NASA Planning and Scheduling Applications: Emerging Technologies and Mission Trends

Mark S. Boddy[1] Bonnie H. Bennett Brian A. Isle Rainer A. Isle
Adventium Labs, Minneapolis

Work performed under NASA Grant NAG-2-1631

FINAL REPORT

May 28, 2004

[1]Corresponding author: mark.boddy@adventiumlabs.org

Contents

Chapter 1

Introduction

Planning and scheduling is an important capability for a wide variety of NASA missions, serving either as an enabler without which the mission cannot be done at all, or serving to improve mission operations through making those operations more robust, more predictable, or more efficient.

The objective of this study was to provide input to strategic planning for planning and scheduling research and development within NASA. The rationale for employing an outside agent in this work was precisely to get that outside perspective, with no agenda or axe to grind regarding research priorities. The output of the study is the current report.

This report is intended to serve several purposes:

- To provide a taxonomy of planning and scheduling functions relevant to NASA missions.
- To provide an analysis of NASA missions currently planned or under study, with regard to the needs of those missions for planning and scheduling, in terms of that taxonomy.
- To identify current or planned research and development within NASA or funded by NASA, addressing those identified needs.
- To identify remaining unmet needs.

A corollary result of having done the analysis described here has been the construction of the methodology used in that analysis. As the set of missions changes, and as current problems become future solutions, leading to more problems to solve, this approach can be re-applied.

The aim of this study was to look at the current portfolio of research in planning and scheduling with an eye to how well it meets NASA's anticipated needs over the next 5-10 years. Specifically, we examined research being conducted by NASA, in particular at Ames and JPL, or funded by NASA under the Intelligent Systems Program.[1]

The rationale for this approach is based on the following points:

- NASA-internal users of IT technologies, including planning and scheduling, think that CICT should be focussing on stable, longer-term technology programs [59].

[1] http://is.arc.nasa.gov/

- The same report makes the assertion that some of NASA's technology needs are either sufficiently specialized or non-commercial due to a small user base that this technology development must be done by or funded by NASA. Hence the focus on NASA internal and NASA-funded programs.
- Planning and scheduling is an area defined more by methods and models than by function. As demonstrated in this report, there is no functional mapping from planning technologies to functional needs: many of the technologies discussed here have potential application in multiple areas.

1.1 Planning and Scheduling Defined

"Planning and Scheduling" is a broad area that bleeds at the edges into several other areas. To make matters worse, the terms are both invidually and collectively used in different ways by different communities. The definitions provided here are closely aligned with the commonly-understood use of these terms with the community associated with conferences such as the "International Conferences on AI Planning and Scheduling" (ICAPS), and are consistent as well with usage within the AI and Robotics research community, including research groups at NASA Ames Reseearch Center (ARC) and Caltech's Jet Propulsion Laboratory (JPL).

Planning describes the process of deciding what to do. Generally, this is taken to mean a process of constructing a sequence or a partially-ordered network of actions to be taken to achieve some goal. There are strong analogies to certain kinds of control systems, specifically model-predictive control, in that the reasoning being done can be described as defining a "trajectory," then using information encoded in a projective model to figure out how to use available perturbations of the system state so as to follow that trajectory. On the other hand, most planning problems are more strongly discrete (the decision is whether to do A or B, not choosing a setpoint for a continuous value). Examples of planning problems relevant to NASA missions include generating a sequence of traversals and sampling operations on a specified set of rocks for the Mars Exploration Rovers (MER), or determining which analysis operations in what order performed on an image will result in the required data product.

Scheduling is the process of figuring out what resources will be used when to execute a set of tasks (which may, for example, be steps in a plan as above). For examnple, scheduling observations using a telescope imposes an ordering on those observations, because the telescope cannot point in two directions at the same time. These orderings must in addition be consistent with other constraints, such as when the observations's target is visible, or keeping the telescope from pointing too close to excessively bright objects.

Most real-world problems involve both planning and scheduling, in some combination. It is rare that the tasks to be done are provided in a completely-specified form, ready to be assigned to resources. It is at least as rare to be able to generate a plan without worrying about the availability of equipment or other resources needed to execute the steps in that plan.

By the nature of the problem being addressed, it is impossible to say that planning and scheduling systems are "solving" a problem that was previously unaddressed. Any operational mission is in some way deciding on what to do, and what resources to use to do those things decided upon. The available methods range from grease pencil on butcher paper, to tools such as spreadsheets or Microsoft Project, to early-generation schedulers such as OMP-26, all the way to more modern planning and scheduling systems, such as Aspen or Europa.

1.2 Relating Autonomy to Planning and Scheduling

In the NASA Information Technology Assessment Study [59] autonomy as an area is both too broad and too narrow for the purposes of this study. Too broad, in that planning is only one technology required for autonomy. Only in combination with other technologies (for example, system diagnosis) and only within the context of a particular "autonomy architecture" can planning be said to provide a solution for autonomy.

On the other hand, autonomy is too narrow as an area in which to discuss planning and scheduling, in that there are many applications where planning and scheduling may be of great value, which involve no autonomous operations at all. For example, generating an observation schedule for Hubble, or defining a sequence of traversals and experiments for MER, are significantly difficult planning and scheduling problems, whether or not the operations, once planned, are carried out autonomously.

1.3 Areas Excluded

There are some areas that will not be discussed in this report. They include:

- Broad surveys of the current state of the art in Planning and Scheduling in the research community as a whole.
- Autonomy architectures (except as they bear on requirements for planning and scheduling).
- Software engineering for planning and scheduling systems.
- Infusion of Information Technology in general or planning and scheduling in particular.

These issues are important, and are in various ways being addressed, but fall outside the scope of this report.

1.4 Information Technology Assessment Study

In 2001 and 2002, NASA conducted an internal "Information Technology Assessment Study" (ITAS) [59, 60]. This study was a broader look at IT investment for the Office of Space Science, not focussed on a particular technical area as is the current report. The conclusions in that study regarding how research is done, how the results of that research can best find their way into operational use, and where the research is best conducted, are relevant to planning and scheduling as a sub-field of their overall area of investigation.

Feedback from multiple sources in addition to the ITAS report (e.g., OMB, NRC PRT report) strongly support the notion that NASA should maintain internal capabilities in critical technologies, especially where those technologies are not yet available in commmercial products. Planning and scheduliing is one such area, and one which the NRC Interim report evaluates current NASA capabilities as being "world class" [1].

In the ITAS study, planning and scheduling is described as part of the autonomy area. Clearly, planning and scheduling is a required function for some types of autonomous behavior. What kind of planning and scheduling is required depends on the behavior required, and upon the architecture in which the planning system is operating (see [18] for a brief summary of NASA-relevant autonomy architectures). However, there are numerous uses for planning and scheduling in applications other than autonomy, for example in the MAPGEN planning tool currently being used to do activity planning for the Mars Exploration Rovers.

1.5 Roadmap for the Report

Chapter 2 presents the criteria according to which the missions were evaluated, and summarizes our findings. More detailed mission analysis data can be found in Appendix A. Chapter 3 maps from the mission criteria to a set of technical capabilities which can be applied to current research to evaluate the relevance of that research to the mission needs identified. Chapter 4 presents the summary analysis of coverage for these requirements, against projects currently active in NASA, or being funded by the Intelligent Systems Program. Detailed information on the projects discussed can be found in Appendix B. Chapter 5 provides a brief discussion of areas identified for further work, given mission trends and current research coverage as discussed previously. Finally, Chapter 6 sums up the report.

Chapter 2

Mission Summaries

We start with 67 missions, and rank them on 11 criteria. The criteria we used and the intended interpretation of the different rankings assigned are discussed, below.

Once collected, these data can be analyzed to identify missions with a large number of critical planning and scheduling challenges. They can identify planning and scheduling challenges which are critical across a variety of missions, and they can be grouped to identify classes of missions that share a common cohort of challenges for planning and scheduling.

According to this analysis, 22 of the missions originally identified for study do not involve difficult planning problems, meaning that they did not achieve a ranking of 9 on any criterion. On the other hand, several of these "easy" missions had rankings of at least 3 on more than half of the set of 11 criteria. A complete mission list with descriptions and rankings, as well as a description of missions removed prior to analysis and why, can be found in Appendix A.

The 45 remaining missions had significant difficulties, denoted by achieving a ranking of "9" in from 1 to 6 different areas.

The criteria we used are intended to span the range of issues that affect both the nature and difficulty of the planning to be solved. These criteria can be grouped as follows:

- Big issues, with no obvious current technical fix:
  - Autonomy
  - Coordinated multi-platform operations
  - Complex system dynamics
  - Uncertain execution
  - Distributed problem solving
- Significant issues, technical fixes may be known but must be implemented (the difficulty may be in integration with planning):
  - Complex, resource-bounded schedules
  - Over-subscribed schedules
  - Complicated, multi-step operations

- These issues do not in-and-of themselves mean you have a hard problem, but can certainly add to the difficulty:
  - Real-time constraints on response
  - Model drift
  - Costs on replanning

More details on the individual criteria and how we ranked them are given below.

2.1 Survey Criteria

Here are the criteria used in evaluating missions, with rankings attached and explained. In general, we are interested in identifying mission applications with features that may result in hard planning problems. The features on which we ranked the missions include:

- Autonomous operations (e.g., Europa cryobot, Venus landers, Mars rovers like MER)
- Crowded, resource-bounded schedules (e.g. Shuttle ops and refurbishment, Station ops)
- Over-subscribed schedules with many stakeholders (e.g., Hubble, SIRTF, SOFIA)
- Coordinated multi-platform operations (e.g. satellite constellations for earth observation, space-based interferometry)
- Complex system dynamics: orbital mechanics, resource limits (fuel, battery, reaction mass), biological systems (long-term lifesupport, e.g. for manned Mars exploration).
- Unpredictability (complex environments, information-gathering actions).
- Distributed problem solving (within the mission). We are interested in distribution as a requirement (driven by political separation of authority, real-time constraints, or limits on comm. bandwidth, e.g.) as opposed to distribution as a solution strategy.
- Real-time constraints on response. This is not the light-speed lag issue, which is addressed either as a need for automous operation, or a requirement for distributed problem solving, depending on other mission parameters. This is limits on local processing power compared to what's needed (so, driven either by limited processing hardware, or a dynamic environment).
- Model drift (equipment wear, seasonal variations, etc) due to extended mission lifetime and/or a harsh environment (e.g. JIMO or Galileo, enduring the near-Jupiter radiation environment over a period of years).
- Costs on replanning. This shows up in missions where 1) the plan is for an extended period, 2) the plan is published to or otherwise synchronized with others, and 3) the environment is unpredictable enough that replanning happens often compared within a reasonable planning horizon.

How missions were evaluated with respect to each of these criteria is shown below:

- Autonomy

  1 : Manned, or near-earth and single-platform mission.

  - **3** : Light-speed lag measured in minutes. Generally, lag small relative to required response times.
  - **9** : Little or no downlink capability. Communication lag of hours, or response times much shorter than lag.

- Complex, resource-bounded schedules
  - **1** : Single execution thread, or not resource-limited.
  - **3** : Multiple execution threads (housekeeping vs. science, e.g.), significant limits on power, etc.
  - **9** : Many simultaneous ops. Many resources, or many subscribers to a few resources.
- Over-subscribed schedules
  - **1** : Fixed set of tasks
  - **3** : Mild contention (low-priority science observations may lose out)
  - **9** : Severe contention: 5X-10X oversubsciption, many different scientists/institutions involved.
- Coordinated multi-platform operations
  - **1** : Single vehicle
  - **3** : Maybe a lander and an orbiter, loosely coupled (comm. relay,e.g.)
  - **9** : Multi-platform observations (e.g. interferometry), coordinating multiple orbits for joint observation.
- Complex system dynamics
  - **1** : Fixed or predetermined orbit, simple power (fixed power, or solar cells in orbit), Simple (or very noisy) model for locomotion.
  - **3** : Interplanetary orbital manuevers as part of planning (e.g., ion thrusters). Nonlinear power response (solar cells on a planetary surface). Orbital constraints on telecope pointing and status.
  - **9** : Airborne telescopes, coordinated orbital constellations, biological processes (long-term life-support).
- Complicated, multi-step operations
  - **1** : Simple physical plant. Few instruments, pre-programmed operations (e.g., Pioneer Venus Lander).
  - **3** : Limited number of setup steps, e.g. telecope pointing and setup for observations, mode reconfiguration for communication.
  - **9** : Multi-system interaction (multi-sensor observation, dumped into a complex datastore for later playback). Path-planning for multiple objectives. Big, complex deep space orbiters.
- Uncertain execution
  - **1** : Orbital mechanics
  - **3** : Conditional execution (quick sample leads to more detailed analysis, or not). Human execution of complex operations.
  - **9** : Driving over variable terrain, balloon travel, ocean navigation, programmed failure (Venus lander).

- Distributed problem solving
  - **1** : Single platform, single controlling authority (not manned, no communications relay through other agencies).
  - **3** : Multiple centers of authority (science vs. vehicle health), coordinated multi-vehicle operations ([3], preferably [9], on that metric), possibly exacerbated by real-time requirements.
  - **9** : Multi-national partnerships, multi-center collaborations, manned space flight.
- Real-time constraints on response
  - **1** : No signifcant interactions with local environment (e.g., satellite in fixed orbit).
  - **3** : Navigating unknown terrain. Flight. Orbital insertion
  - **9** : Control for multi-platform interferometry. Aerobraking and landing. Probes.
- Model drift
  - **1** : Short duration mission (days or weeks)
  - **3** : Extended duration (months), or greater wear (rovers in sand).
  - **9** : Extreme duration (a decade or more) or moderate wear and extended duration.
- Costs on replanning
  - **1** : One stakeholder (or one dictator), actions interchangeable in order (e.g., telescope observations in a forgiving environment), or short-term plans due to unpredictability.
  - **3** : Multiple stakeholders (Hubble), or actions interact in complicated ways.
  - **9** : Space station ops, multinational partnerships.

2.2 Statistical Analysis

The analysis in this section is not rigorous. It is intended to give a qualitative feel for where the hard parts of the problem are, how different missions are related to each other, and how hard problems correlate (for example, that there is a cluster of missions all of which score high on autonomy, uncertain execution, and complex dynamics).

2.2.1 Multi-dimensional Missions

The missions that score highest for having multiple difficulties are:

Missions	Autonomy	Resources	Overscribed	Coordinated	System dynamics	Multi-step	Unpredictability	Distributed	Real-time	Model drift	Replanning
Manned Mars		x	x		x				x	x	x
Constellation-X		x	x	x	x			x	x		
Space Tech 5	x			x	x			x	x		
Mag Constellation	x			x	x	x		x			
Europa Cryobot	x					x	x	x	x		
Shuttle Planning and Ops		x	x			x		x			x
MSL	x					x	x		x		
Jupiter Icy Moons Orbiter	x				x	x				x	
ARES	x					x	x		x		

2.2.2 Analysis of Criteria

As a first-order look at how the criteria served to differentiate the missions, we offer the following table and observations:

	Critical		Significant	
Criterion	All Missions	Top Nine	All Missions	Top Nine
Autonomy	13	6	30	8
Resource-bounded	18	3	34	7
Over-scribed	15	3	29	5
Coordinated multi-platform	9	3	22	5
Complex system dynamics	11	5	24	8
Multi-step	12	6	38	8
Unpredictability	7	3	29	7
Distributed	8	5	29	7
Real-time constraints	8	5	29	7
Model drift	4	2	30	7
Costs on replanning	3	2	28	6

Low ranking of enterprise planning is surprising, possibly driven by just HOW hard the problem needed to be to get a 9.

Also noteworthy that variation across criteria drops considerably when looking at those that show up in high-difficulty missions.

Another interesting measure to look at is correlation: how well one criteria's presence predicts the appearance of another, but we didn't do that.

2.2.3 Clustering

As part of the mission analysis, we looked at various simple clustering schemes. One that appears to group things in useful ways was a K-Means clustering assuming 6 clusters, using an L1 norm

(Manhatten distance between points).

The clusters found and the criteria values were as follows:

Cluster 1

	Criteria										
Missions	Autonomy	Resources	Overscribed	Coordinated	System dynamics	Multi-step	Unpredictability	Distributed	Real-time	Model drift	Replanning
Manned Mars	3	9	9	3	9	3	3	3	9	9	9
Shuttle Mission Planning and Ops	1	9	9	1	1	9	3	9	1	1	9
Hubble SM4	1	9	9	1	1	3	3	9	1	1	3
SOFIA	1	9	9	1	9	1	3	3	1	3	3
Astro-E2	1	9	9	1	1	1	1	1	1	3	3
Earth Observing Mission Planning and Ops	1	9	9	3	1	3	1	3	1	1	3
Herschel	1	9	9	1	3	3	1	1	1	3	3
GLAST	3	9	9	1	1	3	1	1	1	3	3
JWST (NGST)	3	9	9	1	3	3	1	1	1	3	3
WISE	3	9	9	1	1	3	1	3	1	3	3
Deep Impact	3	3	9	1	1	1	3	3	3	1	1
SIRTF	3	3	9	1	1	3	1	1	1	3	3

Large communities (Cluster 1) – These missions are primarily over-subscribed, with a difficult resource problem, and involve a significant degree of complication in replanning. This category includes large airborne and spaceborne telescopes, as well as Mission Ops support for manned space and near-earth spacecraft (which overlap to a significant extent with space telescopes).

Cluster 2

Missions	Criteria										
	Autonomy	Resources	Overscribed	Coordinated	System dynamics	Multi-step	Unpredictability	Distributed	Real-time	Model drift	Replanning
Mars Exploration Rovers	3	3	3	3	3	3	9	9	3	9	3
Shuttle flight No. 120	1	9	9	3	1	3	3	3	3	1	3
Deep Space Mission Planning and Ops	1	3	3	1	1	9	1	3	1	3	3
Shuttle flight No. 114 - Atlantis flight No. 27	1	9	3	3	1	3	3	3	3	1	3
Shuttle flight No. 115 - Endeavour flight No. 20	1	9	3	3	1	3	3	3	3	1	3
Shuttle flight No. 116 - Atlantis flight No. 28	1	9	3	3	1	3	3	3	3	1	3
Shuttle flight No. 117 - Endeavour flight No. 21	1	9	3	3	1	3	3	3	3	1	3
Shuttle flight No. 118	1	9	3	3	1	3	3	3	3	1	3
Shuttle flight No. 119 - Atlantis flight No. 29	1	9	3	3	1	3	3	3	3	1	3

Resource intensive operations (Cluster 2) – These missions are primarily difficult due to resource contention. Their main distinction from the previous category is that over-subscription (fielding requests from a large user or investor community) is less of an issue for these missions. Not a non-issue, necessarily, as can be seen by the inclusion in this category of deep space mission operations. The distinction is a matter of degree.

Cluster 3

Missions	Criteria										
	Autonomy	Resources	Overscribed	Coordinated	System dynamics	Multi-step	Unpredictability	Distributed	Real-time	Model drift	Replanning
SIM	3	3	9	9	3	3	3	3	9	3	3
EUSO	1	3	1	1	1	1	1	1	1	3	9
Mars Scout 2	3	3	1	1	3	3	3	1	9	1	3
SCIM	3	1	1	1	1	3	3	1	9	1	1

Real-time, limited autonomy (Cluster 3) – This group of missions is the smallest, with only three missions included. The missions included can be characterized as those with real-time constraints that are not either highly autonomous, or requiring close coordination among multiple platforms (e.g., for multi-platform inteferometry).

Cluster 4

Missions	Autonomy	Resources	Overscribed	Coordinated	System dynamics	Multi-step	Unpredictability	Distributed	Real-time	Model drift	Replanning
	Criteria										
Jupiter Icy Moons Orbiter	9	1	1	1	9	9	3	1	3	9	3
New Horizons (Pluto)	9	1	1	1	9	3	1	1	3	9	1
Space Tech 6	3	1	1	1	9	9	1	1	3	1	1
Dawn	9	1	1	1	3	9	3	1	3	3	1
Mars NetLander	9	1	3	9	3	3	3	3	3	3	1
THEMIS	9	1	1	9	1	3	3	3	3	3	1
MESSENGER	3	3	3	1	3	9	1	1	3	3	1
MUSES-C	9	1	1	1	1	3	3	1	3	3	1

Long-duration deep space missions (Cluster 4) – These missions include the Jupiter Icy Moons Orbiter (JIMO) and a proposed mission to Pluto and other Kuiper Belt objects, among others. The hallmark issues for this group include autonomy and complex, multi-step operations, coupled with largely standalone operations (so, coordination with other stakeholders is not a major issue in replanning).

Cluster 5

Missions	Autonomy	Resources	Overscribed	Coordinated	System dynamics	Multi-step	Unpredictability	Distributed	Real-time	Model drift	Replanning
	Criteria										
Constellation-X	3	9	9	9	9	1	1	9	9	3	3
Mag Constellation	9	3	1	9	9	9	1	9	3	3	1
Space Tech 5	9	1	1	9	9	3	3	9	9	1	1
LISA	3	1	1	9	9	1	1	3	9	3	3
Mag Multiscale	3	3	1	9	9	1	1	9	3	3	1
GEC	3	3	1	9	9	3	3	3	3	3	1

Constellations (Cluster 5) – These missions first and foremost require coordinated operations in the presence of complicated system dynamics (plans that involve tightly synchronized orbits, for example). Real-time constraints and distributed problem solving (driven in large part by the real-time constraints) are other significant features.

Cluster 6

Missions	Autonomy	Resources	Overscribed	Coordinated	System dynamics	Multi-step	Unpredictability	Distributed	Real-time	Model drift	Replanning
Europa Cryobot	9	3	3	3	3	9	9	9	9	3	1
ARES	9	3	1	1	3	9	9	1	9	3	3
Mars Science Laboratory	9	3	3	1	3	9	9	3	9	3	3
Mars Sample Return Lander	3	1	1	1	3	9	9	3	9	3	1
SAGE	9	3	3	3	1	3	9	1	9	3	1
Titan Aerobot Multisite	9	3	3	3	3	9	9	3	3	3	1

Rovers (Cluster 6) – These missions generally involve a planetary rover, aerobot, or probe. These are the missions with the broadest range of challenges, including unpredictable execution, real-time constraints, multi-step plans, and a high degree of autonomy.

2.2.4 Concluding Remarks

As stated above, the analysis in this section is intended to guide intuitions more than to provide a rigorous analysis. That said, we found it instructive to consider what makes a mission planning problem hard, how missions shared common structure among those criteria, and how criteria were correlated (i.e., when and where missions with one difficulty were likely to be hard in other ways as well).

The analysis performed for these missions can easily be modified to include other missions, other criteria, or other views on appropriate criteria rankings for a given mission.

A couple of points are worth further emphasis. First, missions grouped according to what kinds of planning and scheduling they need do not necessarily line up in predictable ways. A second and more critical point is that no class as a whole, and very few individual missions, are hard in only one way. This has implications for approaches or algorithms that require restrictive assumptions in order to work. For example, a technique that is very promising for uncertain execution, without any prospects for extension to real-time operations (or resource reasoning, or optimization over conflicting goals) may not be all that useful.

Chapter 3

Planning Technology Requirements

The 11 criteria against which the missions have been evaluated are requirements for functional capabilities that must be satisfied for the mission to succeed. Precisely what it means to satisfy these requirements depends on the nature of the mission, including especially and in particular what other functional requirements must be satisfied.

For each of the criteria, we discuss the technical approaches that may be used to implement an appropriate functionality. This discussion is necessarily at a high level, due to the breadth of criteria covered (and thus the size of the design space, when different aspects of a given mission profile are considered in combination), and to the dependence of what technique is most suited on the specific characteristics of that mission.

3.1 Autonomy

The relevance of a requirement for autonomy on approaches to planning and scheduling is indirect: the precise nature of the constraints imposed is dictated by the architecture within which planning and scheduling will be done. The ITAS report [59] and what of the CICT "Collaborative Decisions Systems" wedge planning has been made public both discuss autonomy as a separate requirement, to be satisfied in part by planning and scheduling technology.

What specific technology depends on the architecture within which the planner is to be employed. Claraty [72] is a top-down, three-level architecture with a planner as the top level and thus in ultimate control. IDEA [55] uses a planner as a service, generating an artifact to be employed by the "plan runner" that does the actual execution. Remote Agent [56] had several planners, including a top-level mission planner whose output was then decomposed and further elaborated by the Executive (which was thus doing some planning and scheduling itself), and a "reactive planner" that responded to faults as part of the diagnostic and health management component.

3.2 Resource Bounds

Methods commonly employed for handling resource bounds in planning and scheduling fall broadly into three categories.

Scheduling techniques treat resources as primary. There are a number of different approaches to actually generating a schedule, but there are some common elements among all of them, including an emphasis on modeling resources and finding solutions to resource conflicts, as opposed to reasoning about linking sequences of activities together to achieve a desired effect.

Temporal planning techniques, including constraint-based planning, takes a middle course, representing both state information and resource contention in constraints, then solving the resulting Constraint Satisfaction Problem (CSP) or Constrained Optimization Problem (COP). These approaches tend to do a better job of handling the planning part of the problem (generating sets of activities to achieve a given goal), but are generally not tuned as well for the resource part of the problem, which can lead to problems with scaling up to larger problems.

Planning with Resources is meant to describe a variety of recent developments in classical planning, all intended to add capabilities for reasoning about resource usage, including time elapsed, to the STRIPS-rule representation and goal regression techniques commonly used in classical planning, for example in many of the systems employed in the International Planning Competition.

3.3 Oversubscription

Oversubscription is a major differentiator between two styles of scheduling. The first of these starts with a fixed set of tasks and tries to accomplish the entire set within a given set of constraints, possibly minimizing a function such as makespan (time to complete all the tasks) or tardiness (number of time units by which those tasks with specified deadlines miss their deadlines). The second style of scheduling is the one we would call oversubscribed. In this problem, the available resources are fixed, and the problem is to fit in as many tasks as possible, generally with some kind of weighting or priority as to which tasks are the most important or urgent, and quite frequently also with costs for deviating from the preferred execution of a given task (moving it in time, or changing resource usage).

A similar notion of oversubscription has been investigated in classical planning, with "tasks" that may or may not be completed in this case consisting of goals that either are or are not satisfied [71]. The analogy is not perfect, because tasks and goals are not equivalent.

3.4 Coordinated Operations

Coordination between different entities can take place at any of several levels, depending on the level of autonomy, and the nature of the domain. As a simple hierarchy, we can talk about

- Coordinated control
- Coordinated plan execution ("task-level control")
- Coordinated planning
- Negotiation over roles

These different flavors of coordinated behavior involve very different techniques, ranging from short-latency distributed control, through the use of discrete synchronization mechanisms, to explicit negotiation strategies.

Coordinated operations, as the term is used in this report (and in particular in the discussion of mission requirements) refers to coordinated execution and control. Coordinated planning and negotiation over roles is covered below, under a different term: Distributed Operations.

How coordinated execution and control are implemented depends strongly on other factors. For example, a spacecraft constellation comprising a Very Long Baseline Interferometry (VLBI) telescope needs to maintain relative positioning among the different spacecraft to within a fraction of the wavelength of radiation being sensed. The precision of the required control thus depends on the type of telescope. In addition, due to the nature of the task, the "coordinated execution" aspect of this problem is relative simple: observations need to start and stop at roughly the same time (and be directed at the same field of the sky).

A special case that needs to be addressed (and is being addressed in current work) is when the coordination involves both humans and automated systems ("Adjustable Autonomy" or"Human-Robot Teams"), as for example in the Personal Satellite Assistant, or in the design of automation for a long-term human presence on Mars.

3.5 Distributed Operations

Distributed operations addresses the need to negotiate plans among different entities. "Entity" in this case could refer to multiple rovers exploring a planetary surface, all the way up to negotiations among national space agencies involved in a multi-national mission.

The applicable techniques in this case depend on the granularity of the operations involved (are we negotiating over who will visit a given geological site, or over who is supporting the next resupply flight?), and the degree of human involvement. With specific respect to planning and scheduling, applicable techniques across this spectrum include market- and auction-based approaches to task allocation, collaborative multi-agent planning, asynchronous multi-user access to a centralized plan (more a database issue than a planning and scheduling issue), and for a local plan, limited replanning and plan repair.

3.6 Multi-Step Operations

Planning and scheduling problems involving the generation of multi-step operations range from rover tours closely related to constrained versions of the Travelling Salesman Problem, to scheduling experiments or maintenance tasks as part of Space Station operations, where a relatively small sequence or network of tasks must be fit into a much larger set of existing tasks, all contending for the same resources.

Consequently, there are several approaches that may be taken to resolving such problems, including classical planning augmented to handle metric information, constraint-based planning as in Europa [37], HTN planners like SHOP2, and in some cases simple forward-chaining production systems that generate tasks based on a set of rules.

3.7 Complex System Dynamics

Especially for space applications, there are planning and scheduling problems where the continuous dynamics of the system being simulated (during planning and scheduling) and controlled (during execution) go well beyond reasoning about task durations and deadlines. Orbital dynamics, fuel/energy consumption, driving in loose material, ephemeris computations (for telescope observations, or solar panel pointing and flux calculations) are a few of the areas where complex continuous models may need to be modeled explicitly.

For problems where the projective horizon is a long one, and where simple bounding approximations do not suffice, this will require the use of a hybrid solver or simulator of some form. For problems that are closer to control problems (limited predictive horizon), the use of frequently-consulted simple approximate models as in Model Predictive Control will be a better fit.

3.8 Uncertain Execution

As the term is commonly used, "uncertain execution" can cover both acting with limited information (the outcome of actions would be deterministic, *if* you knew enough about the state in which those actions were executed), and actions with uncertain outcomes (no amount of state information will suffice).

The important distinction from the point of view of planning and scheduling is whether or not the information in question can be observed or tested for. If not, then the action might as well be modeled as having an uncertain outcome. If it can, then the plan may include actions whose rationale in whole or in part is specifically to discover that information.

In this case there are two additional complications, however. First is the fact that it may not be possible to discover the information required at plan time, as opposed to at run time. This means that new information relevant to the plan will be discovered while the plan is being executed, leading to a need for contingent planning, or replanning, if that information is to be taken into account (in some cases this means you ignore the information and construct a *conformant* plan. The other complication is that the actions by which the information is discovered may have other, unintended effects, for example a cost in terms of time or resource usage. In this case, the cost of finding out needs to be weighed against the cost of not knowing.

A wide variety of methods can be applied here, depending on details of the problem being addressed (and the hammer you choose to apply, since few of these are mature technologies), including contingent planning, conformant planning, reactive planning, rapid (thus, generally local) replannning, planning to gather information, stochastic/decision-theoretic planning, and others.

3.9 Real-Time Requirements

Requirements for real-time behavior (meaning time guarantees, not just "fast") can include real-time planning, or generating real-time responses, depending on mission details such as the necessary level of autonomy and the degree to which planning can be done well in advance of execution.

Real-time planning and scheduling is a significant problem, because planning and scheduling are computationally intensive activities, frequently employing algorithms such as heuristic search, which have unpredictable run-times and are thus difficult to provide performance guarantees for. Real-time execution compiling a plan or policy into a form amenable to execution in a bounded amount of time, for example bounded-depth rules in a production system (rule-based executives are a subset of this class).

3.10 Model Drift

One of the significant issues that arose as we studied the range of missions described in Chapter 2 and Appendix A was the need to support long-term missions, for example the decade or more that will be spent by some of the outer-planet orbiters, for example JIMO.

A significant problem with this kind of time-scale is that the model being used to control the spacecraft will change in ways both large and small. A small change might consist of increased drag or increased play in the rotation of an instrument platform. A large change might be the failure of a scientific instrument, or worse, a navigation sensor.

This is a qualitatively different issue than that currently addressed by model-based executives such as Titan [74], which seek to insulate planning, scheduling, and execution from details of the underlying vehicle or mechanism. Such systems typically respond to failure or degradation by attempting to preserve the services expected by the higher-level functions. In other words, they try to keep the plant consistent with the model, rather than concentrating on adjusting the model to fit the current state of the plant.

Biswas' work on Fault-Adaptive Control, funded under the Intelligent Systems program, is relevant here [45].

3.11 Replanning Cost

This criterion was included to capture situations where changing the plan is difficult or expensive. There are several reasons this might be the case. For example, multi-agency or multi-national mission operations will involve complex negotiations among multiple parties, as discussed above. Once the plan or schedule for a particular scope of operations has been set, what if something breaks? Replanning in these cases will be more complex than, for example, for an autonomous exploratory vehicle the operations of which affect nothing other than itself, as long as it keeps to specified communication windows and protocols.

Ways to reduce the cost incurred for replanning fall broadly into two categories. The first is to limit the scope of the replanning needed, the second is to generate the plan initially so as to reduce the frequency and scope of the replanning needed ("flexible plans"). There is, of course, no reason that both approaches cannot be implemented at the same time.

3.12 Mixed-Initiative Solving

This was not one of the mission criteria we looked at directly, but it is relevant, and is related to some of the other criteria. Mixed initiative solving is not quite the same thing as autonomy, but as with autonomy, mixed-initiative planning and scheduling is in significant part an architectural issue. In other words, the algorithmic requirements are driven by the nature of the interaction you want to support.

Another issue arising in mixed-initiative systems is the need for the human interacting with the system to understand what it has done, is doing, or is about to do.

So, replanning, explanation, local plan repair, especially in the presence of user-directed plan modifications that must be maintained, are all techniques that support mixed-initiative planning and scheduling. There are several projects within the Intelligent Systems program that are some measure addressing the issue of mixed-initiative planning and scheduling.

3.13 Resulting Technical Requirements

From the discussion above, we can generate a set of technical requirements, which can then be used to evaluate proposed or existing research programs for relevance to the mission criteria.

- Autonomy architectures
- Integrated planning and execution
- Integrated planning and control
- Scheduling
- Temporal planning
- Planning with resources
- Optimization
- Explanation
- Negotiation
- Coordinated execution
- Classical planning
- Constraint-based planning
- Hierarchical Task Network planning
- Procedural executives
- Complex hybrid models
- High-level control
- Contingent planning
- Reactive planning
- Rapid replanning
- Information-gathering plans
- Stochastic/decision-theoretic planning
- Collaborative planning

- Mixed-initiative planning
- Real-time planning
- Generating real-time responses
- Model updating
- Acting to gain information
- Local plan repair

3.14 Implementation/Engineering Issues

Another set of issues arises in thinking about the implementation and maintenance of planning and scheduling systems. These issues are outside the scope of the current report, but too important not to note.

- Knowledge engineering for planning (construction, validation, and maintenance of planning models)
- Validation and verification of planning models and algorithms
- Integration with health management (diagnosis, prognostics, automated testing)
- Integration with other decision-making functions
- Software architectures, including but not limited to
  - Compilation approaches
  - Reliable systems
  - Real-time systems
- Multi-agent coordination protocols
- Learning
  - Domain models
  - Plans
  - Policies
  - Heuristics

3.15 Summary

The discussion in this section makes it clear how large a set of disparate technical problems are encompassed within NASA missions requiring some form of planning and scheduling. A second issue highlighted here is how much the approach must depend on specific details, and, further, on details relevant to many different criteria: it is not as simple as a "Chinese menu" permitting independent choices from different columns (corresponding in this case to the different criteria).

Several interesting examples of this arise in considering the interaction between the need for autonomy and other criteria. For example, a constellation with only limited autonomy requirements (e.g., exploring the terrestrial magnetosphere, or an L2 VLB interferometry mission) has a coordinated *control* problem, in that once operations are defined for all elements of the constellation they must then be executed in precise synchrony. In order for there to be a distributed planning problem as well, other mission characteristics must be present, for example a need for autonomy for either the constellation as a whole or the individual elements, based on light-speed lag or other communication limitations, or a bureaucratic separation of authority (an interferometry operation using spacecraft owned by different missions, or different countries).

Another interesting combination to look at is real-time requirements coupled with uncertain execution, especially in the presence of coordinated operations.

Chapter 4

Coverage of Requirements by Projects

In this chapter, we summarize in tabular form how work being done at, or being funded by, NASA satisfies the capabilities discussed in Chapter 3. Much though not all of this work has been funded by the Intelligent Systems program. Brief descriptions and citations for these projects can be found in Appendix B.

As the table makes clear, these projects do a very good job of covering the technical needs discussion in previous sections. There is some room for additional work going forward, both in these areas (funded projects does not mean that the last word has been written, but rather that the right objectives are being pursued), and in some new areas. New areas recommended for attention are discussed in Chapter 5.

	Autonomy architectures (incl. plannig)	Planning & execution	Integrating planning & control	Scheduling	Temporal planning	Planning with resources	Optimization	Explanation	Negotiation	Coordinated (planning and) execution (control)	Classical Planning	Constraint-based Planning	HTN Planning	Procedural executives	Complex continuous/hybrid models	High-level control	Contingent planning	Reactive planning	Rapid replanning	Information-gathering plans	Stochastic/decision theoretic planning	Collaborative planning	Mixed-initiative	Real-time planning (not "fast")	Gernerating real-time responses	Model updating (both continuous and discrete)	Acting to gain information	Local plan repair
Agent Development and Verification				✓	✓																							
Continual Team Planning									✓	✓												✓						
Team Sequence Execution		✓							✓	✓									✓			✓						
Autonomy Verification and Validation																								✓				
Human-Automation Interaction																							✓					
Livingstone Diagnostic Agent		✓	✓																✓							✓		
Probabilistic Hybrid Fault Detection															✓						✓							
Spacecraft Mobile Robot																									✓			
Integrated Resource and Path Planning	✓				✓	✓						✓																
Integrated Planning and Execution	✓	✓	✓	✓		✓						✓																
SOFIA Scheduling				✓								✓			✓		✓	✓		✓				✓				✓
Intelligent Specification-Centered Test Case Generation								✓															✓					
Model Usability																		✓										
Constraint-based Planning						✓	✓																			✓		
Multi-Resolution Planning						✓					✓		✓			✓								✓				
Human-Centered Software Development								✓															✓					
Agents for Distributed Team Operations				✓					✓													✓						
Assistant Systems for Mission Control				✓																								
IDEA Autonomy Architecture									✓									✓				✓		✓				✓
Rover Autonomy Architecture	✓		✓			✓			✓			✓																
Probabilistic Reasoning	✓					✓														✓								
MER Mixed-Initiative Planning						✓																		✓				✓
Intelligent Launch and Range Operations	✓		✓		✓			✓				✓										✓	✓					
Onboard Science Understanding						✓														✓							✓	
Distributed Control of Life Support	✓			✓																		✓	✓					
Multi-Rover Coordination			✓	✓		✓			✓							✓						✓						
Combinatorial Optimization Planning	✓			✓												✓												
Concurrent Contingency Planning					✓		✓										✓											
Stochastic Anytime Planning																		✓			✓							
Onboard Rover Autonomy																✓												
Model-based Reactive Control																	✓	✓										
Interleaved Contingent Planning and Execution																✓												
Autonomous Rotorcraft	✓		✓						✓									✓										
Hybrid Health Management and Control				✓	✓					✓																		
Autonomous Rover Command Generation - ASPEN				✓		✓										✓												
Adaptive Problem Solving (APS)				✓																								
CASPER (Continuous Activity Scheduling Planning Execution and Replanning)															✓													
CLEaR (Closed Loop Execution and Recovry)	✓														✓													
Unamanned Air Vehicles (UAVs)			✓												✓													
Citizen Explorer (CX1)				✓		✓																						
Distributed Self-Comanding Robotic Systems																	✓							✓				
Techstat-21																	✓											
Three Corner Sat (3CS)																						✓						

Distributed Rovers/MISUS (Multi-Rover Intergrated Science Understanding Systems)																					✓							
Onboard Planning for Rocky7																												
DSSC (Deep Space Station Controller)									✓												✓							
Scheduling with Resource Enevelopes					✓																							
Spacecraft Mobile Robot Autonomy																												
Antarctic Robotics Traverse																												
Biomorphic Robotics																✓												
Coupled Layer Architecture for Robotic Autonomy (CLARAty) Architecture											✓																	
Executive-Level Decision Making																✓												
Intergrated Technolygy Demonstrations for Mars Science Laboratory (MSL)				✓												✓												
K9 Platform, Architecture and Test Facility																												
2003 Mars Exploration Rovers (MER) Mixed Initiative Plan Generator (MAPGEN)				✓																	✓							
2009 Mars Science Larboratory (MSL) Mission Planning Execution Project																	✓											
Look-ahead Model Based Programming																								✓				
Plan Works																✓												
Advanced Information Systems Technology (AIST) Earth Observing Satellite Scheduling				✓		✓																					✓	
Constraint Based Planning					✓	✓																						
Imagebot																✓												
SOFIA Observation Scheduling				✓											✓													
ScienceDesk																					✓							
Aviation Data Integrations																												
Mars Exporation Rover Human Centered Computing																						✓						
MERBoard																					✓	✓						
Enabling Knowledge Management for Organizations Risk Analysis																												
Postdoc																						✓						
Secure Advanced Feerated Environment (SAFE)									✓												✓							
IxTeT-eXeC		✓			✓											✓		✓										
Dynamic Ontology Refinement		✓	✓															✓	✓							✓		
High-level Robot Programming and Program Exeuction																✓		✓						✓				

Chapter 5

Areas for Emphasis

As shown in Chapter 4, the current coverage of planning and scheduling needs by NASA's IS program and other NASA funding is broadly complete. However, there are some areas where further attention might be appropriate in the near term. Here is a summary of areas we have identified for additional attention in the next five to ten years:

- Updating for drift or failure
- Optimization
- Real-time response
- Planning and Control
- Integrating Planning with Other Capabilities
- Planning and Scheduling as Engineering

These areas are discussed in more detail, below.

5.1 Updating for Drift or Failure

Many future NASA missions will be of long duration. This is true for both manned and unmanned missions (e.g., Mars Exploration, or JIMO, respectively). This has implications both large and small, in terms of ways the underlying system will change.

These changes can be divided as follows:

- Changes that can be handled by the control system, perhaps augmented as by Biswas' Fault Adaptive Control [45].
- Changes that can be handled by some form of "Mode Identification and Recovery," as on Remote Agent.
- Significant changes to the model, affecting the operation of the spacecraft or mechanism itself. This is not easily addressed in current approaches to autonomy, requiring a tighter integration of planning and diagnosis than is usually attempted.

Adaptive solvers are relevant, and are being worked on within NASA, e.g. [15], but only address part of the problem. Adaptive heuristics will not help in cases where the underlying model has changed, such that the predictions made by the model of system behavior for a given plan are significantly in error.

Similarly, model-based planning/exercuion/control is relevant but not the whole answer. The issue is how the planner responds to changes in the model, and that is an issue that these systems have not yet addressed in any detail.

5.2 Optimization

Optimized planning and scheduling is being pursued on several projects within NASA, for example Knight and Smith's work adding combinatorial optimization to planning, Frank's SOFIA planner [40], and Meuleau and Smith's TSP w. rewards [53].

There is ample room for further work in this area. For one, there is a broad set of missions for which optimized plans and schedules are necessary, for example any mission involving *oversubscription*. Oversubscription refers to a situation in which there are more tasks to be accomplished than resources available. Scheduling observations on Hubble is a good, if extreme, example of oversubscription.

Other reasons to optimize including extending mission life, reducing operational risk, and minimizing disruption to a previous schedule in rescheduling, among many others.

As shown in Chapter 2, missions will in general have multiple characteristics affecting the nature of the planning and scheduling problem(s) presented. This is a significant effect when considering optimization as well. For example, in missions involving a considerable degree of autonomy, any solution must address the problem of how to optimize and still preserve predictable, in particular timely, solving behavior. Another difficulty that planning and scheduling in this domain shares with others is the problem of optimizing in the presence of multiple objectives. Mission life and safety, for example, are two objectives that when optimized individually may conflict (in other words, a safer schedule is likely to be a shorter one, and a longer one less safe).

5.3 Real-time response

Real-time behavior in a high-assurance environment (which will be a requirement, again, for either manned exploration or long-duration, high-autonomy missions) is a property that must be implemented and enforced as an integral property of the system architecture. As such, a "real-time planner" is a component of a larger system, and probably not the one in control, 10 years of Three Layer Architectures to the contrary. Something like IDEA [55], in which the executive is in control, is more likely to have the required properties.

However, the planner will need to be provably real-time in the sense of delivering *some* kind of response in a predictable way. See Ben Wah's work on using the calculus of variations [10], or Tony Barrett on U(n) plans [4], or Drummond and Bresina's "Situated Control Rules" from a few years ago [25], for a few of the many ways that have been explored to implement real-time planning.

5.4 Planning and Control

There are two issues, both relevant to NASA applications of planning and scheduling. The first is how to construct planning systems that interact effectively with underlying control systems. This issue has been addressed to a considerable extent in work on autonomy and execution architectures, including a lot of work at NASA [55, 72, 6], and in work on generating *flexible plans* that provide more leeway to an executive to adjust for vagaries in the actual course of events.

The other issue is the explicit treatment of the interleaving of planning and execution as a control problem. The ICAPS-03 Workshop on Planning and Execution contained a number of papers which touched briefly on this area but none that addressed it explicitly, much less systematically.

The argument for explicitly viewing the planning and execution loop as a control loop, or at the least strongly analogous to a control loop. is that this view facilitates the use of further analogies. For example, modern control systems typically employ predictive models to project to a horizon that is well beyond the next few control inputs. The models employed are not very accurate, and do not have to be, because this projection is done for the computation of each new control input, so the error in the models is corrected for by computing from the actual, rather than the predicted, state at each control update.[1] Potential dead ends that may not be avoidable due to the limited predictive horizon can be prevented through the use of "prohibited regions," which are essentially a partial policy.

Finally, the control community has explicitly addressed the tradeoff between sample rate and computational expense. Any autonomous execution loop will have to implement some kind of tradeoff in this regard, and better it be an explicit one, subjected to a principled analysis. For planning and scheduling applications, this analysis will address the tradeoffs among reactive planning, re-planning, and flexible plans, in the context of a particular application domain. To date, this kind of analysis is largely missing.

5.5 Integrated Capabilities

As argued in Chapter 3, broad set of techniques that may apply for an individual piece of the problem (e.g., contingent planning) can be filtered by considering other requirements (e.g., we have to handle uncertain execution in the presence of strict resource bounds on execution, and achieving any given task requires multiple steps).

Several interesting examples of this arise in considering the interaction between the need for autonomy and other criteria. For example, a constellation with only limited autonomy requirements (e.g., exploring the terrestrial magnetosphere, or an L2 very long baseline interferometry mission) has a coordinated *control* problem, in that once operations are defined for all elements of the constellation they must then be executed in precise synchrony. In order for there to be a distributed planning problem as well, other mission characteristics must be present, for example a need for autonomy for either the constellation as a whole or the individual elements, based on light-speed lag or other communication limitations, or a bureaucratic separation of authority (an interferometry operation using spacecraft owned by different missions, or different countries).

Another interesting combination to look at is real-time requirements coupled with uncertain execution, especially in the presence of coordinated operations.

[1] The GPSS scheduler for Shuttle refitting worked in this way, though the developers did not to my knowledge make this connection explicit [78].

A third place where this need for integration is apparent, and not currently being addressed adequately, is in integrated planning and diagnosis, or more generally a tighter integration between reasoning *about* a physical model and reasoning *with* a physical model.

5.6 Planning and Scheduling as Engineering

The final area we suggest for further attention with specific regard to planning and scheduling is the maturation of planning and scheduling as an engineering discipline. Currently, implementing planning and scheduling systems, and applying those systems to particular domains, is very much a boutique art.

Given that much of the current effort is still in the realm of research, this is probably invevitable, but the situation will have to change before planning and scheduling solutions find broad acceptance in mission applications.

The computational complexity of planning and scheduling, the size and complexity of the models that must be generated, and the sensitivity to modeling choices and application characteristics in the performance of the resulting systems combine to make reducing planning and scheduling practice to engineering a significant challenge. However, there are precendents to draw on.

For example, the application of various flavors of mathematical optimization has many of the same characteristics. Work in this area has a longer history, and can thus be examined to see what has worked and what has not. Among the significant features of current practice in mathematical optimization:

- The development of generic, standardized tools. Not a single tool, because there are different models, nor even a single algorithm for any given tool, because different models impose different costs and requirements.
- Standard modeling languages. This process is underway in the planning and scheduling commmunities, but has not yet resulted in a language (or a small number of languages) generally acknowledged to be sufficient in the way that GAMS and AMPL are for mathematical optimization.
- The emergence of domain modeling and tool application as a discipline in its own right. In math. optimization, this process has progressed to the point where there are multiple disciplines, focussed on different classes of applications.

All of these things are achievable for planning and scheduling. None have yet been achieved.

Chapter 6

Summary and Conclusions

This report summarizes a moderately-detailed investigation into mission requirements and how those requirements map into technical requirements for planning and scheduling research over the next 5 to 10 years. Chapter 2 presents the criteria according to which the missions were evaluated, and summarizes our findings. More detailed mission analysis data can be found in Appendix A. Chapter 3 maps from the mission criteria to a set of technical capabilities which can be applied to current research to evaluate the relevance of that research to the mission needs identified. Chapter 4 presents the summary analysis of coverage for these requirements, against projects currently active in NASA, or being funded by the Intelligent Systems Program. Detailed information on the projects discussed can be found in Appendix B. Chapter 5 provides a brief discussion of areas identified for further work, given mission trends and current research coverage as discussed previously.

The overall conclusion of the report is that will limited exceptions, NASA's funding in the area of planning and scheduling is well-matched to the needs of future NASA missions. There is clearly room for further work in both existing and new areas, as one would expect in an active research program.

The analysis in this report makes it clear how large a set of disparate technical problems are encompassed within NASA missions requiring some form of planning and scheduling. A second issue highlighted here is how much the approach must depend on specific details, and, further, on the interaction among different mission characteristics. Several interesting examples of this arise in considering the interaction between the need for autonomy and other criteria.

This study also introduced a set of criteria and a method of analysis over missions to highlight functional needs. This methodology has the potential to prove useful on an ongoing basis, as both mission requirements and the current state of the art evolve.

As a case in point, the ground has changed somewhat since this investigation was started, with the constitution of the new Exploration Enterprise (Code T) and the incorporation of CICT into Code T. While this shift in emphasis towards Exploration implies some readjustment of priorities, most of the analysis here should hold up well, for two reasons. First and most importantly, because we took exploration and manned missions into account—the mission that ranks the highest in our analysis in Section 2 is Manned Mars. The second reason that this work holds up is because the requirements for planning and scheduling do not change all that much.

While the full set of requirements and priorities have yet to be worked out, preliminary indications make it clear that much of what was needed, is still needed. Human presence on the Moon or

Mars makes autonomy more important, not less. Automated or semi-automated housekeeping, maintenance, lifesupport, exploration, mining operations, are all functions that will be crucially important to maintaining mission safety and effectiveness at a reasonable cost (or perhaps at any cost).

From unofficial conversations and one presentation, it also appears as though there will be some movement in the direction of surface exploration rather than orbiters, as well as further development of coordinated multi-spacecraft or multi-robot missions. Both of these tendencies will increase the need for, and the difficulty of, effective planning and scheduling. Adjustable autonomy in operations (which is not a planning and scheduling requirement per se, but has planning and scheduling implications), and mixed-initiative solvers are both clearly going to be relevant.

Chapter 7

Bibliography

Bibliography

[1] *Interim Report of NRC Review of NASA's Pioneering Revolutionary Technology Program*, National Academies Press, 2003.

[2] M. Ai-Chang, J. Bresina, L. Charest, J. Hsu, A. Jónsson, B. Kanefsky, P. Maldague, P. Morris, K. Rajan, and J. Yglesias, "MAPGEN Planner: Mixed Initiative Activity Planning for the Mars Exploration Rover Mission," in *Proceedings of the International Symposium on AI, Robotics and Automation in Space*, 2003.

[3] S. Y. Alan. *Inductive Policy Selection for First-Order MDPs*.

[4] A. Barrett, "Domain Compilation for Embedded Real-Time Planning," in *Proceedings of the ICAPS-03 Workshop on Plan Execution*, 2003.

[5] T. Bedrax-Weiss, J. D. Frank, A. K. Jónsson, and C. McGann, "Identifying Executable Plans," in *Proceedings of the ICAPS-03 Workshop on Plan Execution*, 2003.

[6] D. Bernard, G. Dorais, C. Fry, E. Jr, B. Kanefsky, J. Kurien, W. Millar, N. Muscettola, P. Nayak, B. Pell, K. Rajan, N. Rouquette, B. Smith, and B. Williams, "Design of the remote agent experiment for spacecraft autonomy," in *Proceedings of the IEEE Aerospace Conference*, 1998.

[7] D. S. Bernstein, S. Zilberstein, R. Washington, and J. L. Bresina, "Planetary Rover Control as a Markov Decision Process," in *Proceedings of the 6th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*, Montreal, Canada, June 2001.

[8] G. Brat, D. Giannakopoulou, A. Goldberg, K. Havelund, M. Lowry, C. S. Pasareanu, A. Venet, and W. Visser, "Experimental Evaluation of Verification and Validation Tools on Martian Rover Software," in *Proceedings of the SEI Software Model Checking Workshop*, 2004.

[9] G. Brat and W. Visser, "Combining Static Analysis and Model Checking for Software Analysis," in *16th International Conference on Automated Software Engineering (ASE'01)*, 2001.

[10] Y. Chen and B. W. Wah, "Automated Planning and Scheduling using Calculus of Variations in Discrete Space," in *Proceedings ot the Thirteenth International Conference on Automated Planning and Scheduling (ICAPS) 2003*, 2003.

[11] S. Chien, T. Debban, C. Yen, R. Sherwood, R. Castano, B. Cichy, A. Davies, M. Burl, A. Fukunaga, R. Greeley, T. Doggett, K. Williams, V. Baker, and J. Dohm, "Revolutionary Deep Space Science Missions Enabled by Onboard Autonomy," in *Proceedings of the International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*, 2003.

[12] S. Chien, B. Engelhardt, R. Knight, G. Rabideau, R. Sherwood, E. Hansen, A. Ortiviz, C. Wilklow, and S. Wichman, "Onboard Autonomy on the Three Corner Sat Mission," in *Proceedings of the Sixth International Symposium on Artificial Intelligence and Robotics and Automation for Space (i-SAIRAS 2001)*, 2001.

[13] S. Chien, B. Engelhardt, R. Knight, G. Rabideau, R. Sherwood, D. Tran, E. Hansen, A. Ortiviz, C. Wilklow, and S. Wichman, "Onboard Autonomy Software on the Three Corner Sat Mission," in *Proceedings of the SpaceOps 2002 Conference*, 2002.

[14] S. Chien, R. Sherwood, M. Burl, R. Knight, G. Rabideau, B. Engelhardt, A. Davies, P. Zetocha, R. Wainright, P. Klupar, P. Cappelaere, D. Surka, B. Williams, R. Greeley, V. Baker, and J. Doan, "The Techsat-21 Autonomous Sciencecraft Constellation," in *Proceedings of the Sixth International Symposium on Artificial Intelligence and Robotics and Automation for Space (i-SAIRAS 2001)*, 2001.

[15] S. Chien, R. Sherwood, G. Rabideau, R. Castano, A. Davies, M. Burl, R. Knight, T. Stough, J. Roden, P. Zetocha, R. Wainwright, J. V. Gaasbeck, P. Cappelaere, and D. Oswald, "The Techsat-21 Autonomous Space Science Agent," in *Proceedings of Agents 2002*, 2002.

[16] S. Chien, R. Knight, A. Stechert, R. Sherwood, and G. Rabideau, "Using Iterative Repair to Improve the Responsiveness of Planning and Scheduling," in *Proc. 5th International Conference on Artificial Intelligence Planning and Scheduling (AIPS2000)*, Breckenridge, CO, April 2000.

[17] C. Chouinard, F. Fisher, D. Gaines, T. Estlin, and S. Schaffer, "An Approach to Autonomous Operations for Remote Mobile Robotic Exploration," in *Proceedings of the IEEE 2003 Aerospace Conference*, 2003.

[18] *CLOCKWORK: Requirements Definition and Technology Evaluation for Robust, Compiled Autonomous Spacecraft Executives.* http://www.adventiumlabs.org/publications.htm.

[19] R. Dearden and C. Boutilier, "Abstraction and Approximate Decision Theoretic Planning," *Artificial Intelligence*, vol. 1, pp. 219–283, 1997.

[20] R. Dearden and D. Clancy, "Particle Filters for Real-time Fault Detection in Planetary Rovers," in *International Workshop on Diagnosis (DX)*, 2001.

[21] R. Dearden, N. Meuleau, S. Ramakrishnan, D. Smith, and R. Washington, "Contingency Planning for Planetary Rovers," in *NASA 2002 Planning and Scheduling Workshop*, 2002.

[22] M. B. Dias, S. Lemai, and N. Muscettola, "A Real-Time Rover Executive Based On Model-Based Reactive Planning," in *The 7th International Symposium on Artificial Intelligence, Robotics and Automation in Space*, May 2003.

[23] M. B. Do and S. Kambhampati, "Improving Temporal Flexibility of Position Constrained Metric Temporal Plans," in *Proceedings ot the Thirteenth International Conference on Automated Planning and Scheduling (ICAPS) 2003*, 2003.

[24] G. A. Dorais and K. Nicewarner, "Adjustably Autonomous Multi-agent Plan Execution with an Internal Spacecraft Free-Flying Robot Prototype," in *Proceedings of the ICAPS-03 Workshop on Plan Execution*, 2003.

[25] M. Drummond and J. Bresina, "Anytime Synthetic Projection: Maximizing the Probability of Goal Satisfaction,", pp. 138–144, 1990.

[26] B. Engelhardt, S. Chien, A. Barrett, J. Willis, and C. Wilklow, "The DATA-CHASER and Citizen Explorer Benchmark Problem Sets," in *Proceedings of the 2001 European Conference on Planning (ECP2001)*, 2001.

[27] T. Estlin, R. Castano, B. Anderson, D. Gaines, F. Fisher, and M. Judd, "Learning and Planning for Mars Rover Science," in *IJCAI 2003 workshop notes on Issues in Designing Physical Agents for Dynamic Real-Time Environments: World modeling, planning, learning, and communicating*, 2003.

[28] T. Estlin, F. Fisher, D. Gaines, C. Chouinard, S. Schaffer, and I. Nesnas, "Continuous Planning and Execution for a Mars Rover," in *Proceedings of the 3rd International NASA Workshop on Planning and Scheduling for Space*, 2002.

[29] T. Estlin and D. Gaines, "An Optimization Framework for Interdependent Planning Goals," in *Proceedings of the AIPS 2002 Workshop for Planning and Scheduling with Multiple Criteria*, 2002.

[30] T. Estlin, R. Volpe, I. Nesnas, D. Mutz, F. Fisher, B. Engelhardt, and S. Chien, "Decision Making in a Robotic Architecture for Autonomy," in *Proceedings of the Sixth International Symposium on Artificial Intelligence and Robotics and Automation for Space (i-SAIRAS 2001)*, 2001.

[31] T. Estlin, F. Fisher, D. Gaines, C. Chouinard, S. Schaffer, and I. Nesnas, "Continuous Planning and Execution for an Autonomous Rover," in *Proceedings of the Third International NASA Workshop on Planning and Scheduling for Space*, Houston, TX, October 2002.

[32] T. Estlin and R. Mooney, "Learning to Improve Both Efficiency and Quality for Planning," in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, (IJCAI-97)*, Nagoya, Japan, August 1997.

[33] T. Estlin, G. Rabideau, D. Mutz, and S. Chien, "Using Continuous Planning Techniques to Coordinate Multiple Rovers," in *IJCAI Workshop on Scheduling and Planning*, Stockholm, Sweden, August 1999.

[34] A. B. Fiona McNeill and M. Schorlemmer, "Dynamic Ontology Refinement," in *Proceedings of the ICAPS-03 Workshop on Plan Execution*, 2003.

[35] F. Fisher, D. Gaines, T. Estlin, S. Schaffer, and C. Chouinard, "CLEaR: A Framework for Balancing Deliberation and Reactive Control," in *Artificial Intelligence Planning and Scheduling (AIPS 2002) On-line Planning and Scheduling workshop*, 2002.

[36] F. Fisher, T. Estlin, D. Gaines, C. Chouinard, and S. Schaffer, "Onboard Decision Making for Rover Autonomy," in *CSMISS IT Symposium and Missions Track*, 2002.

[37] J. Frank and A. Jonsson, "Constraint Based Attribute and Interval Planning (in press)," *Constraints*, vol. 8, no. 4, pp. 339–364, 2003.

[38] K. Golden, W. Pang, R. Nemani, and P. Votava, "Automating the processing of Earth observation data," in *Proceedings of the 7th International Symposium on Artificial Intelligence, Robotics and Automation for Space (i-SAIRAS)*, 2003.

[39] E. A. Hansen and R. Zhou, "Synthesis of Hierarchical Finite-State Controllers for Hansen-ICAPS-03POMDPs," in *Proceedings ot the Thirteenth International Conference on Automated Planning and Scheduling (ICAPS) 2003*, 2003.

[40] E. K. Jeremy Frank and, "SOFIA's Choice: Scheduling Observations for an Airborne Observatory," in *Proceedings ot the Thirteenth International Conference on Automated Planning and Scheduling (ICAPS) 2003*, 2003.

[41] B. John, A. Vera, M. Matessa, M. Freed, and r. Remington, "Automating CPM-GOMS," in *Association for Computing Machinery: Proceedings of CHI*, 2002.

[42] A. K. Jónsson and J. Frank. *A framework for dynamic constraint reasoning using procedural constraints*, 2000.

[43] A. K. Jónsson, P. H. Morris, N. Muscettola, K. Rajan, and B. D. Smith, "Planning in Interplanetary Space: Theory and Practice," in *Artificial Intelligence Planning Systems*, pp. 177–186, 2000.

[44] L. P. Kaelbling, T. Lane, and G. Theocharous. *Multi-Resolution Planning in Large Uncertain Environments*. http://is.arc.nasa.gov/AR/slides/HierPln02.pdf, September 2002.

[45] G. Karsai, G. Biswas, S. Narasimhan, T. Szemethy, G. Peceli, G. Simon, and T. Kovacshazy, *Towards Fault-Adaptive Control of Complex Dynamic Systems*, pp. 347–368, Wiley, 2003.

[46] L. Khatib, J. Frank, D. Smith, R. Morris, and J. Dungan, "Interleaved Observation Execution and Rescheduling on Earth Observing Systems," in *Proceedings of the ICAPS-03 Workshop on Plan Execution*, 2003.

[47] P. Kim, B. C. Williams, and M. Abramson, "Executing Reactive, Model-based Programs through Graph-based Temporal Planning," in *Proceedings of the International Joint Conference on Artificial Intelligence*, 2001. Discusses RMPL and Kirk, which does graph search on TPNs.

[48] R. Knight, F. Fisher, T. Estlin, B. Engelhardt, and S. Chien, "Balancing Deliberation and Reaction and Planning and Execution for Space Robotic Applications," in *International Conference on Intelligent Robots and Systems (IROS 2001)*, 2001.

[49] R. E. Korf, "Optimal Rectangle Packing: Initial Results," in *Proceedings ot the Thirteenth International Conference on Automated Planning and Scheduling (ICAPS) 2003*, 2003.

[50] S. Lemai and F. Ingrand, "Interleaving Temporal Planning and Execution: IxTeT-eXeC," in *Proceedings of the ICAPS-03 Workshop on Plan Execution*, 2003.

[51] N. Leveson, "Intent Specifications: An Approach to Building Human-Centered Specifications," *IEEE Trans. on Software Engineering*, vol. SE-26, no. 1, , January 2000.

[52] M. Matessa, A. Vera, B. E. John, R. Remington, and M. Freed, "Reusable templates in human performance modeling," in *roceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society)*, 2002.

[53] N. Meuleau and D. E. Smith, "Optimal Limited Contingency Planning," in *ICAPS-03 Workshoop on Planning Under Uncertainty*, 2003.

[54] N. Muscettola, "Computing the Envelope for Stepwise-Constant Resource Allocation," in *Constraint Programming*, pp. 139–154, 2002.

[55] N. Muscettola, G. A. Dorais, C. Fry, R. Levinson, and C. Plaunt. *IDEA: Planning at the Core of Autonomous Reactive Agents*, 2001. Seems to have been at some AAAI event but unknown which.

[56] N. Muscettola, P. P. Nayak, B. Pell, and B. C. Williams, "Remote Agent: To Boldly Go Where No AI System Has Gone Before," *Artificial Intelligence*, vol. 103, no. 1-2, pp. 5–47, 1998.

[57] I. Nesnas, R. Volpe, T. Estlin, H. Das, R. Petras, and D. Mutz, "Toward Developing Reusable Software Components for Robotic Applications," in *International Conference on Intelligent Robots and Systems (IROS 2001)*, 2001.

[58] E. C. Ong, , and N. G. Leveson, "Fault Protection in a Component-Based Spacecraft Architecture," in *Proceedings of the International Conference on Space Mission Challenges for Information Technology*, July 2003.

[59] J. Peterson, "Information Technology Assessment Study," Technical Report JPL-02-03, Jet Propulsion Laboratory, 2002.

[60] J. Peterson, "Information Technology Assessment Study: Executive Summary," Technical Report JPL-02-02, Jet Propulsion Laboratory, 2002.

[61] G. Rabideau, R. Knight, S. Chien, A. Fukunaga, and A. Govindjee, "Iterative Repair Planning for Spacecraft Operations in the ASPEN System," in *International Symposium on Artificial Intelligence Robotics and Automation in Space*, Noordwijk, Netherlands, June 1999.

[62] R. Remington, B. John, A. Vera, M. Matessa, M. Freed, M. Dalal, R. Harris, and E. Dalhman, "Apex/CPM-GOMS: Modeling human performance in applied HCI domains (tutorial materials," in *Twenty-Fourth Annual Conference of the Cognitive Science Society*, 2002.

[63] R. Sherwood, S. Chien, R. Castano, and G. Rabideau, "Autonomous Operations through Onboard Artificial Intelligence," in *Proceedings of the SpaceOps 2002 Conference*, 2002.

[64] R. Sherwood, S. Chien, R. Castano, and G. Rabideau, "Autonomous Planning and Scheduling on the TechSat 21 Mission," in *Proceedings of the Fifteenth Australian Joint Conference on Artificial Intelligence*, 2002.

[65] R. Sherwood, S. Chien, R. Castano, and G. Rabideau, "Spacecraft Autonomy using Onboard Processing for a SAR Constellation Mission," in *Proceedings of the First International Workshop on Future Intelligent Earth Observing Satellites (FIEOS) at the (ISPRS) Commission 1 Symposium*, 2002.

[66] R. Sherwood, A. Mishkin, S. Chien, T. Estlin, P. Backes, B. Cooper, G. Rabideau, and B. Engelhardt, "An Integrated Planning and Scheduling Prototype for Automated Mars Rover Command Generation," in *Proceedings of the 2001 European Conference on Planning (ECP2001)*, 2001.

[67] R. Sherwood, A. Mishkin, T. Estlin, S. Chien, P. Backes, J. Norris, B. Cooper, S. Maxwell, and G. Rabideau, "Autonomously Generating Operations Sequences For a Mars Rover using AI-based Planning," in *Proceedings of the 2001 Intelligent Robots and Systems Conference*, 2001.

[68] R. Sherwood, S. Chien, D. Tran, B. Cichy, R. Castano, A. Davies, and G. Rabideau, "Next Generation Autonomous Operations on a Current Generation Satellite," in *Proceedings of Reducing the Cost of Spacecraft Ground Systems and Operations 2003*, Pasadena, CA, July 2003.

[69] M. Soutchanski, "High-level Robot Programming and Program Execution," in *Proceedings of the ICAPS-03 Workshop on Plan Execution*, 2003.

[70] R. Trejo, J. Galloway, C. Sachar, V. Kreinovich, C. Baral, and L. Tuan, "From planning to searching for the shortest plan: an optimal translation," *International journal of uncertainty, fuzziness and knowledge-based systems*, vol. 9, no. 6, pp. 827–837, 2001.

[71] M. van den Briel, R. Sanchez, M. B. Do, and S. Kambhampati, "Effective Approaches for Partial Satisfaction (Over-Subscription) Planning," in *Proc. National Conf. on Artificial Intelligence*, 2004.

[72] R. Volpe, I. Nesnas, T. Estlin, D. Mutz, R. Petras, and H. Das, "The CLARaTY Architecture for Robotic Autonomy," in *Proceedings of the 2001 IEEE Aerospace Conference*, Big Sky, Montana, March 2001.

[73] K. A. Weiss, E. C. Ong, , and N. G. Leveson, "Reusable Specification Components for Model-Driven Development," in *Proceedings of the International Conference on System Engineering (INCOSE '03)*, July 2003.

[74] B. C. Williams, M. Ingham, S. H. Chung, and P. H. Elliott., "Model-based Programming of Intelligent Embedded Systems and Robotic Space Explorers," *Proceedings of the IEEE*, vol. 9, no. 1, pp. 212–237, January 2003. Special Issue on Modeling and Design of Embedded Software.

[75] B. C. Williams and P. P. Nayak., "A Reactive Planner for a Model-based Executive," in *Proceedings of the International Joint Conference on Artificial Intelligence*, 1997.

[76] H. L. Younes, D. J. Musliner, and R. G. Simmons, "A Framework for Planning in Continuous-time Stochastic Domains," in *Proceedings ot the Thirteenth International Conference on Automated Planning and Scheduling (ICAPS) 2003*, 2003.

[77] S. Zilberstein, R. Washington, D. S. Bernstein, and A.-I. Mouaddib, "Decision-Theoretic Control of Planetary Rovers," in *Proceedings of the Dagstuhl Workshop on Plan-based Control of Robotic Agents*, Wadern, Germany, October 2001.

[78] M. Zweben, B. Daun, E. Davis, and M. Deale, *Scheduling and Rescheduling with Iterative Repair*, Morgan Kaufman, 1994.

Appendix A

Raw Mission Data

This Appendix contains the mission data referenced and summarized in the body of the report.

First, here are the missions ommitted from the classes reported in Section 2, with their criterion rankings:

	Criteria										
Missions	Autonomy	Resources	Overscribed	Coordinated	System dynamics	Multi-step	Unpredictability	Distributed	Real-time	Model drift	Replanning
AMS	1	3	3	1	1	1	1	1	1	3	3
ASPERA-3	3	1	1	1	1	1	1	1	3	1	1
CINDI	1	3	3	1	1	3	1	3	1	3	3
Geospace	1	1	3	1	1	1	3	3	1	3	1
Gravity Probe B	1	1	1	1	1	1	1	1	3	1	1
Kepler3		1	1	1	1	1	1	1	1	3	1
Lunar-A	1	1	1	1	1	1	1	1	3	1	1
Mars '05 Orbiter	3	3	3	1	1	1	1	1	3	1	1
Mars Express	3	3	3	1	1	3	3	3	3	1	1
Planck	1	1	1	1	1	1	1	1	1	3	1
Rosetta	3	3	1	1	3	3	1	1	3	3	1
SDO	3	3	3	1	1	1	1	1	3	3	3
SELENE	3	1	1	1	1	1	1	1	3	3	1
SMART-1	1	1	1	1	1	1	1	1	3	3	1
Solar B	1	3	3	1	1	1	1	1	1	3	3
Solar Probe	3	1	3	1	1	1	1	1	3	1	3
SPIDR	3	3	3	1	1	1	1	1	1	3	3
STERO	3	3	3	3	1	1	3	3	1	3	1
Swift	3	3	3	1	1	1	1	1	1	3	3
Titan Tethered Aerobot	3	3	3	3	3	3	3	3	3	3	1
TWINS	1	3	3	3	1	1	1	1	1	3	3
Venus Express	3	1	1	1	1	1	1	1	3	3	1

Here is the full set of missions, with rankings and descriptions.

	Date	Missions	Criteria: complex mission operations										
			Autonomy -- a spacecraft or rover is far away and out of touch or with small throughput and/or long transmission lag	Crowded, resource-bounded schedules	Over-subscribed schedules with many stakeholders	Coordinated multi-platform operations	Complex system dynamics	Complicated, multi-step operations	Uncertain execution	Distributed problem solving (within the mission)	Real-time constraints on response	Model drift	Costs on replanning
		examples	Europa cryobot, Venus landers, Mars rovers like MER	Shuttle ops and refurbishment, Station ops.	Hubble, SIRTF, SOFIA	Satellite constellations for earth observation, space-based interferometry.	Orbital mechanics, resource limits (fuel,battery, reaction mass), biological systems (long-term life-support, e.g. for manned Mars exploration).	Cassini operations, or Galileo	Complex environments, information-gathering actions	NB: distinguish distribution as a requirement (driven by political separation of authority, or limits on comm. bandwidth, e.g.), from distribution as a solution strategy (say I decide to use distributed agents as a solution, doesn't mean that's the only way to do it)	Limits on local processing power compared to what's needed (Driven eitherr by liminted processing hardware or a very dynamic environment)	Equipment Wear, seasonal variations, etc. due to expanded mission lifetime and/or a harsh environment	1) plan for extended period, 2) plan published/synchronized with oters, and 3) environment is unpreductable
Manned Near	TBA	Shuttle flight No. 114 - Atlantis flight No. 27	1	9	3	3	1	3	3	3	3	1	3
Manned Near		1. The Multi-Purpose Logistics Module, or MPLM, carries supplies and equipment to the station. 2. Delivers the External Stowage Platform to the station. 3. Remove and replace Control Moment Gyro	Piloted missions	Multiple tasks in limited time	Multiple stakeholders w/mixed objectives, but well defined	Docking w/ISS	Assuming fixed orbit	Assembly of ISS	Human execution	Assembly of ISS	Orbital mission	None significant within this limited mission	
Manned Near	TBA	Shuttle flight No. 115 - Endeavour flight No. 20	1	9	3	3	1	3	3	3	3	1	3
Manned Near		1. Delivers the second port truss segment, the P3/P4 Truss, to attach to the first port truss segment, the P1 Truss. 2. Deploys solar array set 2A and 4A. 3. Activates and checks out Solar Alpha Rotary Joint (SARJ). 4. Deploys P4 Truss radiator	Piloted missions	Multiple tasks in limited time	Multiple stakeholders w/mixed objectives, but well defined	Docking w/ISS	Assuming fixed orbit	Assembly of ISS	Human execution	Assembly of ISS	Orbital mission	None significant within this limited mission	
Manned Near	TBA	Shuttle flight No. 116 - Atlantis flight No. 28	1	9	3	3	1	3	3	3	3	1	3
Manned Near		1. Delivers third port truss segment, the P5 Truss, to attach to second port truss segment, the P3/P4 Truss. 2. Deactivates and retracts P6 Truss Channel 4B (port-side) solar array. 3. Reconfigures station power from 2A and 4A solar arrays. 4. Delivers the Expedition Eight crew to the station and returns the Expedition Seven crew to Earth.	Piloted missions	Multiple tasks in limited time	Multiple stakeholders w/mixed objectives, but well defined	Docking w/ISS	Assuming fixed orbit	Assembly of ISS	Human execution	Assembly of ISS	Orbital mission	None significant within this limited mission	
Manned Near	TBA	Shuttle flight No. 117 - Endeavour flight No. 21	1	9	3	3	1	3	3	3	3	1	3
Manned Near		1. The second starboard truss segment, the S3/S4 Truss, is attached to the first starboard truss, the S1, along with a third set of solar arrays. 2. Four external attachment sites for truss-mounted exterior experiments and research are delivered. 3. Activate and check out S4 Truss Solar Alpha Rotary Joint (SARJ). 4. Channel 1A and 3A solar arrays are deployed and station power supply reconfigured. 5. P6 Truss Channel 2B (starboard) solar array is retracted.	Piloted missions	Multiple tasks in limited time	Multiple stakeholders w/mixed objectives, but well defined	Docking w/ISS	Assuming fixed orbit	Assembly of ISS	Human execution	Assembly of ISS	Orbital mission	None significant within this limited mission	
Manned Near	TBA	Shuttle flight No. 118	1	9	3	3	1	3	3	3	3	1	3
Manned Near		1. The third starboard truss segment, the ITS S5 Truss, is attached to the station 2. A SPACEHAB Single Cargo Module delivers supplies and equipment to the station	Piloted missions	Multiple tasks in limited time	Multiple stakeholders w/mixed objectives, but well defined	Docking w/ISS	Assuming fixed orbit	Assembly of ISS	Human execution	Assembly of ISS	Orbital mission	None significant within this limited mission	
Manned Near	TBA	Shuttle flight No. 119 - Atlantis flight No. 29	1	9	3	3	1	3	3	3	3	1	3
Manned Near		1. Fourth and final set of U.S. solar arrays delivered along with fourth starboard truss segment, the S6 Truss. 2. Relocate P6 Truss from atop Z1 Truss to final assembly location attached to P5 Truss (becomes final port-side truss segment). 3. Redeploy and activate P6 Truss Channel 2B and 4B solar arrays.	Piloted missions	Multiple tasks in limited time	Multiple stakeholders w/mixed objectives, but well defined	Docking w/ISS	Assuming fixed orbit	Assembly of ISS	Human execution	Assembly of ISS	Orbital mission	None significant within this limited mission	
Manned Near	TBA	Shuttle flight No. 120	1	9	9	3	1	3	3	3	3	1	3
Manned Near		1. The second of three station connecting modules, Node 2, attaches to end of U.S. Lab and provides attach locations for the Japanese laboratory, European laboratory, the Centrifuge Accomodation Module and later Multipurpose Logistics Modules. 2. Primary docking location for the shuttle will be a pressurized mating adapter attached to Node 2. 3. ISS U.S. Core complete.	Piloted missions	Multiple tasks in limited time	Multiple stakeholders w/mixed objectives, w multi nationals, but well defined	Docking w/ISS	Assuming fixed orbit	Assembly of ISS	Human execution	Assembly of ISS	Orbital mission	None significant within this limited mission	
Manned-Far		Manned Mars	3	9	9	3	9	3	3	3	9	9	9
Manned-Far		Manned exploration of Mars and return with scientific data and samples	High degree of autonomy of support systems	Multiple tasks in limited time	Multiple stakeholders w/mixed objectives, w multi nationals, but well defined	Main vehicle, Lander, Rover	Human life support for up to 2 years, perservation of samples	Travel, Orbit, land, explore, return	Human execution	Likely, due to multiple tasks and goals	Not for comm, but for remote operations which need to proceed W/OUT communications	Probably, given mission length	Extended mission
IN DEVELOPMENT FOLLOWS													
Earth Obs (payload)	Late 2003	CINDI	1	3	3	1	1	3	1	3	1	3	3
Earth Obs (payload)		CINDI will provide measurements of the neutral atmosphere wind velocity and the charged particle drifts in the equatorial upper atmosphere at altitudes from 400 to 700 km. CINDI seeks to discover how the neutral gas motions. become part of the payload for the Communication and Navigation Outage Forecast System		assume multiple experiments	Assume multiple stakeholders					shared between on-board automation and mission control		assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Earth Obs.		TWINS	1	3	3	3	1	1	1	1	1	3	3

Earth Obs.		Image the Earth's magnetosphere in energetic neutral atoms from two widely-spaced, high-altitude spacecraft		less then Hubble	Less then Hubble							assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Orbiter	2005	Mars '05 Orbiter	3	3	3	1	1	1	1	1	3	1	1
Orbiter		measure thousands of Martian landscapes at 20- to 30-centimeter (8- to 12-inch) resolution	Time lag	less then Hubble	Less then Hubble								
Orbiter	2003	Mars Express	3	3	3	1	1	3	3	3	3	1	1
Orbiter		explore Mars atmosphere and surface from Ploar Orbit	Time lag	less then Hubble	Less then Hubble			well understood operations	Lander	lander & orbiter			
Orbiter	2004	MESSENGER	3	3	3	1	3	9	1	1	3	3	1
Orbiter		MErcury Surface, Space ENvironment, GEochemistry and Ranging mission to orbit Mercury following two flybys of that planet. The orbital phase will use the flyby data as an initial guide to perform a focused scientific investigation	multiple flybys 8+ year mission mostly autonomous	less then Hubble	Less then Hubble	simple one platform	planning and orbital determination decoupled	decision for study based on info gathered in fyl-bys				6 year mission life	
Orbiter	2006	New Horizons (Pluto)	9	1	1	1	9	3	1	1	3	9	1
Orbiter		spacecraft would use a remote sensing package that includes imaging instruments and a radio science investigation, as well as spectroscopic and other experiments, to characterize the global geology and morphology of Pluto and its moon Charon, map their surface composition and characterize Pluto's neutral atmosphere and its escape rate	20+ year mission, autonomous				Planning and orbital determination tightly tied.	two step, Pluto then Kupier Belt				assume drift on sophisticated sensors	
Orbiter	Nov. 2005	STERO	3	3	3	3	1	1	3	3	1	3	1
Orbiter		Study solar ejections - Two identical spacecraft with identical instrument complements - Heliocentric orbit 20-30 degrees from Sun-Earth line		less then Hubble	Less then Hubble	2 craft	coordinate 2 crafts thru Heliocentric orbit 20-30 degrees from Sun-Earth line		Response to unexpected events		RT alerts back to Earth	assume drift on sophisticated sensors	
Orbiter	2003	SMART-1	1	1	1	1	1	1	1	1	3	3	1
Orbiter		To flight test the new Solar Electric Propulsion technology and to thoroughly investigate and map the Moon. SMART-1 will also search for ice on the lunar poles	Simple mission with comm capability										
Orbiter	2005	SELENE	3	1	1	1	1	1	1	1	3	3	1
Orbiter		after a year of mapping, the orbiter's propulsion unit will separate and land on the Moon to send a signal from the lunar surface to provide another reference point											
Orbiter	2005	Venus Express	3	1	1	1	1	1	1	1	3	3	1
Orbiter		study the Venus atmosphere in great detail										assume drift on sophisticated sensors	
Orbiter	2011	Jupiter Icy Moons Orbiter	9	1	1	1	9	9	3	1	3	9	3
Orbiter		Jupiter Icy Moons Orbiter would make detailed studies of the 3 moons' makeup, history and potential for sustaining life	long mission,				establish 3 separarte orbits on the mission		Extended operations at a distance, new science objectives from obs.				
Orbiter (Payload)	2003	ASPERA-3	3	1	1	1	1	1	1	1	3	1	1
Orbiter (Payload)		study the interaction between the solar wind and the Martian atmosphere. Travel on Mars Express, polar orbit											
Rover	2004	Deep Impact	3	3	9	1	1	1	3	3	3	1	1
Rover		A radical mission to excavate the interior of a comet. A camera and infrared spectrometer on the spacecraft, along with ground-based observatories, will study the resulting icy debris blasted off the comet, as well as the pristine interior material exposed by the impact.	Simple crash landing	not as crowded as Hubble	many analyst and experimenters		Impact into with out going thru the comet		surface of comet is not well understood	coordinate between grnd and craft	RT data collection		
Rover	?	Gravity Probe B	1	1	1	1	1	1	1	1	3	1	1
Rover		Gravity Probe B telescope is the key reference for the Earth-orbiting relativity experiment. The GPB gyro spin directions are referenced to the line-of-sight direction to a guide star, as seen through the telescope. As the gyroscopes undergo relativistic drift, caused by the distortion of space-time near our massive planet, their spin axis orientation is compared to this reference direction.		single mission			needs a highly stable platform for telescope						
Rover	2003	Mars Exploration Rovers	3	3	3	3	3	3	9	9	3	9	3
Rover		search for and characterize a wide range of rocks and soils that hold clues to past water activity on Mars	Has comm thru-out space flight Human in the loop travel on surface (could be automated)	Multi-experiments, but less then Hubble	Multiple stakeholders	Two Rovers acting independently	space flight has 3 correction points, fairly well understood	landing, deploying Rover, several camera and sensor systems, lift off, mostly human in the loop but could be more automated	terrain navigation is unknown and un-predictable	2 rovers			
Rover	2004	Rosetta	3	3	1	1	3	3	1	1	3	3	1
Rover		Note: mission postponed due to missed window First effort to land on a comet. study the origin of comets, the relationship between cometary and interstellar material and its implications to the origin of our solar system	long mission, lander will need to deploy and gather info				Manuvering close to comet will be difficult	deploy lander	Fixed set of (initial) observations	coordinate between craft & lander		assume drift on sophisticated sensors	
Rover	2004, 2006	Space Tech 6	3	1	1	1	9	9	1	1	3	1	1
Rover		Autonomous Sciencecraft Experiment (Sciencecraft) and Inertial Stellar Compass (Compass). Sciencraft will enable a spacecraft to decide what science observations to make, and then process and return data——all on its own. Compass will enable a spacecraft to continuously sense its position and recover after a temporary malfunction or power loss	developing autonomous technology, thus under supervision				Technology to be aware of position	Technology to allow autonomic behaviour					

Rover	2003	Lunar-A	1	1	1	1	1	1	1	1	3	1	1
Rover		smash through the lunar surface and study's the Moon's interior with seismometers and heat-flow probes											
Rover	2009	Mars NetLander	9	1	3	9	3	3	3	3	3	3	1
Rover		An orbiter spacecraft will send four landers to the surface of Mars prior to entering Mars orbit. The four landers will establish the first network of science stations on Mars. The stations and orbiter will study the surface, subsurface, interior, atmosphere, ionospheric structure and the size and shape of the planet			many stakeholders	4 netlanders and 1 orbiter	Orbits and landing only		Landers, but no movement.	multi-vehicles and coordinated landing			
Rover	2011	Mars Scout 2	3	3	1	1	3	3	3	1	9	1	3
Rover		Phoenix lander. Multiple instrument packages, multiple science teams, multi-country											
Rover	2013	Titan Aerobot - Multisite (TAM) mission	9	3	3	3	3	9	9	3	3	3	1
Rover		For a surface-oriented mission, a balloon using argon as a reversible fluid is the preferred approach and would permit visits to hundreds of sites well distributed over the surface of the satellite. This vehicle could make numerous visits to the surface with duration of hours to days before rising to altitude and drifting to another location. Unlike Venus, temperature variations in the Titan atmosphere are not large and thermal control considerations do not limit the duration of surface stay time.	assume modify search based on sensor data	limited plant	assume multiple stakeholders, but not commercialy driven	Assume communication repeater orbiter working in conjunction with rover		long flight, difficult to deploy Aerobot, Aerobot's course will depend on discovery during exploration	Mobile in a dynamic environment	multi-vehicles	assume swarm behaviour		
Rover	2013	Titan Aerobot -Singlesite (TAS) Mission	3	3	3	3	3	3	3	3	3	3	1
Rover		For a focus on the atmosphere, a higher altitude capability is needed and this would utilize a superpressure balloon. The vehicle would descend to the lower atmosphere of Titan for inflation and then float at an altitude of 115 to 125 km to implement its atmospheric mission. The vehicle would then vent gas descend to near the surface and skim the terrain using a guide rope. It would conduct observations at one surface site for a distance of several tens of kilometers.	assume modify search based on sensor data	limited energy resources and many experiments	assume multiple stakeholders, but not commercialy driven	Assume communication repeater orbiter working in conjunction with rover		long flight, difficult to deploy Aerobot, Aerobot's course will depend on discovery during exploration	Tethered, in a dynamic environment	multi-vehicles	assume swarm behaviour		
Sample & return	2011	Mars Sample Return Lander	3	1	1	1	3	9	9	3	9	3	1
Sample & return		first mission to return samples of Martian rock and soil to Earth	autonomous sample collection					return vehicle	sample gathering, packaging,& return	redocking manuever			
Sample & return	2009	Mars Science Laboratory	9	3	3	1	3	9	9	3	9	3	3
Sample & return		Undefined - develop and to launch a roving long-range, long duration science laboratory that will be a major leap in surface measurements and pave the way for a future sample return mission. NASA is studying options to launch this mobile science laboratory mission as early as 2007. This capability will also demonstrate the technology for "smart landers" with accurate landing and hazard avoidance	smart landers and hazard avoidance	scheduling for Lab experiments	many stakeholder	unknown if coordinated with flying Scouts	sample preservation systems (including bio)	difficult to send samples back/coordination	high degree of autonomy in manuvering and sample gathering	coordinated behaviour	search based on data	assume drift on sophisticated sensors	plans published
Sample & return	2003	MUSES-C	9	1	1	1	1	3	3	1	3	3	1
Sample & return		The spacecraft and sample return vehicle will travel to asteroid (25413) 1998SF36. The trip will take about 17 months. MUSES-C will initially study the asteroid from a distance of 20 km (12.4 miles) and then move closer for a series of soft landings	48 month rnd trip Autonmous nav.	simple one ship	single stakeholder	simple one platform	3 sample collection of asteroid	series of soft landings	ability to retry is sample collection fails				
Unv. Obs	Jan/Feb 2005	Astro-E2	1	9	9	1	1	1	1	1	1	3	3
Unv. Obs		Astro-E2 will be Japan's fifth X-ray astronomy mission, and is being developed at the Institute of Space and Astronautical Science (ISAS) in collaboration with U.S. (NASA/GSFC, MIT) and Japanese institutions. Astro-E2 will cover the energy range 0.4 - 700 keV with the three instruments, X-ray micro-calorimeter (X-ray Spectrometer; XRS), X-ray CCDs (X-ray Imaging Spectrometer; XIS), and the hard X-ray detector (HXD)		assume multiple experiments	Assume multiple stakeholders		Cryo-cooling, limited life time, optimize use and graceful degradation					assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Unv. Obs	2007	Herschel	1	9	9	1	3	3	1	1	1	3	3
Unv. Obs		Will use infrared astronomy to solve the mystery of how stars and galaxys and stars were born		Hubble like scheduling	Hubble like		Optimizing cryogenic use.	double mission launch that separates in space				assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Unv. Obs	Late 2004	Hubble SM4	1	9	9	1	1	3	3	9	1	1	3
Unv. Obs		Service mission to upgrade to Wide Field of Vision Camera 3, and Cosmic Origins Spectrograph		many experiments	many stakeholders			need to deliver and install new equipment in Hubble	Human execution			assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Unv. Obs	2007	Planck	1	1	1	1	1	1	1	1	1	3	1

Unv. Obs		CMB measurement		Dedicated mission			unknown	unknown				assume drift on sophisticated sensors	
Unv. Obs	August. 2003	SIRTF	3	3	9	1	1	3	1	1	1	3	3
Unv. Obs		the Space Infrared Telescope Facility, SIRTF will obtain images and spectra by detecting the infrared energy, or heat, radiated by objects in space between wavelengths of 3 and 180 microns (1 micron is one-millionth of a meter). Most of this infrared radiation is blocked by the Earth's atmosphere and cannot be observed from the ground		less then Hubble			use Earth trailing orbit to provide "shade" to extend life of Cryo-cooling					assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Unv. Obs	2004	SOFIA	1	9	9	1	9	1	3	3	1	3	3
Unv. Obs		SOFIA is an airborne (747) observatory that will study the universe in the infrared spectrum							Terrestrial weather			assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Unv. Obs	Sep. 2006	Solar B	1	3	3	1	1	1	1	1	1	3	3
Unv. Obs		Single spacecraft, sun-synchronous polar Earth orbit for full time solar observation . It will determind the solar orgins of space weather and global change by studying stellar magnetic fields.		less then Hubble	Less then Hubble							assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Unv. Obs		Space Tech 5	9	1	1	9	9	3	3	9	9	1	1
Unv. Obs		ST5 is currently designing and building miniaturized components and technologies that can be integrated into a small satellite. Known as nanosats or small-sats, each satellite will weigh approximately 47 pounds	Able to operate as constellation	depends on mission	depends on mission	Swarm capable, shared mission	small size, limited energy resources	multi-step deployment	new technology for constellation behaviour	swarm behaviour	swarm behaviour		
Unv. Obs	Dec. 2003	Swift	3	3	3	1	1	1	1	1	1	3	3
Unv. Obs		Swift is a first-of-its-kind multi-wavelength observatory dedicated to the study of gamma-ray burst (GRB) science		less then Hubble	Less then Hubble				React to events, but no plan!			assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders
Unv. Obs (payload)	2005	AMS	1	3	3	1	1	1	1	1	1	3	3
Unv. Obs (payload)		AMSAn experiment to search in space for dark matter, missing matter & antimatter on the international space station.		Multi-experiments, but less then Hubble								assume drift on sophisticated sensors	long term test and published schedule with multiple stakeholders

		Shuttle Mission Planning and Ops	1	9	9	1	1	9	3	9	1	1	9
		Ground support for shuttle missions											
		Deep Space Mission Planning and Ops	1	3	3	1	1	9	1	3	1	3	3
		Earth Observing Mission Planning and Ops	1	9	9	3	1	3	1	3	1	1	3
		Training Ops	1	9	9	1	1	3	3	3	1	1	9
		Astronaut training, e.g.											
Earth Obs.	1999	AIM	1	1	1	1	1	1	1	1	1	9	3
Earth Obs.		The overall goal is to resolve why Polar Mesospheric Clouds (PMCs) form and why they vary. AIM will measure PMCs and the thermal, chemical and dynamical environment in which they form										Assume drift on sophisticated sensing eqp.	Assume multiple users
Earth Obs.	2009	GEC	3	3	1	9	9	3	3	3	3	3	1
Earth Obs.		consisting of a cluster of 4 satellites, combined with ground-based observations will make systematic multi-point measurements to delineate and bring to closure our understanding of key roles the IT plays in the Sun-Earth connection	Capable of changing orbit autonomously based on data			4 space crafts	change flight based on data	coordinated flight	flight based on data	shared mission across 4 Sats & gnd station	Varying orbit based on measurements	Assume drift on sophisticated sensing eqp.	single mission
Earth Obs.	2008 - 2010	Geospace	1	1	3	1	1	1	3	3	1	3	1
Earth Obs.		Understanding and characterizing the effect of solar variability on those geospace phenomena that most affect life and society											
Earth Obs.	N/A	Mag Constellation	9	3	1	9	9	9	1	9	3	3	1
Earth Obs.		MC will answer the fundamental question: "How does the dynamic magnetotail store, transport, and release matter and energy?"MC Mission Description: - A constellation of 50 small satellites distributed in 3x7 Re to 3x40 Re, low inclination, nested orbits. - "Nearest neighbor" average spacing 1.0-2.0 RE between satellites, in the domain of the near-Earth plasma sheet	swarm of 50 sats.	assume single mission, but need scheduling for coordinating swarm		swarm of 50 sats.	nano scale has fuel resource issues, controlability issues	deploying the sats in space		multiple nano sats, assume networked	assume RT control to stay in "tail"	Assume drift on sophisticated sensing eqp.	
Earth Obs.	2009	Mag Multiscale	3	3	1	9	9	1	1	9	3	3	1
Earth Obs.		Broad regions of the magnetosphere are connected by fundamental processes operating in thin boundary layers. Processes of vastly different scale sizes can interact strongly. Understanding these fundamental processes requires multipoint measurements that uniquely separate temporal and three-dimensional spatial variations. Magnetospheric Multiscale gives us this essential capability.	assume modifying mission based on data	single mission			tracking magnetic boundaries			multiple nano sats, assume networked	assume RT control to stay in "tail"	Assume drift on sophisticated sensing eqp.	
Orbiter	2006	Dawn	9	1	1	1	3	9	3	1	3	3	1
Orbiter		Delve into the orgins of our solar system through intense study of Ceres and Vesta, two minot planeTs that reside in the cast asteroid belt between Mars and Jupiter. The mission marks the first time a spacecraft will orbit two planetary bodies on a single voyage.	10+ yr mission					multi- step mission					
Orbiter	2007	SCIM	3	1	1	1	1	3	3	1	9	1	1

Orbiter		SCIM will gather invaluable Martian dust and atmosphere samples without descending to the surface or even entering Mars orbit. Instead, SCIM will depart from Earth and fly through the Martian atmosphere at high speed, gather its samples, and return directly to Earth											
Orbiter	?	SDO	3	3	3	1	1	1	1	1	3	3	3
Orbiter		Understand the solar cycle, identify the role of the magnetic field in delivering energy to the solar atomosphere and its many layers, study how the outer regions of the Sun's atmosphere evolve over time, and monitor the radiation levels of solar output											
Orbiter	?	SIM	3	3	9	9	3	3	3	3	9	3	3
Orbiter		SIM will be an optical interferometer operating in an Earth-trailing solar orbit	long duration but in comm distance. Able to handle targets of opprotunity	multiple experiments (but less then Hubble)	multiple stakeholder		highly accurate orbit/positioning	shuttle to earth orbit to solar orbit	able to handle "targets of opprotunity"			Assume drift on sophisticated sensing eqp.	assume published schedule with multiple stakeholder
Orbiter	?	Solar Probe	3	1	3	1	1	1	1	1	3	1	3
Orbiter		Determine the acceleration processes and find the source regions of the fast and slow solar wind at maximum and minimum solar activity; Locate the source and trace the flow of energy that heatsthe corona; Construct the three-dimensional density configurationfrom pole to pole, and determine the subsurface flow pattern, the structure of the polar magnetic field and itsrelationship with the overlying corona; and Identify the acceleration mechanisms and locate thesource regions of energetic particles, and determine therole of plasma turbulence in the production of solar wind and energetic particles.	Time lag										Due to solar proximity, this mission has a larger cost of replanning if an error is made (and, for example, the orbiter travels too close to the sun).
Rover	2010 - 2015	SAGE	9	3	3	3	1	3	9	1	9	3	1
Rover		Venus orbiter and lander	Probe, limited lifetime, very limited bandwidth	Craft would have a limited time on surface because of Harsh environment		Orbiter and probe	Orbital insertion		Operating in the context of planned but not predictable failure			Short term mission, but some gradual degradation possible anyway.	
Rover	2008	ARES Science	9	3	1	1	3	9	9	1	9	3	3
Rover		ARES will use a OAV to gather & return critical science data across up to 680 km of diverse terrain in one of the most scientifically intriguing regions of Mars: the Southern Highlands.	UAV flight to gather data	optimize data gathering w/limited resource			UAV has limited fuel resources and large area to cover		Rover in dynamic terrain	coordinate between land rover and orbiter	vary flight path based on data sensed and received from ground and space	Assume drift on sophisticated sensing eqp.	change in flight plan will change ground rovers search.
Rover	2030	Europa Cryobot	9	3	3	3	3	9	9	9	9	3	1
Rover		exobiological exploration of Europa, E3 spacecraft comprised of a science/relay orbiter and a mapper/lander/cryobot/hydrobot. The latter incorporates two robotic probes including a cryobot, which moves through ice by melting it. A hydrobot is a self-propelled underwater vehicle.	Robots will modify course based on exploration, very limited com to orbiter	short battery life and many experiments	multiple stakeholder	Orbiter, cryobot, hydrobot - bot to bot coordination	resource limitation for Bot exploration below ice.	find thin ice, melt thru, deploy hydrobot, explore, report back	Rover in dynamic and unknown "terrain"	coordinate between land rover and orbiter	vary underwater search based on sensor data	Assume drift on sophisticated sensing eqp.	
Univ. Obs	?	Constellation-X	3	9	9	9	9	1	1	9	9	3	3
Univ. Obs		The Constellation-X Mission will place in orbit an array of X-ray telescopes that will work in unison to improve our view of the Universe by a hundredfold.		optimize schedule for tasks	multiple stakeholders	multiple coordinated platforms	multiple coordinated platforms			coordination & closed loop control around matrix of telescopes	require timely coordination across matrix of telescopes	Assume drift on sophisticated sensing eqp.	long term test and published schedule
Univ. Obs	2001	EUSO	1	3	1	1	1	1	1	1	1	3	9
Univ. Obs		Investigate the nature and origin of extreme-energy cosmic rays (EECRs) -the window on the extreme energy universe. ISS-based experiment										Assume drift on sophisticated sensing eqp.	long term test and published schedule
Univ. Obs	2006	GLAST	3	9	9	1	1	3	1	1	1	3	3
Univ. Obs		Gamma-ray Large Area Space Telescope (GLAST) will open this high-energy world to exploration of Supermassive black holes, merging neutron stars, streams of hot gas moving close to the speed of light		many experiments	multiple stakeholders		coordination of ship and pointing of telescope					Assume drift on sophisticated sensing eqp.	long term test and published schedule
Univ. Obs	2011	JWST (NGST)	3	9	9	1	3	3	1	1	1	3	3
Univ. Obs		An orbiting infrared observatory to determine the shape fo the univers, explain galaxy evolution, understand the birth and formation of stars, determin how planetary systems form and interact, determine how the universe built up its present chemical/elemental compsition, and probe the nature and abundance of Dark Matter		Assume multiple experiments	assume multiple stakeholders		coordination of ship and pointing of telescope					Assume drift on sophisticated sensing eqp.	long term test and published schedule
Univ. Obs	Get date	Kepler	3	1	1	1	1	1	1	1	1	3	1
Univ. Obs		The scientific goal of the Kepler Mission is to explore the structure and diversity of planetary systems. The Kepler instrument is a 0.95-meter aperture differential photometer with a 105 deg2 field of view. It continuously and simultaneously monitors brightnesses of 100,000 A-K dwarf (main-sequence) stars brighter than 14th magnitude. The experiment is not biased by preselection of stellar type or single versus multiple star system. The orbit and size of the planets can be calculated from the period and depth of the transit	searches for specific stars	multiple experiments, narrower scope then telescopes like Hubble	multiple stakeholder, narrower scope then telescopes like Hubble							Assume drift on sophisticated sensing eqp.	long term test and published schedule
Univ. Obs	2011	LISA	3	1	1	9	9	1	1	3	9	3	3

Univ. Obs		the first dedicated space-based gravitational wave observatory, LISA will detect gravitational waves generated by binaries within our Galaxy(the Milky Way) and by massive black holes in distant galaxies	searches for specific stars	multiple experiments, narrower scope then telescopes like Hubble	multiple stakeholder, narrower scope then telescopes like Hubble							Assume drift on sophisticated sensing eqp.	long term test and published schedule
Univ. Obs	2005	SPIDR	3	3	3	1	1	1	1	1	1	3	3
Univ. Obs		The Spectroscopy and Photometry of IGM's Diffuse Radiation (SPIDR) will spectrally image a large (~20%) portion of the sky in the 103 – 120 nm and 154 – 156 nm bands. These measurements will be obtained by six single-element imaging spectrographs in combination with a novel observation strategy. SPIDR contains no expendables and has a nominal operational life of three years											
Univ. Obs	2006	THEMIS	9	1	1	9	1	3	3	3	3	3	1
Univ. Obs		THEMIS is a study of the onset of magnetic storms within the tail of the Earth's magnetosphere. THEMIS will fly five microsatellite probes through different regions of the magnetosphere and observe the onset and evolution of storms											
Univ. Obs	2007	WISE	3	9	9	1	1	3	1	3	1	3	3
Univ. Obs		find the most luminous galaxies in the Universe. find the closest stars to the Sun. detect most main belt asteroids larger than 3 km. extend the 2MASS survey into the thermal infrared. enable a wide variety of studies ranging from the evolution of protoplanetary debris disks to the history of star formation in normal galaxies. provide the essential catalog for the James Webb Space Telescope					IRAS/ COBE-type, circular, 500 km, Sun synchronous polar orbit on a Taurus 2210 expendable launch vehicle					Assume drift on sophisticated sensing eqp.	long term test and published schedule
		Space Tech 7	3	1	1	1	3	3	9	1	3	3	3
		NASA's New Millennium Program has selected two organizations to lead the work on sensor and thrust-producing technologies to control a space vehicle's flight path so the payload responds only to gravitational forces. The Disturbance Reduction System technology is scheduled to fly in 2006 as the Space Technology 7 project. Space Technology 7 is designed to test and validate advanced technologies for future use on NASA missions that have never been flown in space.	Likely, due to various test cases, and alternate subsequent test cases based on outcomes.	Possibly, due to various test cases, and alternate subsequent test cases based on outcomes.		Possibly due to various test scenarios.			The crticial component in testing new platforms and technologies	Possible			

Appendix B

Current Research

In this Appendix, we present brief summaries of projects underway at Ames, JPL, and funded by the IS program. Much of this information can be found on the web, at the following URLs:

IS http://is.arc.nasa.gov/

Ames http://www.arc.nasa.gov/

JPL http://www.jpl.nasa.gov/

Agent Development and Control Verification Using Dual Characterizations Arizona State University Chitta Baral (ASU) Vladik Kreinovich (UTEP), Tran Cao Son (NMSU-LC)

An autonomous agent must generate complex plans dealing with incomplete information and with events that have variable duration, delayed and continuous effects, changing resources (fluents), and complex constraints (temporal, spatial, procedural, and hierarchical). This work [70] will use declarative logic programming and domain constraints to generate and verify control programs for autonomous agent planning, scheduling, monitoring, and diagnosis. An English-like language will be used to describe a spacecraft and its abilities, mission goals, environment, observations, and the effects of actions. The language will be used in a planning, scheduling, and plan-verifying (PSV) system component, and will be relatively easy for mission controllers to understand and specify. Domain experts will have progression and regression tools to help construct verifiable plans and control modules, and will develop domain-specific planning constraints to aid the autonomous reasoner. Planning/scheduling, diagnosis/repair, and mission-related domain knowledge will be formulated in logic programs and proved correct. Logic-based verification will be fully automatic, even for causal models and plans that include sensing actions and conditionals. Such techniques will permit rapid development of reliable software agents.

Continual Coherent Team Planning NASA Jet Propulsion Laboratory Tony Barrett (JPL/AI) Milind Tambe (USC/ISI) Bradley Clement (JPL/AI), Hyuckchul Jung (USC)

Future NASA missions will require coordinated, self-reconfiguring teams of autonomous agents (such as a fleet of spacecraft or rovers) acting in dynamic, partially understood environments. This research [4] will explore negotiation strategies for distributed autonomy, enabling closely coordinated spacecraft to manage their local plans while performing team activities. Strategies will be chosen dynamically to satisfy real-time deadlines, with roles negotiated to optimize the team's collective performance. Techniques will be developed and validated through spacecraft cluster simulation using the Spacecraft Control Toolbox (SCT). Potential NASA missions include interferometers and signal-isolating constellations, synthetic apertures, coordinated spectral or spatial coverage, and robotic explorer and outpost teams.

Team Sequence Execution for Cluster Operations NASA Jet Propulsion Laboratory Tony Barrett (JPL/AI) Paolo Pirjanian (JPL/Robotics), Seung Chung (MIT)

Spacecraft clusters must closely coordinate in conditions that are only partially known in advance. An integrated cluster management system will permit control of multiple agents via a single team plan instead of explicit command sequences for each agent. Team members will coordinate via a shared team state within a hierarchical plan, and will use negotiation-based distributed diagnosis to correct any detected problems [4]. Potential NASA missions include interferometers and signal-isolating constellations, synthetic apertures, coordinated spectral or spatial coverage, and robotic explorer and outpost teams.

Analytic Verification and Validation for Space Missions NASA Ames Research Center Guillaume Brat (Kestrel/ARC) Arnaud Venet (Kestrel/ARC), Allen Goldberg (Kestrel/ARC), Klaus Havelund (Kestrel/ARC) Group Lead: Michael Lowry (ARC/IC)

Large software systems with real-time decision capabilty are difficult to develop and validate. Automated formal verification can be applied early in the software development process, catching errors before they become costly to find and fix. The methods must be user-friendly, precise, and scaleable. This task [8] will develop static analysis and runtime analysis tools for automated verificaton and validation of autonomy software. The tools will be benchmarked on Mars rover executives and NASA code, and will be developed for direct use by mission engineers.

Formal Analysis of Human-Automation Interaction NASA Ames Research Center Guillaume Brat (Kestrel/ARC) Willem Visser (RIACS), Everett Palmer (ARC/IHI), Seungjoon Park (RIACS), Oksana Tkachuk (ARC/IC)Group Lead: Michael Lowry (ARC/IC)

One of the main tenets of most current models [Brahms Simulation, Task Learning] is that human behavior is heavily constrained by the structure of the physical and information environment. This task [8, 9] extends formal model-checking from automated software engineering to determine how accurately they can model operator displays and procedures. The goal is to investigate and demonstrate how well formal methods and tools can be used to automate the verification and design of systems with human-machine interactions. Formal verification is no substitute for cognitive models [Task Learning, Model Usability], but rather complements them with automation and formalization of their work.

Livingstonee Diagnostic Agent NASA Ames Research Center Lee Brownston (QSS/ARC) James Kurien (PARC), Pandu Nayak (RIACS), David Smith (ARC/IC), Will Taylor (QSS/ARC) Group Lead: Mark Shirley (ARC/IC)

Software systems for vehicle monitoring, diagnosis, and control can be very complex and difficult to develop. Construction of accurate simulators for software validation is also a bottleneck. The Livingstone [6] fault diagnosis and recovery kernel has proved its value. A new generation called L2 includes temporal trajectory tracking, and is being extended to reason about continuous systems. The improved autonomous monitoring and maintenance of vehicle health may enable better mission science returns with increased safety and smaller operations support crews. Results will be incorporated in many NASA missions and systems, including the ISS Command and Data Handling system. Model-based rapid prototyping tools for high-fidelity simulators are also being developed.

Probabilistic Fault Detection for Hybrid Discrete/Continuous Systems NASA Ames Research Center Richard Dearden (RIACS) Vandi Verma (CMU/RI) Reid Simmons (CMU/RI), Rich Washington (RIACS), Dan Clancy (ARC/IC), Thomas Willeke (QSS/ARC), Frank Hutter (ARC)

A rover or autonomous spacecraft must constantly monitor for faults and unexpected conditions, based on models of the system configuration, environment, and current action. Dangerous state changes must be detected very quickly, but with few false alarms and minimal computational power. Probabilistic models (based on particle filters) can help distinguish sensor noise and expected state changes from conditions that are unexpected and mission threatening. This will permit tight tolerances and close monitoring of critical systems [6, 21]. Global decisions will be handed off to a higher-level reasoner (within the CLARAty autonomy architecture model), which will feed back situation knowledge for adjusting resource estimates and thresholds. Experiments will be conducted with a Marsokhod rover wheel model that has 23 discrete modes (including 14 fault modes), plus four continuous state variables. Field test demonstration of onboard health management will be conducted with an Ames K9 rover. Other applications may include Space Shuttle engine or life support systems.

Spacecraft Mobile Robot NASA Ames Research Center Gregory Dorais (ARC/IC) Yuri Gawdiak (ARC/IC)

A small mobile robot could attend to many housekeeping chores in space. The Personal Satellite Assistant (PSA) [24, 6] is to be a softball-sized astronaut support device that can monitor environmental factors, document routine events, detect failures, and help with maintenance operations. The prototype system uses commercial components and software, plus intelligent systems technology for navigation and autonomous action.

Integrated Resource and Path Planning NASA Jet Propulsion Laboratory Tara Estlin (JPL/AI), Caroline Chouinard (JPL/AI)

Future NASA missions will require smarter rovers that can traverse long distances and perform intelligent onboard decision-making. In particular, rovers must decide whether the potential benefits of a science observation are worth the costs and risks of traveling to the observation site. Typically the decision of what waypoints to visit, and in what order, is made by a high-level activity planner/scheduler. A lower-level path planner then chooses a specific route. This task will improve communication between the modules, resulting in better plans. The high-level planner will be able to ask for (and choose between) plans satisfying various constraints and priorities, and the path planner will have these global directives to help it choose specific routes. Implementing constraint communication between the two planners is a first step toward such flexibility. This work will develop a communication interface between the CASPER resource planner and the Tangent Graph and D* path planning systems, within JPL's CLARAty architecture for multiple rovers [17, 31, 28, 35, 36, 72].

Integrated Planning and Execution NASA Jet Propulsion Laboratory Forest Fisher (JPL/AI) Reid Simmons (CMU/RI), Tara Estlin (JPL/AI), Daniel Gaines (JPL/AI) Steve Schaffer (JPL/AI), Caroline Chouinard (JPL/AI)

Strategic planning systems are not fast enough to reason at a stimulus-response level, and reactive planners are ignorant of long-term or system-wide considerations. This results in fragile plans and frequent delays for replanning. Execution systems can handle a broader range of exceptions in real time if told which constraints can be relaxed. High-level, declarative planners could often benefit from procedural planning techniques, as in planning conditional (if-then) actions and iterative loops or when reasoning about exception conditions and plan changes. Incorporation of procedural concepts could improve communication with reactive planners. This research [17, 31, 28, 35, 36, 72] will consider a hybrid control framework with a unified deliberation and execution layer. Greater communication between the planners will be implemented, with reconsideration of how problems should be decomposed. Both static and dynamic decision policies will be investigated. The final planning and execution system will be demonstrated on a Rocky 8 rover at JPL.

Flight Planning for SOFIA NASA Ames Research Center Jeremy Frank (ARC/IC) David Smith (ARC/IC), Elif Kurklu (QSS/ARC)

The Stratospheric Observatory for Infrared Astronomy (SOFIA) is NASA's next-generation airborne astronomical observatory. It consists of a 747-SP aircraft modified to accommodate a 2.7-meter telescope, plus various other astronomical instruments. Manual flight planning cannot meet the expected demand for observation scheduling. This task [40, 21, 42] will model the flight-planning problem, develop constraint-based search techniques for single-flight planning, and extend the solution to instrument selection for multi-flight planning. Plan search works well for complex problems with many variables and relationships, though it may not always find the mathematically optimal problem solution. The planner is partially funded by SOFIA, and may be integrated with on-going operations. Its approach to planning and scheduling of science observations on a mobile platform may be relevant to traverse planning for other autonomous vehicles.

Intelligent Specification-Centered Test-Case Generation University of Minnesota Mats Heimdahl (UMN/CriSys) Willem Visser (RIACS) Mike Whalen (UMN/CriSys), Sanjai Rayadurgam (UMN)

Validation and verification for autonomy software and other critical control systems is enormously expensive, yet may still miss critical faults. Automatic test-case development offers dramatic reduction in the time and cost. This task [9, 8] will develop automated specification-based and code-based generation of test cases by using a model checker to report input sequences that can lead to forbidden execution states. This approach should generate test suites to user-defined levels of specification test coverage. A variation may provide test cases generated from implemented code. Models of flight-control logic will be used to test scalability. Verification capabilities will be integrated into the NIMBUS requirements engineering environment.

Model Usability Carnegie Mellon University Bonnie John (CMU) Mike Freed (SJSU)

Human factors design depends on accurate models of human task performance, but development of task models is time-consuming and requires special expertise. This task will develop user-friendly extension of the Apex human-system modeling framework for designing efficient and easy-to-use control systems and data-entry interfaces. Apex [41, 52, 62] is descended from the popular GOMS modeling system, with inclusion of reactive planning. It can represent complex expert performance and can also serve as an intelligent control system.

Constraint-based Planning NASA Ames Research Center Ari Jonsson (RIACS) Tania Bedrax-Weiss (QSS/ARC), Will Edgington (QSS/ARC), Jeremy Frank (ARC/IC), Conor McGann (QSS/ARC), Paul Morris (ARC/IC), David Smith (ARC/IC) Group Lead: Nicola Muscettola (ARC/IC)

Observation scheduling in ground-based, airborne, and space-based observatories is a difficult problem of growing importance for NASA. Constraint-based planners and schedulers should be able to quickly generate complex, flexible, concurrent plans, leaving maximum flexibility for intelligent run-time choices when faced with faulty equipment, changing observing conditions, changing sets of observation requests, and other likely complications. This research [5, 43] will extend the Remote Agent planner to handle metric domains, interacting concurrent actions, uncontrolled external processes, and resource usage. A combination of planning, scheduling and operations research will be used in a constraint-based interval planning (CBIP) framework. Advances will be integrated with JPL's Mission Data Systems (MDS) project and with planners for ground operations, spacecraft, rovers, and other applications.

Multi-Resolution Planning in Large Uncertain Environments Massachusetts Institute of Technology Leslie Pack Kaelbling (MIT)

In complex domains, an agent's real-time planning problems may exceed its computational resources. Hierarchical models and task abstractions can help teams of robots function in unpredictable and changing environments. As world conditions and agent beliefs change, new abstractions will focus attention on different aspects of the domain. Partially observable Markov decision process (POMDP) models will be used to find best-possible action plans within the available computation time [44].

Heuristic Control of Planning and Execution in Metric/Temporal Domains Arizona State University Subbarao Kambhampati (ASU)

NASA missions will need efficient, scalable planning systems that handle metric and temporal constraints. Distance-based search control has dramatically scaled up plan synthesis in recent years. Most of this work has been for classical state-space planning (such as NASA's Remote Agent Experiment planner), and has yet to be extended to metric/temporal planning. Distance-based heuristic search techniques for metric and temporal partial-order planning. Heuristic control techniques from classical planning will be generalized to least-commitment and metric/temporal planners. An existing planner with heuristic search control [23] will be extended to handle conjunctive partial-order plans with temporal planning graphs. A decision-theoretic approach will be developed for execution monitoring and control, with loose coupling to the planner for superior scale-up potential. Derivation of distance-based search control heuristics from planning graphs (currently implemented in AltAlt) will be extended to temporal planning graphs and conjunctive partial-order planning. Automated reachability analysis may be used to control the planner's search process. UCPOP, a classical partial-order planner, is being extended with distance-based heuristics for RAX-style problems. (Experience with the HSTS planner showed partial-order planning to be attractive for metric/temporal problems, but with a need for better search control.) A novel decision-theoretic approach for execution monitoring will also be developed, to control execution while exploiting the advantages offered by the deterministic plan synthesis techniques. Automated planning will enable onboard autonomy for future space missions. This work will extend the most promising ideas from the current planning and scheduling research, developing a scalable modular planning architecture for metric and temporal domains.

Approaches to Human-Centered Software Development Massachusetts Institute of Technology Nancy Leveson (MIT) John Hansman (MIT), Margaret Storey (MIT)

Automated tools can help design robust mission systems and propagate expertise from one mission to another. This task will develop formal methods for designing mission systems that include human experts. Formal modeling and model-checking [73, 58, 51] will help to ensure completeness and safety of mission systems, while simulation and visualization methods will help validate the designs. The tools will enhance situation awareness, minimize human errors, optimize allocation of tasks, enhance learnability, and simplify training.

Software Agents to Support Distributed Team Operations NASA Johnson Space Center Jane Malin (JSC/ER) Hongbin Wang (UT/HSCH), Debra Schreckenghost (Metrica), David Kortenkamp (Metrica), Jiajie Zhang (UT/HSCH), Jack Smith (UT/HSCH and JSC), Kathy Johnson (UT/HSCH and JSC)

NASA must increasingly rely on experts who divide their time among multiple projects. Agent-based systems can help support the activity of such distributed teams, providing better situational awareness and prompt and appropriate information dissemination and coordination. This reserach will develop an information architecture to support nominal and off-nominal operations by teams of cooperating humans and software agents. Delivery of information will accommodate changing tasks, roles, preferences, and locations of crew and flight controllers. The agent-based architecture will be demonstrated for simulated operation of a life support system.

Human-Centered Advisory and Assistant Systems for Mission Control NASA Johnson Space Center Jane Malin (JSC/ER) Carroll Thronesbery (SKE/JSC), Kathy Johnson (UT/HSCH and JSC), David Overland (JSC), Debra Schreckenghost (Metrica/JSC), Land Fleming (Hernandez/JSC), Lou Flores, Arthur Molin (SKE/JSC), Grace Lei (SKE/JSC), Dan Smith (SKE/JSC), Patrick Oliver (JSC), Gene Peter (JSC), Kevin Taylor (JSC)

This research is jointly developing an agent-based anomaly response management system and a methodology for such development. The system will include an electronic console log and workspaces for tracking and responding to issues, with flexible support for changing situations, tasks, roles, and user locations. Investigators will focus on communication protocols and on prototype-specification and knowledge-capture tools for user proxy/assistant agents, based initally on software application development team roles, methods, and tools. Applications include a life-support testbed system and an issue-tracking system for Space Shuttle Mission Control.

Intelligent Distributed Execution Architecture NASA Ames Research Center Nicola Muscettola (ARC/IC) Chuck Fry (QSS/ARC), Rich Levinson (QSS/ARC), Chris Plaunt (ARC/IC), Greg Dorais (ARC/IC), Baskaran Vijayakumar (QSS/ARC), Felix Ingrand (LAAS/CNRS), Bernardine Dias (CMU/RI), Solange Lemai (LAAS/CNRS)

Levels of autonomy vary from simple reactivity to the complex reasoning of a rover science mission. Advanced autonomous agents are typically built from several types of reasoning systems and communications protocols, which complicates design, integration, and validation. The Intelligent Distributed Execution Architecture (IDEA) [24, 6, 55] is a control system architecture that uses similar execution machinery and task-network communication protocols for all reasoning agents and components. This scaleable reactive planner approach should accommodate various functional responsibilities, reactivity requirements, and needs for problem solving power. Each agent may use a different type of planner, to match its function. At the lowest level, the sense-plan-act loop will offer real-time guarantees on the order of a few milliseconds per execution cycle.

Rover Autonomy Architecture Automated Reasoning NASA Jet Propulsion Laboratory Issa Nesnas (JPL/Robotics) Tara Estlin (JPL/AI), Caroline Chouinard (JPL/AI)

Future rover executives must model, estimate, and predict resource usage in order to make autonomous decisions about best actions. This may require recursive queries to various functional subsystems, which may need to query other subsystems to form their estimates. The CLARAty architecture for autonomous rovers will add such capabilities to JPL's Mission Data System (MDS) for spacecraft command and control. Current development includes intelligent reasoning about resource estimation for model-based real-time control of various robotic and rover platforms. This will require communication between a Decision Layer of planning and execution algorithms and a Functional Layer for physical reasoning and control. Functional object prototypes representing real robotic hardware will maintain and report resource needs, in the same way that state information is handled [27, 17, 35, 28, 28, 30, 57, 31, 72]. Planning software will be able to query these modules to make resource usage predictions.

Probabilistic Reasoning for Complex Dynamic Systems Harvard University Avi Pfeffer (Harvard) Leon Peshkin (Harvard), Brenda Ng (Harvard)

Agents that monitor dynamic systems must integrate multiple sources of evidence over time, reason under uncertainty, and maintain beliefs. Probabilistic reasoning provides a coherent framework, but has not been scaled to large dynamic systems. Hierarchical decomposition will be used to implement probabilistic techniques for monitoring complex dynamic systems (such as spacecraft or life support systems). Reasoning can then be performed by a hierarchy of reasoning agents with limited intercommunication. Algorithms should aid diagnosis, prediction, and real-time monitoring.

Mixed-Initiative Planning and Scheduling for the Mars'03 Mission NASA Ames Research Center Kanna Rajan (RIACS) John Bresina (ARC/IC), Len Charest (JPL/AI), Will Edgington (QSS/ARC), Ari Jonsson (RIACS), Bob Kanefsky (QSS/ARC), Pierre Maldague (JPL), Paul Morris (ARC/IC) Mitchell J. Ai-Chang (QSS/ARC), Jennifer Hsu (FCCD/ARC), Jeffrey Yglesias (QSS/ARC), Alan Baba (JPL), Adans Ko (JPL)

Science activities by Mars rovers will require local, real-time decisions based on opportunities, environmental conditions, and resource availability. With mixed-initiative planning for semi-autonomous execution, JPL's ground-based APGEN planning tool will be able to schedule activities that allow for Mars-based decisions [2, 7, 6, 43]. The rover's planner and executive will then reason about opportunities and resource limits while enforcing ground-based constraints and mission/flight rules.

Distributed Crew Interaction with Advanced Life Support Control Systems NASA Johnson Space Center Debra Schreckenghost (Metrica) David Kortenkamp (Metrica), Carroll Thronesbery (SKE), David Woods (OSU)

Advanced life support systems will require intermittent supervisory control without labor-intensive monitoring. This reserach will develop a task-management interface that supports distributed team activities by humans and software agents. Benefits include both distributed control capability for life support systems and an integrated suite of agent-based communication and coordination tools for mixed-initiative control of other complex, automated systems.

Heterogenous Multi-Rover Coordination for Planetary Exploration Carnegie Mellon University Reid Simmons (CMU/RI) Anthony Stentz (CMU/RI), Stephen Smith (CMU/RI), Pradeep Khosla (CMU/RI), Tucker Balch (CMU/RI), Howie Choset (CMU/RI), Alfred Rizzi (CMU/RI), Jeff

Schneider (CMU/RI), Sebastian Thrun (CMU/RI), David Wettergreen (CMU/RI) Dani Goldberg (CMU/RI), Bernardine Dias (CMU/RI), Vincent Cicirello (CMU/RI), Trey Smith (CMU/RI)

Operations in remote and unpredictable environments will require teams of software agents and robots. This reserach [76] will explore market negotiation for team formation, resource allocation, and coordination, with decision-theoretic strategies for tradeoffs between planning, negotiation, sensing, and action. Market-based distributed scheduling will be combined with more global planning, with robots negotiating appropriate levels of coordination for each task. Distributed task synchronization and monitoring will enable teams to perform tasks that no one robot can do individually. The research emphasis will be on choosing best coordination methods, for tasks that involve sophisticated reasoning about geometry, uncertainty, perceptual abilities, and hybrid discrete/continuous control.

Using Combinatorial Optimization Algorithms to Improve Automated Planning and Scheduling NASA Jet Propulsion Laboratory Ben Smith (JPL/AI) Richard Korf (UCLA), Steve Chien (JPL/AI), Russell Knight (JPL/AI)

Automated planning systems can reduce mission-planning effort, improve mission quality, and reduce operations costs. However, many important NASA problems – such as instrument observation scheduling for celestial surveys and planetary mapping – are too difficult for general-purpose planners and too diverse for special-purpose optimizers. A combination of the two approaches may solve large, complex problems that have strongly interacting combinatorial optimization sub-problems [49]. This would enable breakthrough improvements in the speed and solution quality of AI plan optimization.

Limited Contingency Planning for Concurrent Activities NASA Ames Research Center David Smith (ARC/IC), John Bresina (ARC/IC) Paul Morris (ARC/IC), Richard Dearden (RIACS) Nicolas Meuleau (QSS/ARC), Sailesh Ramakrishnan (QSS/ARC), Betty Lu (FCCD/ARC), Rich Washington (RIACS), Tony Barrett (JPL/AI), Steve Chien (JPL/AI), Barbara Engelhardt (JPL), Jeremy Frank (ARC/IC), Ari Jonsson (RIACS)

A contingency planner can generate highly robust plans, permitting intelligent autonomous operations in uncertain environments . Mission plans that allow for the most likely and important contingencies will enable appropriate real-time response to foreseen but unpredictable events. Wait states will be greatly reduced, and spacecraft and rovers will be able to accomplish more science in less time and with less risk. This research [21, 20, 19, 7] will extend the Remote Agent planner to handle temporal and metric resource constraints under overlapping concurrent activities of varying durations.

Stochastic Anytime Search With Applications in Autonomous Planning and Scheduling University of Illinois at Urbana-Champaign Benjamin Wah (UIUC)

Planning and control of spacecraft operations involves generating low-level spacecraft commands from mission goals. Search-based function optimization is often used in continuous domains such as navigation, orbital mechanics, and thruster control. This reserach [10] will extend such techniques to discrete (or hybrid) domains and non-differentiable, stochastic, or non-closed form function spaces. This will be applied to "anytime" planning/scheduling systems for most time-critical planning and control problems, including achievement of high-level science and engineering goals with the constraints of hardware capabilities and mission flight rules .

Onboard Rover Autonomy Architecture NASA Ames Research Center Rich Washington (RIACS) John Bresina (ARC/IC), Howard Cannon (ARC/IC)

A flexible but standard autonomy software architecture will be developed, to integrate Mars Rover development at ARC, JPL, and partner institutions. Teams need to understand where there is agreement and where different approaches are still being followed. A joint ARC/JPL field demonstration of autonomous science rovers will motivate standards and specifications for a modular autonomy architecture that spans multiple abstraction levels. The standards effort will create a continuing exchange of ideas and software between teams, and should facilitate technology transfer to missions. Resolution of design conflicts will also clarify architectural issues for other autonomous systems [7, 21, 77] .

Compilation of Model-based Programs for Reactive Autonomous Control JHU Applied Physics Laboratory David Watson (JHU/APL) Brian Williams (MIT)

Precompiled decision policies enable model-based control to scale up to complex fault monitoring, diagnosis, and repair tasks. Diagnostic reasoning and failure recovery can be very fast, with limited onboard computation. Flight engineers can review generated autonomy rules, and the approach is compatible with current NASA mission architectures. This reserach [74, 47, 11, 47, 6, 47, 75] will create development tools and an integrated system for planning and reactive execution using the Reactive Model Programming Language (RMPL) framework. The Titan model-based executive system will be capable of estimating current spacecraft operating modes, detecting and repairing failures, and executing commands in real time while remaining responsive to other spacecraft needs. This will be demonstrated on a realistic problem set.

Interleaved Contingent Planning and Execution University of Washington Daniel Weld (UW/CSE)

Situations are seldom fully known and actions do not always produce their intended consequences. Contingent plans allow for uncertainty by modeling likely divergence points. Flexible planning improves execution time if a modeled contingency occurs, at the cost of having to do extra planning work. This one-year reserach explores ways to combine relational logic with Markov decision process models for efficient contingent planning. The planner will run just a little ahead of execution, thinking about what is likely to happen next and how to deal with it.

Autonomous Rotorcraft Project NASA Ames Research Center Matt Whalley (ARC/ARH), Dan Christian (ARC/IC) Ann Patterson-Hine (ARC/IC), Mike Freed (SJSU/ARC), Greg Schulein (SJSU/ARC), Marc Takahashi (QSS/ARC), Robert Harris (SJSU/ARC), Yiyuan Zhao (UMN)

Autonomous rotorcraft can enhance national security and public service support, provide user-friendly personal transport, and perhaps offer vertical lift capability for planetary exploration. This new class of vehicle is also an ideal platform for developing and demonstrating automated reasoning software for Mars landers, aircraft or satellite clusters, and other NASA flight applications. The challenge – and opportunity – is to embed rotorcraft-specific flight control and vehicle health maintenance within a planning and execution framework shared with other autonomous vehicle types. Onboard decision making will incorporate vision-based processing, sensor monitoring, spatial reasoning, communications, mission constraints, and high-level planning . A simulated scout mission will drive development of real-time reactive control and intelligent mission planning.

A Hybrid Discrete/Continuous System for Health Management and Control Massachusetts Institute of Technology Brian Williams (MIT) David Kortenkamp (Metrica), Michael Hofbaur (MIT/AI) Jonathan How (MIT), Dave Miller (OU), Jesse Leitner (GSFC), Maria Zuber (MIT), Robert Goldman (Honeywell)

NASA's Livingstone diagnostic system on Deep Space 1 was very successful, but could not reason about continuous dynamics. It can be extended with hierarchical, probabilistic models to track component parameters and detect multiple faults via noisy data from complex hybrid system behaviors. Applications include planetary rovers, spacecraft, in situ propellant production (ISPP) plants, life support systems, and other devices with rich continuous and discrete behaviors. Automated learning of the hierarchical models may enable rapid prototyping of simulators from flight hardware data, for validation of monitoring and diagnosis software.

Autonomous Rover Command Generation - ASPEN NASA JPL Planning and Scheduling Group, Planning Technologies Rob Sherwood, Andrew Mishkin, Tara Estlin, Steve Chien, Scott Maxwell, Barbara Engelhardt, Brian Cooper, Gregg Rabideau

This system is a proof-of-concept prototype for automatic generation of validated rover command sequences from high-level science and engineering activities [68, 11, 15, 65, 63, 13, 12, 66, 14, 67, 65, 16] . This prototype is based on ASPEN, the Automated Scheduling and Planning Environment. This AI-based planning and scheduling system will automatically generate a command sequence that will execute within resource constraints and satisfy flight rules. Commanding the rover to achieve mission goals requires significant knowledge of the rover design, access to the low-level rover command set, and an understanding of the performance metrics rating the desirability of alternative sequences. It also requires coordination with external events such as orbiter passes and day/night cycles. An automated planning and scheduling system encodes this knowledge and uses search and reasoning techniques to automatically generate low-level command sequences while respecting rover operability constraints, science and engineering preferences, and also adhering to hard temporal constraints.

Adaptive Problem Solving (APS) NASA JPL Planning and Scheduling Group, Planning Technologies Steve Chien, Barbara Engelhardt, Gregg Rabideau, Robert Sherwood, Darren Mutz, Forest Fisher, Ben Smith, Tara Estlin, Russen Knight, Tony Barrett

Techniques from machine learning and statistics are being applied to learn heuristics to improve automated scheduling [15, 65, 63, 13, 15, 12, 14, 30, 26, 48, 15, 16, 33, 61] . This is a collaborative effort with Operations Mission Planner - 26m team at JPL which is fielding a scheduler for the 26 meter antenna subnetwork of the Deep Space Network (DSN). The goal is development of learning techniques to improve scheduling speed and quality by learning heuristics specialized to a distribution of problems.

CASPER (Continuous Activity Scheduling Planning Execution and Replanning) NASA JPL Planning and Scheduling Group, Planning Technologies Steve Chien, Russle Knight, Gregg Rabideau, Robert Sherwood, Darren Mutz, Forest FIsher, Tara Estlin

To achieve a higher level of responsiveness in a dynamic planning situation, this research utilizes a continuous planning approach and implemented in a system called CASPER (for Continuous Activity Scheduling Planning Execution and Replanning) [15, 65, 63, 13, 15, 12, 14, 30, 26, 48, 15, 16, 33, 61] . Rather than considering planning a batch process in which a planner is presented with goals and an initial state, the planner has a current goal set, a plan, a current state, and a model of the expected future state. At any time an incremental update to the goals or current state may update the current state of the plan and thereby invoke the planner process. This update may be an unexpected event or simply time progressing forward. The planner is then responsible for maintaining a consistent, satisficing plan with the most current information. This current plan and projection is the planner's estimation as to what it expects to happen in the world if things go as expected. However, since things rarely go exactly as expected, the planner stands ready to continually modify the plan. Current iterative repair planning techniques enable incremental changes to the goals and the initial state or plan and then iteratively resolve any conflicts in the plan. After each update, its effects will be propagated through the current projections, conflicts identified, and the plan updated (e.g., plan repair algorithms invoked).

CLEaR (Closed Loop Execution and Recovery) NASA JPL Planning and Scheduling Group, Planning Technologies

CLEAR is an integrated planning and execution framework for autonomous control of robotic entities [35] . The CLEaR system currently utilizes CASPER and TDL, and focuses on on the use of both near-term reactive behavior and long-term deliberative decision making.

Unmanned Air Vehicles (UAVs) NASA JPL Planning and Scheduling Group, Autonomous Aerial Vehicles Forest Fisher, Sandeep Gulati, Mark James, Ryan Mackey, Robert Koneck

(Unmanned Air Vehicles) consist of integrating JPL planning, diagnostics, and prognostics systems as part of the control architecture for UAVs .

Citizen Explorer (CX1) NASA JPL Planning and Scheduling Group, Spacecraft Autonomy Steve Chien, Gregg Rabideau, Robert Sherwood, Colette Wilklow, Jason Willis

Citizen Explorer (CX1) is a small earth orbiting satellite built and managed by the Colorado Space Grant Consortium launching in December 1999 [13, 12, 26] . The ASPEN planning and scheduling system was used in the design of the spacecraft evaluate power and other engineering requirements. The ASPEN planning and scheduling system is also being used in the ground operations system. ASPEN will be used to automatically generate validated command sequences to command CX1.

Balancing competing mission needs (e.g., power, data storage) during mission design is a complex problem. Designing a mission operations strategy and hardware configuration to achieve mission goals while staying within cost, mass, and volume constraints is a challenging optimization problem. Automated planning technology can aid in this process by enabling mission designers to layout mission operations plans and analyze science return in the context of different hardware configurations and mission operations policies (e.g., commanding frequency, pass lengths, etc.).

The same planning system can then be used during operations to automate generation of validated command sequences. This reduces commandding turnaround time and effort, and can also increase overall science return by increasing utilization of scarce resources. Automated planning/scheduling technologies have great promise in reducing operations cost and increasing the autonomy of aerospace systems. By automating the sequence generation process and by encapsulating the operation specific knowledge, we hope to allow spacecraft commanding by non-operations personnel, hence allowing significant reductions in mission operations workforce with the eventual goal of allowing direct user commanding (e.g., commanding by scientists).

The CX1 satellite will be launched on a Delta-II launch vehicle December 15, 1999. As a technology validation and demonstration, the Colorado Space Grant Consortium has been using the ASPEN automated planning and scheduling system in the mission design for the CX1 satellite. Plans are also to use the ASPEN system for automated command generation of the CX1 spacecraft.

Distributed Self-Commanding Robotic Systems NASA JPL Planning and Scheduling Group, Spacecraft Autonomy Tony Barrett, Tara Estlin, Greg Rabideau, Steve Chien

In general, autonomous spacecraft and rovers must balance long-term and short-term considerations. They must perform purposeful activities that ensure long-term science and engineering goals are achieved and ensure that they each maintain positive reso urce margins. This requires planning in advance to avoid a series of shortsighted decisions that can lead to failure. However, they must also respond in a timely fashion to a dynamic and unpredictable environment. In terms of high-level, goal-oriented activity, the robotic systems must modify their collective plans in the event of fortuitous events such as detecting scientific opportunities like a Martian hydrothermal vent or a sub-storm onset in Earth's magnetosphere, and setbacks such as a spacecraft losing attitude control.

Techsat-21 NASA JPL Planning and Scheduling Group, Spacecraft Autonomy Steve Chien, Rob Sherwood,Becky Castano,Ashley Davies,Gregg Rabideau,Daniel Tran,Ben Cichy,Nghia Tang,Rachel Lee,Russell Knight,Steve Schaffer

The Techsat-21 project consists of flying the Autonomous Sciencecraft Experiment flight experiment (ASE) onboard the Air Force Techsat-21 constellation (an unclassified mission scheduled for launch in 2004) [68, 63, 65, 64] . ASE will use on-board science analysis and replanning to redically increase science return by enabling intelligent downlink selection and autonomous retargeting. A typical ASE demonstration scenario involves monitoring of active volcano regions such as Mt. Etna in Italy.

Three Corner Sat (3CS) NASA JPL Planning and Scheduling Group,Spacecraft Autonomy E. Hansen, C. Koehler, J. Michels, S. Wichman, B. Sanders, C. Wilklow, University of Colorado, Space Grant S. Chien, R. Knight, R. Sherwood, B. Engelhardt, G. Rabideau, Artificial Intelligence Group, JPL

The Three Corner Sat (3CS) satellite project is a mission being developed jointly by Arizona State University (ASU), The University of Colorado, Boulder (CU), and New Mexico State University (NMSU). 3CS consists of three coordinated satellites that will be deployed in a stack configuration from the Space Shuttle and will then separate to form a "virtual formation. [39] "The goals of the 3CS mission include the demonstration of stereo imaging, formation flying and innovative command and data handling, including on-board autonomy.

Distributed Rovers/MISUS (Multi-Rover Integrated Science Understanding System) NASA JPL Planning and Scheduling Group, Rover Autonomy Tara Estlin, Daniel Gaines, Forest Fisher, Brad Clement, Gregg Rabideau

While it is up to mission designers to determine the optimal number of rovers for a given mission, multiple rovers have three types of advantages over single rover approaches: force multiplication, simultaneous presence and system redundancy [29, 28, 28, 31, 31, 32, 17, 17, 35] .

- Force multiplication. Multiple rovers can perform certain types of tasks more quickly than a single rover, such as: performing a geological survey of a region or deploying a network of seismographic instruments. We call these cooperative tasks.

- Simultaneous presence. Multiple rovers can perform tasks that are impossible for a single rover. We call these coordinated tasks. Certain types of instruments, such as interferometers, require simultaneous presence at different locations. Rovers landed at different locations can cover areas with impassable boundaries. Using communication relays, a line of rovers can reach longer distances without loss of contact. More complicated coordinated tasks can also be accomplished, such as those involved in hardware construction or repair.
- System redundancy. Multiple rovers can be used to enhance mission success through increased system redundancy. Several rovers with the same capability may have higher acceptable risk levels, allowing one rover, for example, to venture farther despite the possibility of not returning. Also, because designing a single rover to survive a harsh environment for a long periods of time can be difficult, using multiple rovers may enable missions that a single rover could not survive long enough to accomplish.

DSSC (Deep Space Station Controller) NASA JPL Planning and Scheduling Group, Rover Autonomy Leslie Paal, Forest Fisher, Mark James, Barbara Engelhardt, Han Park

DSSC is an extension to the work performed for the Deep Space Terminal (DS-T) task, part of the Deep Space Network, in the area of track automation. This work utilizes the CLEaR system to provide the capability for robust dynamic desision making and execution management for autonomous DSN ground station operations .

Scheduling with Resource Envelopes NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Nicola Muscettola

Development and full computational study of resource envelope scheduler with comparison to state of the art flexible schedulers on challenging benchmark problem sets [54] .

Spacecraft Mobile Robot Autonomy NASA Ames, Computational Science Division/Autonomy and Robotics (ARA) Greg Dorais

Development of an adjustably autonomous control system for a free-flying, remote sensing vehicle capable of autonomously navigating in three dimensions and interacting with local and remote users in a manner that is useful, easily understood and easily commanded in a minimally time-consuming manner [24] .

Antarctic Robotic Traverse NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Liam Pedersen

Development of a Robotic Antarctic Traverse mission. This will advance the sciences of astrobiology, climate change, and glaciology; and will push the development of ruggedized autonomy for long duration robotic missions.

Biormorphic Robotics NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Silvano P Colombano

The development of robotic hardware and control strategies that mimic biological systems (e.g. snakebot, scorpion) .

Coupled Layer Architecture for Robotic Autonomy (CLARAty) Architecture NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Anne Wright

This research leverages CLARAty to enhance software architecture, interoperability, and maintainability for the K9 Rover. This effort contributes to CLARAty development to increase robustness and applicability to the advancement of NASA robotic technologies .

Executive-Level Decision Making NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Rich Washington

Development of technologies that increase science productivity and scientific return by enabling rovers to effectively respond and react to the inherent uncertainty in planetary exploration missions [21, 7, 77, 3] .

Integrated Technology Demonstrations for Mars Science Laboratory (MSL) NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Liam Pedersen

Develop integrated demonstrations of technologies for robotic Mars exploration, pertinent to the 2009 MSL Mars rover mission. These demonstrations include single cycle instrument deployment, contingent ground planning, robust conditional execution, and 3D science data visualization . A sequence of integrated demonstrations onboard the K9 rover, both in the NASA/ARC Marscape Test Facility and in the field, are in progress to aid in the infusion of these technologies onto the MSL rover mission .

K9 Platform, Architecture and Test Facility NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Maria G Bualat

Development and integration of enabling rover technologies on the K9 rover testbed for NASA missions.

2003 Mars Exploration Rovers (MER) Mixed Initiative Plan Generator (MAPGEN) NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Kanna Rajan

ARA, in collaboration with the Jet Propulsion Laboratory, is supporting the 2003 MER in the development of a science planning and scheduling tool, the MAPGEN [2, 43, 6] .

2009 Mars Science Laboratory (MSL) Mission Planning & Execution Project NASA Ames, Computational Science Division/Autonomy and Robotics (ARA), Mark Drummond

ARA, in collaboration with the Jet Propulsion Laboratory, is supporting the MSL Project by providing technology for integrated model-based diagnosis, plan execution, plan recovery and model-based safing. This work is being done within the framework of JPL's Mission Data System (MDS) .

Look-ahead Model Based Programming NASA Ames, Computational Science Division/Autonomy and Robotics (ARA) Sriram Narasimhan

Development of scenarios that illustrate the need for look ahead in execution to update the expected utility of different branches in the control programs and plans. This is based on the belief about the current state(s) as diagnosed by the model-based diagnosis engine (Livingstone). The suitability of application to different execution languages will be assessed .

PlanWorks NASA Ames, Computational Science Division/Autonomy and Robotics (ARA) Conor McGann

Creation of a development environment that makes it easier to: build, validate and debug models; understand the plans created by automated planners; and understand the process by which those plans were found [5] .

Advanced Information Systems Technology (AIST) Earth Observing Satellite Scheduling NASA Ames, Computational Science Division/Autonomy and Robotics (ARA) Robin Morris

Development of an autonomous system for scheduling and rescheduling requests for sensing data on earth observing satellites [46, 43] .

Constraint Based Planning NASA Ames, Computational Science Division/Autonomy and Robotics (ARA) Ari Jonsson

Development of planning techniques and software for domains with complex temporal and resource constraints [43, 42] .

Imagebot NASA Ames, Computational Science Division/Autonomy and Robotics (ARA) Keith Golden

Development of an agent-based system for processing and tracking scientific data, including spacecraft image data.[38]

SOFIA Observation Scheduling NASA Ames, Computational Science Division/Autonomy and Robotics (ARA) Jeremy Frank

Development of observation scheduling and flight planning techniques for the Stratospheric Observatory for Infrared Astronomy (SOFIA) airborne observatory [40] .

Mars Exploration Rover Human Centered Computing NASA Ames, Computational Science Division/Collaborative and Assistant Systems (CAS), Jay Trimble

Develop technologies and procedures to increase productivity of surface operations for Mars Exploration Rover 03 missions based on human centered computing techniques .

MERBoard NASA Ames, Computational Science Division/Collaborative and Assistant Systems (CAS), Ubiquitious Computing and User-Centered Design Group Jay Trimble

A large screen interactive worksurface to support collaboration during surface operations for the Mars Exploration Rover 03 Missions.

IxTeT-eXeC Solange Lemai and Felix Ingrand, LAAS/CNRS, Toulouse, France

IxTeT-eXeC is an extension of a temporal planner which allows exeution control, plan repair and replanning when necessary [50, 22].

Dynamic Ontologies Fiona McNeill, Alan Bundy and Marco Schorlemmer, Centre for Intelligent Systems and their Applicatios, School of Informatics, Universty of Edinburgh

This work sees to dynamically refine agents' representation ontologeis during execution to improve system robustness [34].

High-level Robot Programming and Program Execution Mikhail Soutchanski, Ryerson University, Toronto, Ontairo, Canada

This work proposes a logical framework for robot programing that accommodates multiple leveles of access control [69] .