 132, 135, 143, 158, 161, 240 testing, 138 system, 25, 26, 33, 34, 41-45, 55, 132forecasting, 180 formal language, 103 methods, 91, 218, 219 verification, 91 Formal Approaches to Swarm Technology (FAST), 219 formation flying, 165 gamma ray, 59, 62 General Theory of Relativity, 66, 226genetic algorithms, 107, 181, 217 GenSAA/Genie, 80, 81 geocentric inertial, 231 Georgia Tech, 169, 172, 215 Global Positioning System (GPS), 33, 63, 121, 125, 126, 140, 194, 231, 240goals, 85, 95, 98, 100, 154, 173 directed, 70 science, 157, 162, 164

gravitational radiation, 65 grid computing, 184

groupware, 162

health and safety, 10, 16, 25, 29, 32, 34, 39, 40, 46, 49, 50, 52, 53, 80, 88, 89, 115, 116, 118–120, 122, 124, 138, 139, 143, 144, 157, 162, 189, 192, 197, 205, 241, 244, 247 health management, 145 IBM, 11, 171, 174, 177, 179, 181 AlphaWorks Autonomic Zone, 181 image detector, 109 multi-spectral, 108 processing, 109 information systems, 230 intelligence, 180 ambient, 184 collective, 197 control, 170 distributed, 200 economical, 190, 196 emotion, 180 machine design, 11 interferometry, 63, 64, 226, 227 laser, 226 UV-Optical Fizeau, 226 Internet, 171, 191, 216 invisible computing, 184, 185 Java, 171 Jet Propulsion Laboratory (JPL), 52, 61, 144, 168

Kalman filter, 56, 233 Kennedy, James, 215 knowledge acquisition, 146 base, 70, 79, 81, 84, 170, 178, 200 capture, 181 Interchange Format (KIF), 170 management, 196 Query Manipulation Language (KQML), 76, 170, 172 representation, 181 sharing, 243 Sharing Effort, 103, 170

Lagrange point, 49, 63, 121, 136, 142, 212, 226, 229, 235 celestial pointer, 139, 142 Laufmann, Steve, 71, 76 learning, 71, 72, 102, 106, 107, 181, 202, 243 adaptive, 169, 176 case-based, 169 deliberative, 169

offline, 169 reactive, 169 similarity-based, 169 verification, 219 legacy systems, 182, 183, 186 Lights Out Ground Operations System (LOGOS), 69, 78-81, 90, 91 logic Boolean, 61 fuzzy, 103, 104, 137, 168, 217 tree. 42magnetometer, 5, 61, 232 magnetosphere, 5, 7, 89, 191, 195 radio bursts, 226 Mars, 4, 18, 207 Mathematica, 105 middleware, 179, 182 mission concept, 31, 207, 209, 210, 212, 214, 225, 230 control, 3-5, 10, 17, 25-27, 29, 31-35, 37-39, 41-43, 45-51, 53, 56, 58, 59, 62, 67, 69, 72, 81, 89, 90, 115, 121, 123, 128-131, 133, 135, 137-139, 141-144, 146, 147, 159, 161, 163-165, 167, 189, 192, 194, 197, 202-205, 214, 220, 223-225, 227, 230-232, 236-242, 244, 245 autonomy, 203 lights out, 5, 72, 115, 124, 139–143, 238, 239, 244, 245 operations personnel, 29, 32 operator-to-spacecraft ratio, 15 costs, 15, 144, 146, 223, 232 design, 25, 198 future, 225, 228 heterogeneous, 15 homogeneous, 15 management, 156 manager, 163, 164, 168 negotiation, 161, 162 operations, 13, 69, 192, 223 risks, 146, 232 types deep space, 144 robotic, 5 survey, 139, 142, 239 Mission Operations Control Center, 90
Mission Operations Planning and Scheduling System, 80 missions Application Lunar Base Activities, 210Autonomous Nano Technology Swarm (ANTS), 3, 209, 210, 212, 214, 216-221 Big Bang Observer, 226 Black Hole Imager, 226 Cassini, 3 Cluster, 191 Compton Gamma Ray Observatory, 57.58 Constellation X, 189 Dawn, 4 deep space, 139 Deep Space One, 61, 144, 168, 169 Earth Observing Spacecraft (EOS), 62, 63 Earth Observing-1, 33 Enceladus, 227 Extreme Ultraviolet Explorer, 58, 78 Geospace Electrodynamic Connections, 8 Global Precipitation Measurement, 125, 232 High Energy Astronomical Observatory, 42, 55, 56 High Energy Astronomical Observatory-2, 42 Hubble Space Telescope, 27, 55–58, 63, 133, 224, 227, 247 International Ultraviolet Explorer, 49, 55, 141 James Webb Space Telescope, 49, 54, 63, 135–137, 143, 227, 228 Lander Amorphous Rover Antenna (LARA), 210 Landsat, 60, 63, 140 Laser Interferometer Space Antenna (LISA), 8, 46, 65, 225, 226 Magnetospheric Multiscale, 7, 191 Magnetotail Constellation, 8, 14, 191 Medium-Class Explorers, 51, 58, 226 Microwave Anisotropy Probe (MAP), 14Orbiting Astronomical Observatory-3, 55

Orbiting Solar Observatory-8, 56 Prospecting Asteroid Mission (PAM), 209, 210, 212, 216 Rossi X-ray Timing Explorer (RXTE), 54, 59, 60, 62, 63, 127, 232Saturn, 215 Titan, 215, 227 Saturn Autonomous Ring Array (SARA), 209, 210 Small Explorer, 30, 61, 121 Sojourner, 18 Solar Anomalous and Magnetospheric Particle Explorer, 60, 61 Solar Dynamics Observatory, 63 Solar Imaging Radio Array, 226, 227 Solar Maximum Mission, 55, 56, 58 Solar-Terrestrial Relations Observatory, 6 Space Technology 5, 5, 14, 189 Stellar Imager, 226, 227 Swift, 16, 60, 62, 63, 130 Burst Alert Telescope, 62 Techsat21, 200 Titan Explorer, 227 Total Internal Reflection Optical System, 191 Tracking and Data Relay Satellites (TDRS), 8, 51, 54, 60, 62, 64, 121, 195Demand Access System, 62 Triana, 62 Tropical Rainfall Mapping Mission, 60 Two Wide-angle Imaging Neutralatom Spectrometers, 8 Upper Atmosphere Research Satellite, 58Venus, 215 Willsinson Microwave Anisotropy Probe, 62, 63, 125, 232 mobile code, 224 computing, 185 software, 184 model checking, 219 modeling look-ahead, 120, 125, 127, 128, 140, 141, 144, 238, 239, 244, 245

statistical. 180 modelwebs, 229 monitoring and trending, 122, 129 moon, 27, 126, 210 outpost, 228 nanotechnology, 220 NASA, 3-5, 8, 13, 15, 37, 38, 55, 69, 85, 111, 125, 140, 145, 147-149, 162-166, 168, 169, 171, 173, 183, 186, 189, 190, 207, 209, 210, 212-215, 220 Goddard Space Flight Center, 38, 39, 42, 52, 55, 61, 63, 64, 116, 121, 128, 209 Institue for Advanced Concepts, 214, 215Langley Research Center, 209 New Millennium Program, 5, 61, 168 research and development, 37 Standard Spacecraft Computer, 55 strategic plan, 223, 228, 230 natural language, 75, 77 Navy, 220 networking, 19, 113, 180, 185, 191, 193, 197 ambient, 184 wireless, 184 neural networks, 106, 119, 122-124, 127, 168, 185, 215, 217 Nomadic Computing, 185

object-oriented design, 35, 138 development, 217 on-orbit sparing, 193 operations concept, 134 operations research, 180 optimization, 38, 148 combinatorial, 216 network, 216 query, 179 orbit, 26, 231 control, 231 correction, 242 determination, 25, 28, 33, 120, 121, 123, 125, 126, 129, 130, 139-141, 143, 231, 232, 237, 241 generators, 78

geostationary, 49, 229 geosynchronous, 28, 49, 122, 138, 215, 243 - 246celestial pointer, 138, 141 earth pointer, 139, 141 interpolator, 231 low earth, 27, 28, 41-43, 46, 136, 138-141, 143, 190, 194, 229, 231, 240, 244-246 celestial pointer, 138-140 earth pointer, 138, 140, 141, 240 sun-synchronized, 243 maintenance, 123 maneuver, 33, 120, 122, 129, 130, 134, 140, 141, 143, 157, 200, 232, 238, 241planning, 33, 46 medium earth, 229 perturbations, 232 prediction, 34, 231 propagation, 56, 63, 231 stationkeeping, 19, 28, 46, 122, 136, 140, 141, 232, 240, 241 time processor, 58

persistence, 72 Personal Computing, 185 pervasive computing, 184 planner, 85, 99 architecture, 96 case-based, 101, 102 decision-theoretic, 70 model-based, 100-102numeric, 99 symbolic, 98-100, 102 planning, 13, 95–98, 140, 146, 149, 152-154, 160, 165, 166, 200 command sequence, 78context, 100cycle, 162 deliberative, 97, 154 execution, 84 globally optimal, 97 hierarchy, 161 high level, 168 iteration time, 162 mission, 13, 53, 156-163, 165, 168 reaction time, 161 reactive, 97, 99, 100, 154, 168

replanning, 42, 97, 158 science, 156, 158-160, 162, 163, 165, 166 sequence, 156, 158, 160 sub-optimal, 97, 101 target, 128 time, 161 planning and scheduling, 25–28, 35, 42, 46, 63, 66, 70, 80, 83, 84, 88-90, 115, 119-123, 125, 127-130, 139-144, 163, 168, 232, 236-238, 241, 242, 244-247 architectures, 168 distributed, 195 science, 66, 123 power fuel cell, 228 management, 33 nuclear, 228 solar, 228 process algegra, 219 progressive autonomy, 90, 138, 190, 202, 203, 205 protocol agent, 75, 85, 138, 186 efficiency, 113 Internet, 171 space, 28, 190 publish-and-subscribe, 83, 86, 182 race conditions, 218, 219 radar. 109 Rapid Spacecraft Development Office, 64 reasoning, 81, 196 case-based, 29, 101, 102, 119, 122-124, 127, 130, 169 expert system, 88 heuristic, 214 model-based, 168 partial information, 103 probabilistic, 181 reflection, 11, 180 reflexive, 11, 180, 196, 199 routine, 11, 180 rule development, 181 engine, 183 rule-base, 70, 71, 88, 185

reconfiguration, 230 **Recovery Oriented Computing**, 179 redundancy, 47, 61, 183, 191, 193, 207 reengineering, 25, 26 reliability, 193, 224 remote sensing, 229 resource constraints. 189 limited. 204 management, 39, 44, 49, 50, 128, 134, 137 robots, 3, 17, 18, 21, 22, 26, 100, 112, 113, 115, 145, 146, 149, 154, 169, 172, 173, 205, 224, 228, 230 actuator, 22, 108 cooperative, 18 exploration, 38 immobile, 23 mobile, 169 navigation, 22, 145 reactive control, 22 sensors, 22 Sojourner, 18 space-based, 18 underwater, 169 robustness, 183 root-cause analysis, 181 Saturn, 209, 227 scheduling, 103, 125, 127-133, 135, 139, 140, 162, 227 absolute-time, 135, 136, 139 adaptive, 132, 135-137 calibration, 128, 131 dynamic, 126, 128 event-based, 135–137, 139 generation, 41 goal-driven, 130, 131 ground, 41, 129 onboard, 49, 140 opportunistic, 131 predictive, 161 rescheduling, 42 science, 26, 49, 129 short-term, 137 slective target, 42 spacecraft, 41 target, 128

Schetter, Thomas, 199

science coordinated, 191, 194 data, 43, 59, 148 data processing, 31, 33, 34, 115 execution, 39, 41, 132 goals, 156 negotiation, 156 observation, 50, 129, 132, 241 opportunistic, 15, 27, 38, 54, 56, 59, 60, 62, 125–131, 137, 139, 140, 143, 162, 205, 224, 236, 238 optimization, 38 planning, 13 schedule, 28 support activities, 26, 28 targets, 27, 57 semantics declarative, 170 grammar, 75 sensor calibration, 33 distributed, 183 network, 18, 23 sun, 232 web, 229, 230 sigma-editing, 122, 127 signal processing, 109 simulation, 110-112, 122 distributed interactive, 113, 171 environments, 113 forces, 171hardware, 111 networked, 113 onboard, 243 servers, 113 SimNet, 171 vehicle, 171 software cost, 51, 54, 143 development, 180 engineering, 174, 176, 186 hot swapping, 179 solar array, 21, 45, 227 flare, 56 magnetic activity, 226 position, 237 sail, 212, 217, 220 wind, 5, 7

sonar, 109 South Atlantic Anomaly, 27, 43, 58, 60, 121, 123, 241, 242 space physics, 191 space-based processing, 189 spacecraft anomaly, 32, 35, 39, 42, 44, 52, 53, 62, 79-81, 89, 117, 119, 122, 124, 128, 131, 134–137, 142, 158, 192, 195, 204, 205, 239, 244, 247 antenna, 54, 59, 60 aperature, 198 attitude actuator. 233 control, 21, 25, 28, 30, 32, 35, 43, 50, 53, 56, 58, 61, 117, 120, 121, 125, 133, 134, 140, 141, 143, 228, 231, 232, 237, 238 determination, 10, 28, 33, 56, 120, 125, 129, 130, 140, 141, 143, 231, 232error, 232 maneuvering, 129, 130, 132 measurement, 231 sensor, 141, 232, 233 autonomy, 38, 203 battery, 45, 53 bus, 33-35 calibration, 25, 26, 30, 56, 66, 117, 119-121, 125-129, 131, 132, 139, 140.142celestial pointer, 28, 42, 45, 50, 117 commands, 41 absolute time, 41, 42, 49, 56, 130, 137conditional, 41-43, 56 configuration, 141 delta-time, 42 execution, 44, 48, 49 loading, 26, 28, 66 management, 34 processing, 33, 34 real-time, 43, 137 relative-timed, 41, 42, 56, 58 scripts, 238 sequencing, 158, 165 timeline, 42 timing, 247 validation, 48, 241

verification, 160 communications, 21, 227 computer, 31, 33, 46, 55, 59, 62, 116, 125, 231, 247 DF224.57 NSSC-I, 55, 57 design, 3, 64, 145 engineering support, 26, 28 ephemeris, 19, 20, 27, 51, 53, 57, 60, 65, 121, 123 guidance, 57, 58, 136 Guidance, Navigation, and Control, 232guide star, 56–58, 63, 136 gyro drift, 53, 56, 121, 232 hardware, 41 inertial fixed pointing, 116, 117 hold, 35, 58, 117, 120, 121, 135 instrument, 28, 31 adjustment, 241 calibration, 25, 28, 30, 53, 61, 66, 117, 120, 121, 129, 133, 140, 141, 143, 236, 237 commanding, 33, 34, 116, 120, 124, 125, 128-130, 139, 140, 144, 237, 241, 242 communications, 241, 245 configuration, 42, 43, 62, 65, 120, 124, 125, 137, 139, 140, 144, 237, 241data storage, 241, 245 diagnostic, 119 Narrow Field, 62 performance, 239 pointing control, 42 reconfiguration, 241, 242 safemode, 35 smart, 132 survey, 63 intelligence, 199 launch, 193 mass storage, 126 microsatellites, 5, 227 mirrorsats, 226 momentum angular, 45, 46, 49, 136 dump, 28, 56, 63, 136

management, 33, 45, 57, 117, 145, 235, 236 monitoring, 69 multi-, 14, 15 nano-class, 8, 88, 89, 190, 191, 198 operating system, 126 operations, 13, 20, 21, 26, 37, 38, 40, 42, 46, 47, 51, 52, 134, 168, 199 Optical Telescope Assembly, 28, 34, 133orientation, 231 performance monitoring, 26, 29, 66 pico-class, 207, 209, 210, 212, 215, 220pointing control, 43, 44, 56, 57 power electrical, 45, 134 management, 21, 34, 35, 45, 56, 116, 118, 145 processing, 34, 67, 241 propulsion, 21, 117 management, 33, 45, 46 propellent, 134 subsystem, 33 reaction wheel, 45, 233, 237 safemode, 193, 239 schedule, 41 simulations, 132 spin-stabilized, 35, 58 state, 52 storage, 34, 60, 231 subsystems, 205 support activities, 66 support functions, 27 systems access, 51 telemetry, 9, 29, 30, 32, 34, 40, 44, 45, 50-53, 55, 57, 58, 60, 62, 72. 80, 81, 90, 134, 143, 144, 156, 158, 160, 192, 195, 239, 245 filter table, 53 formats, 53 monitor, 52, 57, 58, 60, 192 thermal management, 21, 33-35, 50, 53, 116, 118 thruster, 33, 118 control, 115, 122 firing, 116 management, 117, 134, 232 optimization, 56

uplink-downlink card, 35 standard deviation, 122, 127 star catalog, 232 star tracker, 10, 56, 60, 120, 125, 126, 232, 233 lost in space, 62, 232quaternion, 61, 237state modeling, 30, 119, 122–124, 127 state-based systems, 196 state-space, 219 stochastic methods, 185 strategy-schema, 71 string of pearls, 189 Sumpter, David J.T., 219 sun acquisition, 35 centroid, 63 coronal mass ejections, 227 point, 56 radio bursts, 226 sunpoint, 58 Super Miniaturized Addressable Reconfigurable Technology, 210 Surka, Derek, 199 swarm, 14, 173, 198, 207-210, 212, 214-221 Ant Colony Optimization, 215 applications, 208 autonomic, 216 birds, 215boids, 215 definition, 208 insect behavior, 215 intelligence, 208, 209 optimization, 208, 215 particle, 215 quantitative structure activity relationships (QSAR), 215 robotics, 209 simulation, 208 Super Miniaturized Addressable Reconfigurable Technology (SMART), 209 team, 212, 216, 217 tetrahedral walker, 207, 209-211 verification, 218, 219 systems engineering, 176

management, 180 of systems, 174, 182 target acquisition, 32, 35, 43, 58, 59, 125, 126, 132, 143, 237, 241 identification, 43, 229, 230 observation, 43quaternion, 59, 62 targeting, 130 task management, 33, 34 tasks, 98 TCP/IP, 73, 78 terrestrial databases, 193 testing, 95, 110, 111, 113, 218 environments, 111, 112 plan, 110 simulation, 110, 111 software, 171 time, 111 time management, 33, 34 total cost of ownership, 174, 186 tracking station, 192 trajectory planning, 200 traveling salesman problem, 216 trend analysis, 193

UK Computer Science Grand Research Challenges, 174
University of California Berkeley, 216
University of Maine, 169
Unmanned Aerial Vehicle, 207, 215
Unmanned Underwater Vehicles, 216
user modeling, 69
utility computing, 184

validation, 48, 132, 202–204, 225, 238, 243 verification, 90, 202–204, 218, 220, 224, 225 command, 160 emergence, 219 formal methods, 218, 219 learning, 219 swarm, 219 Virginia Tech, 214 virtual environment, 171 lens, 194 platform, 147, 149, 164–167, 189 presence, 168 telescopes, 191 world, 17, 19, 171 Visual Analysis Graphical Environment (VisAGE), 80 Weighted Synchronous Calculus of Communicating Systems, 219 Weiser, Mark, 184 Workplace, 85 world computing, 184

X-Machines, 219

Zadeh, Lotfi, 104