Attaining Program Affordability Through Integration Of Logistical Operations and Health Maintenance

Andre Goforth*
NASA Ames Research Center, Moffett Field, CA, 94035-1000, USA

Gabor Karsai†
Vanderbilt University Institute for Software Integrated Systems,Nashville,TN 37235

Michael Orosz‡
Robert Neches§
Pedro Szekely**
University of Southern California Information Sciences Institute, Marina del Rey, CA 90292

Naveen Ashish††
USRA Research Institute for Advanced Computer Science at NASA Ames Research Center, Moffett Field, CA, 94035

Program affordability needs to be built-in at the initial concept formulation stage. For NASA's space exploration vision this is critical for long range sustainability of human presence in space. What is often overlooked in the initial concept formulation of a large scale system endeavor such as NASA's Constellation program is the hidden cost of maintaining requisite operational safety margins and redundancy through adequate supply chain logistics. Ensuring adequate supply chain logistics necessitates the integration of operations and maintenance cycles. What enables this integration is the coordination and reconciliation across multiple equipment types of system health features and logistical information such as: (1) prognostic drivers from Integrated Vehicle Health Monitoring (IVHM) systems producing proactive condition-based "maintain me" demands, (2) maintenance management systems tracking usage and producing scheduled maintenance demands, (3) unscheduled maintenance demands resulting from any trouble reports entered by human observers of conditions missed by the IVHM system, and (4) implicit maintenance demands resulting from mission plans which require assignment of vehicular/robotic assets and consequently require assurance of the assigned assets' fitness for the intended tasks. In this paper we discuss how Coordinated Multi-source Maintenance on Demand (CMMD) technology, which is being is transitioned to the USMC Coherent Analytical Computing Environment (CACE) program and the Joint Strike Fighter Program, can be applied to the NASA domain, and its benefits in terms of mission affordability, operations efficiency and system health effectiveness. Using concepts derived from CMMD, we discuss the kind of IVHM capabilities needed to optimize multiple, parallel, yet inter-linked, operations-maintenance cycles, thereby optimizing program affordability while meeting specific mission supportability requirements across a broad range of mission scenarios.

---

[*] CMMD Co-PI & Ames Project Manager, Code TC, MS269-4
[†] Associate Professor of Electrical and Computer Engineering
[‡] Member of Technical Staff, Distributed Scalable Systems Division
[§] Division Director, Distributed Scalable Systems Division
[**] Member of Technical Staff, Distributed Scalable Systems Division
[††] Computer Scientist

I. Introduction

Program affordability is one of the cornerstones of long-term sustainability of NASA's vision for Space Exploration Systems. At the very least, this vision depends upon 1) public interest and support, 2) programmatic infrastructure, and 3) application of emerging technologies that create unprecedented capabilities. Efficiency and supply-chain logistics are key contributors to program affordability, and are dependent on infrastructure and application of emerging technologies. News of inefficiency and failure to maintain adequate supply chain logistics sufficient for realization of mission goals will rapidly erode public interest and support. Therefore, our development of emerging coordinating and reconciliation of operations and maintenance technology clearly addresses the whole of the sustainability concept. Effective harnessing of emerging technologies through innovative systems engineering concepts and programs optimizes the cost of ownership to the lowest possible level while maintaining requisite safety margins and redundancy, which, in turn, contributes to the maintenance of public interest and support.

The issue of return on investment from harnessing emerging technologies to ensure program affordability is important for determining the size and scheduling of investments. Recently our proposal: Coordinated Multisource Maintenance On Demand (CMMD) was selected as one of the external Broad Area Announcement investments by the NASA Exploration Systems Mission Directorate (ESMD). While we are not yet in a position to provide metrics such as how much a given technology will reduce the cost of ownership for ESMD with it when compared to a solution without it, we are in a position to discuss the technical aspects of the technology, what applications it has been used in and the benefits and metrics realized in another domain. We can also provide a roadmap of the programmatic infrastructure that should be put in place by NASA in order to fully reap the benefits of its technological investments in the areas of diagnostics, prognostics and systems health management. With these in place it will be feasible to derive optimal ownership metrics based on usage and experience through the application of technology.

Human presence in space is one of the core principles of the space exploration vision of the Exploration Systems Mission Directorate. Space as a new frontier is not forgiving to explorers. Without a revolutionary maintenance logistics scheduling capability, ESMD may find that operational and maintenance logistics for its Constellation Transportation System (CTS) consumes a major portion of human resources available, to the degree that there may be chronic and severe impacts on overall operational efficiency and hence mission supportability. ESMD needs the capability of revolutionary maintenance scheduling to manage supply chain logistics so that operational margins and redundancy are satisfied in all critical circumstances.

The rationale for the CMMD project is that if its results are applied in the development of Project Constellation's WBS 4.6.3.3.3 Logistics and 4.6.3.3.4 Flight Operations[1] then the Directorate will have a major requisite that enables a revolutionary health maintenance scheduling capability across multiple campaigns and across multiple life cycle spirals.

One element of the ESMD's Human & Robotic Technology concept of *sustainability*[2] is the Strategic Technical Challenge (STC) of flexibility. The proposed health maintenance technology will directly address this challenge by enhancing exploration missions' supportability through flexible logistical operations and maintenance integration using on-demand and just-in-time (JIT) scheduling. In addition, the capability for revolutionary maintenance scheduling will also address the STC of *margins and redundancy* and *as safe as reasonably achievable* (ASARA) human presence in space. CMMD project's technology will provide the ESMD's exploration missions with critical next-generation logistical capability that will cover system health maintenance scheduling and situational awareness for maintenance personnel and crewmembers. By sponsoring this research now, within the Initial Design, Advanced Development phases, the ESMD will be able to use the provided capabilities to inform the mission design process, and ensure that the resulting mission and hardware designs fully leverage the revolutionary maintenance logistics scheduling capability. Furthermore, the ESMD will be in a position to continually leverage this ability to manage margins and redundancy requirements in subsequent applications across Constellation's many logistics supply chains.

Using ISS situations and data, CMMD will demonstrate immediate feasibility of on-demand maintenance driven by scheduled and unscheduled demands coordinated with the implicit demands of operations activities, with attention to EVA suits (a critical driver). It will serve as a *pathfinder* for operational activity demands on maintenance, coordinated across multiple exploration sites, for multiple critical systems, using probabilistic prognostic failure predictions driving maintenance wherever possible.

II. Background

A. Goals, Objectives

In a future environment in which manned space exploration teams are supported by a large cadre of partly but-not-necessarily-fully interchangeable mobile manned and robotic devices, maintenance demands will stem from multiple competing sources which need to be reconciled. Operational activities themselves fall into several domains, including crew activity planning, science planning, spacecraft command and control and robot commanding, all of which have very different characteristics and will likely be coordinated using separate software systems, but all of which interact with maintenance activities. For *sustained* and *affordable* exploration integration with all these systems is needed to adapt the maintenance schedule to other activities and to adapt the other schedules to critical maintenance activities. Additionally, equipment to be maintained will need to monitor and communicate status. For equipment that cannot monitor its own state, robotic inspection is an option. Since these robots may be performing tasks other than inspection, negotiation will be especially important to balance operational requirements with safety. The proposed technology has a system-of-systems-level impact, influencing the design of equipment (hardware and software to monitor and communicate status), robots (ability to perform inspections and suggest maintenance), and planning and scheduling systems used throughout the mission (integration into a distributed collaborative system). In addition to scheduling maintenance activities, it is also necessary to alert crewmembers to potential problems that cannot be immediately corrected, so they can work around problems as needed and avoid using equipment in ways that would be unsafe. This communication must take into account the urgency of the problem, to avoid inundating the crew with minutia, which could distract them from serious problems.

B. Approach

Our approach balances the concerns driving the competing demands for maintenance within and across assets through a task planning and scheduling system for vehicle maintenance that uses negotiation-based coordination to manage communications between operations scheduling and maintenance scheduling systems. The result is a distributed collaborative system which plans and schedules maintenance for assets driven by the needs of the operational activities. By assuring that maintenance is coordinated with operational requirements our system simultaneously optimizes operational readiness and effectiveness while reducing sustained engineering costs.

Our approach utilizes negotiation technology from USC Information Sciences Institute and Vanderbilt, developed under DARPA's Autonomous Negotiating Teamware[‡‡] (ANTS) program. This is an alternative approach to resource allocation, planning and scheduling problems. Negotiation is preferable to traditional optimization or constraint satisfaction because it focuses on surfacing tradeoffs and conflicts, and on facilitating "pushback" in which users have an opportunity to reformulate their requests as the negotiation process reveals unintended consequences and implications of original requests. ISI and Vanderbilt have built separate negotiation-based schedulers for aviation based operations and maintenance, and interfaced them to coordinate their planning. This work is transitioning to the USMC CACE program and the Joint Strike Fighter Program. The proposed effort would gain significant leverage from that work.

The technical approach for CMMD leverages our experience from the CACE project. CACE is a suite of tools to coordinate flight operations and maintenance activities to support both wartime and training operations/missions in USMC aviation. CACE takes as input guidance on training goals for pilots, missions to be flown, aircraft configuration and availability and produces monthly, weekly and daily flight and maintenance schedules. CACE is a practical application project where technologies developed under the ANTS research program have been evaluated and used.

We propose to develop and demonstrate a capability for on-demand maintenance: a system-of-systems-based capability for coordinating on-demand maintenance schedules for fleets of robots and manned vehicles. The capability will also provide situational awareness of maintenance tasks in a prioritized manner to mission exploration and maintenance personnel.

III. Technical Challenges

The primary technical challenges addressed in the CACE project were computational complexity and usability[§§]. CACE problems are a combination of planning and scheduling, which involve a large number and variety of constraints. This leads to problems that are infeasible to solve optimally, therefore CACE focused on approximation

---

[‡‡] See: http://dtsn.darpa.mil/ixo/programdetail.asp?progid=41

[§§] See problem definitions in http://www.isi.edu/~szekely/antsebook/ebook/index.html techniques that produce good solutions quickly. Coupled to the complexity issues are usability issues where the challenge is to make the automatically derived solutions understandable to users and to allow users to control and steer the solvers to produce customized solutions, if need be.

We believe the CACE methodology is immediately applicable for CMMD because, though the domains are different, the features of the problems to be solved are very similar. For example, in CACE the "operations" side of the problem focused on flight missions. Similarly, in CMMD the operations consist of *activities* to be carried out by astronauts and/or autonomous vehicles. The constraints in these two domains are comparable. In CACE there were constraints concerning the skills required to perform missions; in CMMD there are constraints concerning the skills of different astronauts, i.e. mission specialists, and the skills required for performing different activities. Similarly, in space exploration, maintenance of supporting equipment must be balanced against and scheduled to strategically support the specific operational goals. As managed resources become less specialized, they become candidates for many more types of activities – resulting in much larger search spaces to be traversed. On the other hand, while there are many similarities between the CACE and CMMD domains, NASA's plans for sustainable space exploration where humans presence in space is permanent presents some unique challenges. For example, differences in the roles, responsibilities and division of labor of the crew in space will not be the same in CMMD as those of pilots and maintainers in CACE.

IV. Technical Approach: CMMD Advanced Concepts

A. Integrated planning and scheduling

In CMMD the need to integrate planning and scheduling arises primarily from the need to integrate operations and maintenance activities. The consequences of not integrating these two functions is technology that produces plans and schedules that, at best, are inefficient and, at worst, lead to catastrophic failure. For example, if operations produces a schedule without considering the maintenance needs and requirements of required equipment, there is a high probability that one or more missions will be delayed or re-scheduled due to a lack of required equipment or aborted due to equipment failure.

Consider an operations task O that requires an instance of resource R (e.g., conducting a spacewalk activity requires a spacesuit, one of several that may be available). When the lifetime of such resources goes to zero, or prognostics report an impending problem, a maintenance task M must be inserted into the schedule to restore the lifetime of the resource. Determining when to perform M relative to all the tasks that use instances of resource R is a problem that is beyond the capabilities of scheduling systems, and requires integrating some level of planning.

The pure scheduling approach would involve explicitly modeling the interaction between the tasks that use a resource and the tasks that maintain it (e.g., including a "start-after" link between O and M so that O starts after M finishes). Such approaches fail because they cannot address the following issues:

- If there are multiple instances of the resource R, task O can start before M if it uses another instance that does not need maintenance.
- The time to schedule M depends on the aggregate use of the resource, not just on a specific task O.
- If there is more than one option to perform the maintenance, M1 and M2, only one of the options need to be scheduled.

CMMD focuses upon reasoning about the implicit dependencies among tasks that arise from their use of resources.

In our prior CACE work performed under DARPA funding we collaborated with Selman's group at Cornell University[***] to develop techniques to solve a very similar problem using Over-Constrained Integer Programming techniques (including Pseudo-Boolean Constraints)[3]. We tackled the problem of reasoning about skills that pilots acquire as a result of flying missions which in turn enable them to participate in missions that they were not originally qualified to perform. The problem is similar to NASA's problem, because the essence of it is that the ability of a resource to participate in a task changes as a function of the tasks it participates in. In particular, if necessary, the system would schedule a pilot to participate in a mission that grants him the prerequisites to participate in other missions. This is similar to scheduling a maintenance task (if necessary) to enable a resource to participate in other tasks.

The novel contribution of the approach is that it allows encoding the constraints that capture the implicit relationships between tasks without increasing the number of variables to encode the problem. This is important because it means that these constraints could be encoded without increasing the size of the state space, which has a

---

[***] http://www.cs.cornell.edu/home/selman/ants/

dramatic influence on the performance of the solver. Surprisingly, the additional constraints had no noticeable influence on the performance of the solver. Even though the encoding is more complex because it models the increase of pilot skills over time, it allows a larger number of solutions to the problem because it also increases the number of resources that can perform each task.

In CMMD we will adapt the techniques used to model the increase of pilot skills to model the reduction of resource lifetime as a result of usage or prognostics. We will also extend the techniques to model optional tasks so that the scheduler can select among different alternatives for performing a task.

These techniques will have a large impact on the quality of the schedules. Maintenance tasks will not be gratuitously scheduled before they are needed, and conversely, they will be scheduled if the operations that need those resources cannot use a replacement. The increases in quality will be measured by running the same problems with and without the planning extensions to the scheduling software.

B. Coordination with external scheduling systems

CMMD will facilitate the creation and mending of coordinated maintenance and operational schedules, within available resources and according to user-specified margins of acceptable risk and degrees of redundancy. However, this system will ultimately be subject to constraints imposed by external support systems. For example, as operational limits on the ISS due to lack of EVA suits illustrates, an over-constrained problem may be unsatisfiable due to the lack of a required tool. In the event of such a situation, CMMD should pinpoint all such external dependencies and explore appropriate alternative courses of action by interrogating external systems (e.g., ground supply) through well-defined interfaces.

Solution techniques that explicitly track reasons for failure, combined with simple models of external dependencies, can be used to integrate the CMMD planning and scheduling components with external systems.

A major drawback of traditional, automated scheduling techniques is the inability of these computational algorithms to explain the shortcomings of their computed solutions. Model-checking algorithms either find a solution or simply proclaim that there is no solution. Algorithms based on local search sacrifice completeness in return for the ability to at least provide partial solutions. Optimization-based schemes can be used to improve upon this *good-enough* solution approach by iteratively refining the results to reach better and better local optimums (or a global solution for computationally tractable problems). However, in each case, these schedulers do not provide explanations of the underlying causes for a non-optimal solution. We plan to extend our previous work in the CACE project that (a) differentiates between managed and non-managed resources and (b) supplements traditional solution techniques with the ability to associate violated constraints (and their dependencies) to domain-specific reasons behind the failure.

For example, from our CACE work, figure 1 highlights in red the report of a failure to schedule an inspection that requires an external resource. Since pilots (i.e. skill 7509) are an unmanaged resource for the particular solvers that handle this aspect of the problem, this external dependency is reported to the maintenance planner, who is expected to schedule the HOVER CHECK through personal communication with the operations department.

Similarly, highlighted in blue is a failed 10 ENGINE HOUR task required between two flights, where the first flight generates a maintenance task whose duration extends beyond the subsequent planned launch time. In contrast to the first type of error, which the CACE customer requested be resolved manually, this is an example of a failure that can be negotiated away automatically by the CACE tools. The maintenance advisor tool recognizes that this dilemma is the result of an uninformed decision made by a complementary solver, which assigns appropriate aircraft to missions. By selectively divulging such cross-cutting constraints between neighboring solvers as needed, a coordinated set of plans can be automatically negotiated.

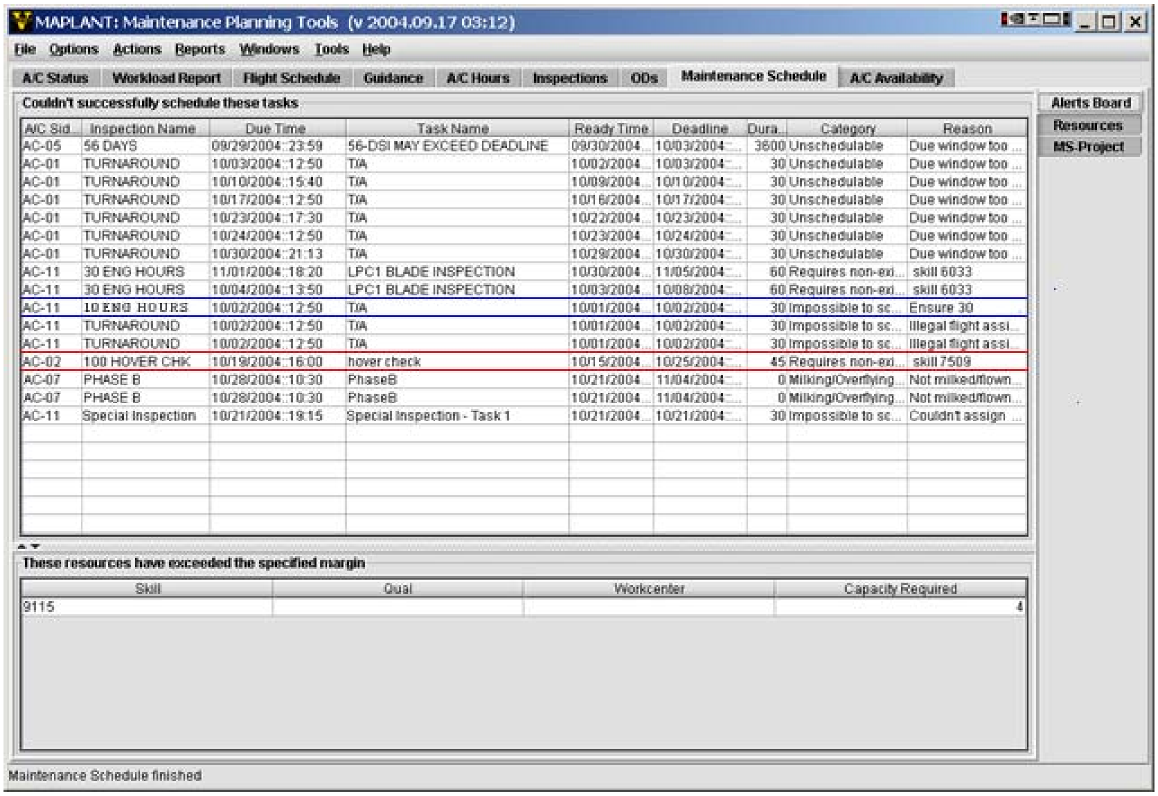

Couldn't successfully schedule these tasks

A/C Sid...	Inspection Name	Due Time	Task Name	Ready Time	Deadline	Dura...	Category	Reason
AC-05	56 DAYS	09/29/2004::23:59	56-DSI MAY EXCEED DEADLINE	09/30/2004...	10/03/2004:...	3600	Unschedulable	Due window too ...
AC-01	TURNAROUND	10/03/2004::12:50	T/A	10/02/2004...	10/03/2004:...	30	Unschedulable	Due window too ...
AC-01	TURNAROUND	10/10/2004::15:40	T/A	10/09/2004...	10/10/2004:...	30	Unschedulable	Due window too ...
AC-01	TURNAROUND	10/17/2004::12:50	T/A	10/16/2004...	10/17/2004:...	30	Unschedulable	Due window too ...
AC-01	TURNAROUND	10/23/2004::17:30	T/A	10/22/2004...	10/23/2004:...	30	Unschedulable	Due window too ...
AC-01	TURNAROUND	10/24/2004::12:50	T/A	10/23/2004...	10/24/2004:...	30	Unschedulable	Due window too ...
AC-01	TURNAROUND	10/30/2004::21:13	T/A	10/29/2004...	10/30/2004:...	30	Unschedulable	Due window too ...
AC-11	30 ENG HOURS	11/01/2004::18:20	LPC1 BLADE INSPECTION	10/30/2004...	11/05/2004:...	60	Requires non-exi...	skill 6033
AC-11	30 ENG HOURS	10/04/2004::13:50	LPC1 BLADE INSPECTION	10/03/2004...	10/08/2004:...	60	Requires non-exi...	skill 6033
AC-11	10 ENG HOURS	10/02/2004::12:50	T/A	10/01/2004...	10/02/2004:...	30	Impossible to sc...	Ensure 30
AC-11	TURNAROUND	10/02/2004::12:50	T/A	10/01/2004...	10/02/2004:...	30	Impossible to sc...	Illegal flight assi...
AC-11	TURNAROUND	10/02/2004::12:50	T/A	10/01/2004...	10/02/2004:...	30	Impossible to sc...	Illegal flight assi...
AC-02	100 HOVER CHK	10/19/2004::16:00	hover check	10/15/2004...	10/25/2004:...	45	Requires non-exi...	skill 7509
AC-07	PHASE B	10/28/2004::10:30	PhaseB	10/21/2004...	11/04/2004:...	0	Milking/Overflying...	Not milked/flown...
AC-07	PHASE B	10/28/2004::10:30	PhaseB	10/21/2004...	11/04/2004:...	0	Milking/Overflying...	Not milked/flown...
AC-11	Special Inspection	10/21/2004::19:15	Special Inspection - Task 1	10/21/2004...	10/21/2004:...	30	Impossible to sc...	Couldn't assign ...

These resources have exceeded the specified margin

Skill	Qual	Workcenter	Capacity Required
9115			4

Maintenance Schedule finished

Figure 1: Unresolved External Dependencies

Collaboration with external systems would fall between these two extremes of user-intervention and negotiated refinements of multiple plans between loosely coupled but cooperating schedulers using complex negotiation protocols. Using a simple interrogation API, CMMD could explore possible local replans based on limited information from the external environment, where the extent of these limits depends on the impedance mismatch between the concepts and the capabilities of CMMD and those of the external system.

In order to demonstrate the ability to carry out system-of-systems coordination, we will interface to external schedulers, such as EUROPA, and demonstrate the ability of CMMD to reason about the implications of off-station supply requirements on previously planned operations.

C. Handling plan disruptions

Plans and schedules get disrupted for many reasons. Tasks may take longer than expected tying up resources needed for other tasks; equipment may fail or require maintenance sooner than planned; tasks may fail to achieve their objective and need to be redone; priorities can change. The key to handling these disruptions is to make plan repairs with minimum disruption to the existing plan or schedule. Disruptions are costly to an operation and in many instances the hidden costs can be substantial. This is particularly true in the space exploration domain where substantial resources and planning are required up-front just to pre-position personnel and equipment for future missions. The more stable a plan, the less investment in stockpiling, repositioning of personnel and equipment, and transportation of material. Thus, repairing a plan with minimum disruption is always preferable to re-planning from a clean slate.

Our approach to plan repair involves three techniques capable of tackling repair problems of increasing difficulty. The *continuous technique* repairs schedules by sliding the start times and durations of tasks, taking advantage of slack time. The *local repair* technique repairs schedules by making local changes to the schedule, substituting available resources for resources that have become unavailable and by swapping tasks and resources to accommodate more complex changes. The *replacement* technique involves re-running the scheduler to tackle disruptions that cannot be achieved with the less disruptive techniques.

Continuous repair will be engaged when tasks slip. A constraint network will be set up with the earliest start and latest finish times of all tasks. Constraint propagation techniques will be used to propagate the effects of slippage on the start and finish times of other tasks. If the constraint network has a solution it means that the schedule can absorb the slippage, and no further changes are necessary. If the constraint network has no solution it means that a more radical change to the schedule is needed, and one of the other two techniques must be engaged. The advantages of using constraint propagation are that (1) it is very efficient, and (2) it can determine conclusively whether more disruptive changes are needed to repair the plan. The user can thus be informed about the severity of the problem.

Local repair works by selecting a broken task and making a change that minimally disrupts other tasks. The process is repeated for a number of iterations until all tasks are scheduled or a certain limit is reached. When the limit is reached, the process is restarted and repeated, including randomization to select different tasks and different changes to the tasks (e.g., picking a different resource). This technique is similar to the very successful WALKSAT techniques used to solve Boolean satisfaction (SAT) problems.[4] Our preliminary experiments with WALKSAT to solve scheduling problems confirm the effectiveness of this technique compared to systematic or heuristic search techniques.

Replacement is the last resort. The local repair technique is not guaranteed to find a solution if one exists, and cannot conclusively prove that no solution exists. From a practical point of view this is unimportant: scheduling problems are NP-hard so no technique can provide a conclusive answer in a practical amount of time. If the local technique fails to find a solution, it will be assumed that the given collection of tasks cannot be rescheduled using the given resources. In that case, the system will redo the schedule, thus considering changing the collection of tasks to be performed.

The benefit of using the three techniques to address plan disruptions is that repairs to schedules are minimally disruptive. Repairs are addressed first by sliding tasks, then by swapping resources and tasks, and finally by considering a different collection of tasks.

The main risk of this approach is that the local repair techniques may not be effective when the effect of global constraints such as crew-day and crew-rest is strong. In such cases a local change to a task may have far-reaching effects on many other tasks, making it difficult for the local search techniques to converge. In our experience with CACE we observed that on typical schedules the effects of such constraints were weak in the sense that these constraints, although very important, were easy to meet. However, on surge schedules that involved scheduling many more missions than usual, the pilots became critical resources and the crew-day and crew-rest constraints became difficult to meet.

We will run experiments with different types of schedulers to determine the effectiveness of the different techniques. Our expectation is that the local techniques will be effective in the typical cases where there is reasonable slack and redundancy to accommodate typical failures. In the extreme cases, when resources are pushed to the maximum there is no easy way to repair plans, and drastic revisions are necessary.

D. Probabilistic planning and scheduling

If you knew when equipment was going to fail, it would be relatively easy to schedule maintenance within a mission plan. The problem is that we don't know when equipment will fail. Prognostics help; however, they can only predict when a future failure will occur within a pre-defined time range with a given probability. The challenge is to help decision makers manage the choice between not doing maintenance or doing very little of it and thus greatly risking mission failure, or undertaking a high level of maintenance, along with its associated costs, but increasing the chance of mission success.

When equipment fails, the tasks that use the equipment will either fail or produce degraded results. With prognostics, one can assign probabilities to the equipment failures and consequently assign probabilities for task failure or degradation. One can, in principle, compute contingency plans that insert required maintenance tasks at different points in time in order to minimize the probability of failure. One could then compute an expected measure of merit to compare the different contingency plans. The contingency plans would probably produce a lower measure of merit compared with the original plan given the added maintenance burden. However, the likelihood that the contingency plans could be carried out would be higher. The situation is further complicated since different equipment may have different failure probabilities, leading to a combinatorial explosion of contingency plans. In addition, boiling down the comparison of contingency plans to expected values of some measure of merit hides essential qualitative elements of the decision from the user.

We propose to take a conservative and qualitative approach to this problem. Our approach produces safe and easy-to-understand recommendations for users. We propose to classify the consequences of equipment failure into different categories and act as follows:

Catastrophic consequences are those where equipment failure could harm people or valuable equipment. In such cases the affected resource will be marked as non-operational and a maintenance task added to make it operational again. The plan repair techniques explained in the previous section will be invoked resulting in a new safe schedule.

Non-catastrophic consequences are those that cause a task to be delayed or rescheduled at a later time without incurring any safety issues. The issue is to determine when to schedule the corresponding maintenance task. The maintenance task could be inserted before the task that uses the resource, or at a later time, running the risk that the task fails during execution. The system will investigate both possibilities as follows: It will run the repair scheduler with the maintenance task before the task that uses the resource and determine the level of disruption on the schedule—continuous, local repair, or complete reschedule. It will run the repair scheduler with the maintenance task after the task in question but in addition book a spare resource in case the resource in question fails. It will again determine the level of disruption and then pick the option that causes the least level of disruption. Ties will be broken in favor of inserting the maintenance task before the task that uses the resource.

The result is a conservative system that avoids unsafe behavior. It performs maintenance operations as soon as possible unless there is enough excess resource capacity that allows postponing maintenance activities that would otherwise be disruptive to the schedule. The different options can also be presented to the user, who may override the system selection. An additional benefit of this scheme is that it is computationally tractable. The number of contingency plans that the system investigates is linear with the number of tasks that use resources for which the system has generated prognostics.

E. Federated reasoning about location

CMMD tools must incorporate spatiotemporal reasoning across organizational structures. The vision for the future of space exploration requires that many different operational entities cooperate to carry out the overall mission. Sometimes, these entities will be separated by physical location. Consider a lunar or Martian base station surrounded by several satellite sites, each performing specific missions. In all cases, these units will be reliant upon one another in order to successfully carry out their goals, sometimes for their very survival.

This inter-unit collaboration requires the occasional exchange of tasks and/or resources between units. Local schedules derived from local goals must be supplemented with or refined to accommodate such external support requirements. Many times these interdependencies will involve the concept of physical location, and their feasibility rests upon the costs and difficulty of an action based on the relative location of the resources involved.

This is to some degree a modeling problem; obviously, the system's model of the world must include representations of an object's current location and its associated organizations as first class concepts. This capability also involves aspects of planning. If for example, one organization lends another organization a high-utility resource that it eventually expects to be returned, there is the logistics problem of how best to plan on transporting the resource to and fro.

We propose an architecture based on a network of nearly-independent decision support tools that can each reason about the possibility of forming partnerships or *federations* with neighboring entities, along with the implications of any constraints imposed by the distances separating them.

In our previous work, limited support for symbolic location was integrated into the CACE tool suite. For example, in response to a request for an aircraft in a sortie originating from a remote location, the system recognizes that this sortie must be associated with a series of related sorties that define an effective means of accomplishing local goals, plus transporting the aircraft to the target location. Furthermore, the aircraft selected for assignment to this aggregation of missions will be in locations with limited maintenance capabilities and therefore, the system must ensure that no significant scheduled or predicted maintenance activities will become due while it is away from base. This may entail *grooming* the aircraft to perform several maintenance items before the initial launch. Also, depending on the duration of the itinerary (both in calendar days and in flight-hours) some candidate resources will be deemed either ineligible or non-preferred candidates. Once an assignment is decided upon, the system respects that the designated air vehicle is no longer able to contribute to local production until after its scheduled return and possibly some resultant maintenance is performed.

We propose to extend this basic concept with more complex models of location and modes of transportation. In CACE, a predefined set of commonly used airbases allowed for the simplification of modeling through symbolic locations. In CMMD, we expect the requirements to be more general (e.g., please send a rover to this location tomorrow with the drill bit you agreed to let me have yesterday). Also, in CACE there was only one mode of transportation available for consideration by the scheduling algorithms. This system cannot, for example, plan to ship an engine by freight. A requirement to move an engine could only be realized by scheduling a flight with the aircraft that houses that engine. This is an obvious shortcoming, as the available modes of transportation will be much richer in CMMD. Finally, we plan to integrate a limited planning capability that, using models of modes of transportation, selects the best transportation mode and best teaming arrangement with external organizations.

Our current thinking is to implement this feature using an arbitration scheme. Individual CMMD nodes would have the capability to recognize failures and/or undesirable levels of risk in the local plan, which could possibly be resolved through cooperation with certain external sites. An arbitrator could then be invoked. These arbitrators would determine, request and consider the relevant constraints, goals, and intermediate plans from each contributing organization. The scheme will be evaluated by comparing the effectiveness of coordinated plans and plans developed in isolation.

V. Attaining Affordability

Our scope is not to consider program affordability of the Space Exploration endeavor as a whole but rather:

1) To have CMMD do for NASA what CACE has done for USMC and JSF in terms of affordability enhancement through optimization of mission operations and maintenance cycles.
2) To leverage to the fullest extent possible and to influence the design and development of the system health and logistics infrastructures available within ESMD's CTS thereby maximizing sustainability.
3) To serve as a pathfinder in developing operational and maintenance metrics that will allow program managers to make informed choices in trades between cost of technology investments in systems health features and logistical infrastructure support and cost of ownership (operations costs).

We have addressed how we are going to accomplish 1) and 2) in the previous section on technical approach through the application of five key technologies. Through yearly tech maturity demonstrations we will have metrics on the "with" and "without" cases of use of these technologies for the targeted NASA scenario, ISS EVA ops and maintenance. From this experience the impact of the use of such technology for attaining program affordability in the area of meeting requisite supply chain logistics margins and redundancy will be possible. The table in figure 2 below is a roadmap of beneficial impacts of our technologies on system health features such as IVHM prognostic drivers and logistical information/infrastructure features such as system-wide usage based tracking of components. The Xs identify where leverage of these features, of which examples are found in DoD's JSF and USAF's Defense Repair Information Logistics System (DRILS), provide a leverage or "tipping points" that if coordinated and reconciled as discussed in our technical approach could lead to revolutionary increases in mission supportability over what these features provide separately.

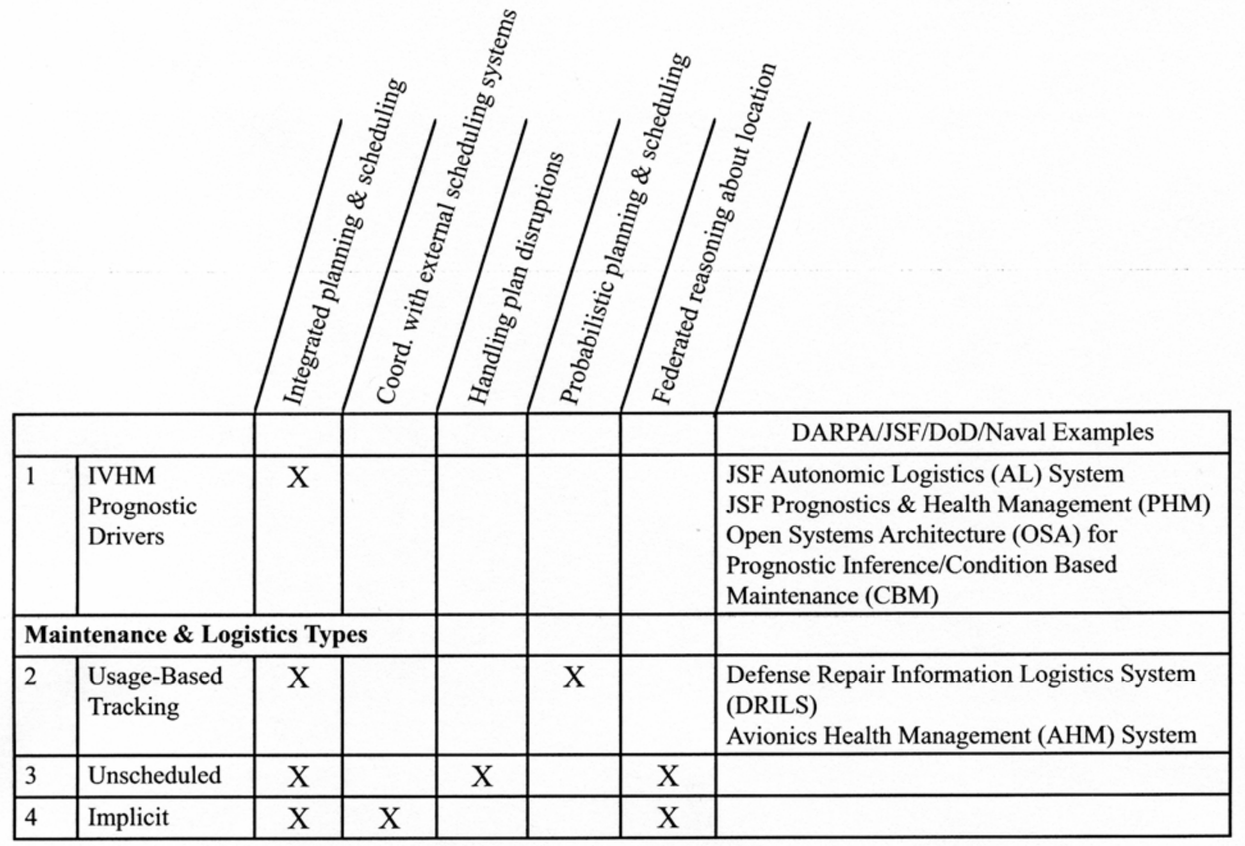

		Integrated planning & scheduling	Coord. with external scheduling systems	Handling plan disruptions	Probabilistic planning & scheduling	Federated reasoning about location	DARPA/JSF/DoD/Naval Examples
1	IVHM Prognostic Drivers	X					JSF Autonomic Logistics (AL) System JSF Prognostics & Health Management (PHM) Open Systems Architecture (OSA) for Prognostic Inference/Condition Based Maintenance (CBM)
Maintenance & Logistics Types							
2	Usage-Based Tracking	X			X		Defense Repair Information Logistics System (DRILS) Avionics Health Management (AHM) System
3	Unscheduled	X		X		X	
4	Implicit	X	X			X	

Figure 2: Roadmap of Beneficial Impacts of CMMD technologies

The CACE software has been fielded at USMC Marine Aircraft Group 13 (MAG-13) and aboard ships participating in operations in the Middle East and the Pacific. CACE handles the real-life issues of practical and complex operations. It efficiently balances a vast number of competing constraints ranging from pilot skills, suitability of available equipment, pilot skill acquisition and degradation, crew day/rest and a variety of other safety constraints, lunar light constraints, aircraft turnaround restrictions, calendar, usage and corrective maintenance activities, maintainer skills, and the tools required for maintenance operations. The CACE software supports construction of schedules for future operations as well as the schedule-repair activities necessary during schedule execution.

In addition to addressing a variety of competing constraints when generating operation and maintenance schedules, the CACE software also addresses the key areas of operational efficiency, safety, and enhanced planning opportunity. Improvements in these areas were measured through the following five metrics:

- Task elimination
- Task reduction
- Increased situation awareness
- Improved decision support
- Enhanced planning

Figure 3 below summarizes performance data collected during on-site testing and fielding of CACE in USMC Marine Aviation environments.

CACE Component	Metric	Result	How Evaluated
Ops Advisor (Operations Scheduling)	Task Elimination	80% reduction in manual tasks	Measured
	Task Reduction	Approx 30 minutes to generate a schedule that used to take 6 hours	Observed
	Situation Awareness	Continuous, iterative analysis, 100x faster	Observed
	Decision Support	Alternative futures, expanded planning horizon	User feedback
	Enhanced Planning	Rapid repair & risk reduction	User feedback
Maintenance Advisor (Maintenance Scheduling)	Task Elimination	80% reduction of manual tasks	Measured
	Task Reduction	97% of scheduling	Measured
	Situation Awareness	Continuous, iterative analysis	User feedback
	Decision Support	Alternative futures, expanded planning horizon	User feedback
	Enhanced Planning	Early warnings, rapid repair, minimize risk	User feedback

Figure 3: CACE Performance Data

For CMMD to have as compelling an impact for ESMD's CTS as CACE's impact for DoD's programs it is crucial to recognize that CACE has had a rich system health and logistics infrastructures within which to operate. For instance, the Joint Strike Fighter (JSF) program was driven to create this rich infrastructure out of necessity to lower long-term operational (ownership) costs and to maximize mission effectiveness.

The question that remains is how to meaningfully gauge the potential return on investment of support technologies such as IVHM prognostic drivers that keep trending information and generate "maintain me" and remaining life information. It is crucial to recognize that such features are not part of CMMD; CMMD's role is to leverage such features so as to maximize coordination and reconciliation across multiple system boundaries. To assume that these systemic attributes can be incorporated as part of CMMD will diminish its flexibility and create application bloat in deployments of CMMD on account of the systems engineering work-arounds typically found when the initial set of requirements defined and implemented are not adequate and are not created with extensibility in mind. This leads us to 3) above, that the experience gained through the use of CMMD in the NASA domain could be used to gauge NASA's need to invest in these features which are already found throughout DoD programs. Our success with CACE depended on such system health features and logistics information infrastructures in no small way. Simply put, without such features there will not be much to leverage synergistically. This need to create rich systems health and logistics infrastructures for NASA's Space Exploration program now exists, although historically it hasn't been the case.

For example, NASA's use of IVHM in human flight systems as part of its systems health portfolio has been limited to hardware instrumentation and sensors instead of fault location and quantification algorithm development.[5] Furthermore, the power allocation for IVHM instrumentation often is very small and the risk of a failure in IVHM instrumentation that could lead to cascading failures in other systems has often resulted in the IVHM component being designed and implemented as a second class citizen to other flight critical system components. This leads to the perception that humans-in-the-loop are the only way to perform systems health monitoring and status estimation adequately. Also the complexity and high cost of the vehicles and their payloads begs for human-in-the-loop flight controller monitoring and supervision. Now that the concept of *sustainability* is central to the development and deployment of NASA's space exploration systems complexity and high cost of space exploration vehicles will have to be designed to be more manageable, more affordable through the use of intelligent, fail-safe, autonomous components, systems, and system-of-systems. We now discuss the kind of health system features already found in practice both in government and industry and how CMMD's technologies could leverage these. If NASA had the following health system features suitable for its exploration needs it would allow CMMD to achieve a system-of-systems impact comparable to, if not greater than, that which CACE has done for DoD.

A. IVHM prognostic drivers

1. Autonomic Logistics System

The DoD's Joint Strike Fighter program has matured a far-reaching system-of-systems supportability concept in its Autonomic Logistics (AL) System.[6] The core idea of the AL system is based upon the autonomic nervous system of the human body where intelligent response to internal stimuli is embedded within the system itself. One of the key enablers of realizing this concept is JSF's version of IVHM which is referred to as Prognostics and Health Management (PHM). One of the cornerstones of the AL system is smart and reliable aircraft which means these are designed from ground up with PHM as a first-class-citizen component. PHM capabilities include built-in automatic (a) anomaly detection, (b) fault identification and fault isolation, and (c) prognostics while aircraft are operating in their theatre. One affordability tradeoff noted by Hess[6] is that in trying to achieve mission effectiveness, which for JSF translates into mission reliability and sortie generation rate goals, through reliability engineering alone leads to unaffordable aircraft. Affordable mission effectiveness may be met by a combination of equipment reliability, reconfigurability and redundancy. The AL system consists of five elements: reliable aircraft, autonomic logistics information system, advanced training and maintenance technician environments, and infrastructure.

2. Avionics Health Management (AHM) System

Advanced Health Management Systems (like the ones described by Byington[7] et al., and Atlas[8] et al.,) provide a functional framework for anomaly detection, diagnostics (i.e., fault detection and isolation), and prognostics. Individual components and subsystems may have their own Built In Tests (BIT) but these are typically insufficient to decide what maintenance needs to be done. An on-board AHM system combines evidence from multiple sources to determine what maintenance action needs to be performed. If the failure mode is successfully isolated by the diagnostics or predicted from precursors by the prognostics system, then the needs for a maintenance action is clear. If anomalies are detected but the diagnostics/prognostics is unable to pinpoint the source, then maintenance needs to be scheduled that requires active diagnostics testing. Thus, an on-board AHM system that combines multiple sources of evidence and takes into consideration cross-subsystem interactions is an invaluable resource for the maintainer and thus for the decision support system scheduling the maintenance.

3. Open Systems Architecture (OSA) for Prognostic Inference/ Condition Based Maintenance (CBM)

Another key building block to leverage will be the use of the Open Systems Architecture for Condition-Based Maintenance (OSACBM) for prognostic inference. The OSA consists of a seven-layer architecture consisting of sensor modules, signal processing, condition monitoring, health assessment, prognostics and presentation.[9] These layers have a detailed framework specification that have been worked out by a consortium of industry and government partners. The framework serves as a way for technologists and technology providers to organize system requirements and to architect a system design in such a way as to mitigate the usual integration issues found in systems with intense information-sharing needs. Some examples of these issues are 1) stove-piping of information within a system/enterprise, 2) the unintended creation of non-standard standards driven by near term expediency, and 3) cost overruns driven by iterative work-arounds necessary to overcome technology component incompatibilities uncovered late in the systems integration phase. The first two of these impact long-term flexibility, thereby diminishing the benefits of a spiral lifecycle development approach, and the last one destroys affordability.

The use of such architectures would provide clean interfaces to underlying system health features that would lower the cost of deploying CMMD technologies which would be responsible for coordinating and reconciling the big picture, or at more of a strategic system-of-systems level, of maintenance and operations across the mission theater level.

B. Maintenance usage-based tracking and logistics infrastructure

Conventional reasoning leads us to the hypothesis that as systems age their performance degrades. In terms of program affordability the desirable situation is to anticipate, or predict this process as it occurs so as to budget enough provisions necessary to maintain repair and replacement parts at rates that meet mission safety margins and redundancies. The Defense Repair Information Logistics System (DRILS) provides the capability of tracking possible detailed usage at the lowest repair level, Line Replace Unit (LRU).[10] The histories of an individual LRU's repair and test are kept in databases from Automated Test Equipment (ATE). Kirkland[10] makes the point that by adding real-time environmental stress monitoring capabilities in combination with these histories it would then be possible to create trends and degradation characterizations that can anticipate which units are likely to fail soon. Creating such a system for space exploration systems would allow for long-term optimization of cost of ownership, hence improving program affordability, especially when coordinated and reconciled with other system-wide needs for multiple equipment types as we have discussed via CMMD technologies.

C. Unscheduled maintenance management

Unscheduled maintenance occurs in large-scale systems regardless of how much oversight has been used to anticipate the unexpected. Current approaches deal with unscheduled maintenance by providing extra maintenance margins in anticipation of unscheduled events—clearly a sub-optimal approach. Keller lists unscheduled maintenance as one of the seven deadly sins against *supportability*.[11] The authors raise this point as part of a systems engineering process for determining the cost/benefit of prognostics, i.e., the impact of prognostics, on operational system ownership of, for example, aircraft components. Using this process one can in principle design the likelihood of unscheduled maintenance for a component to be a very small amount, given cost is no object. However, as noted in step two of the paper by Keller, et al., the list of candidate critical components is based on requirements such as safety. The assessment of what engineering features are needed to address safety margins, in particular for crewed systems where ASARA reigns, rests to some degree on engineering judgment rather than some guaranteed solution derived from closed analytical calculation. Given this uncertainty and the uncertainty inherent in large-scale physical systems found, for instance, in sensors as signal processing noise, unscheduled maintenance management will occur in the form of caution and warning notifications sent to the crew where they will need to decide what course of action to take. As a consequence, smooth integration of human–machine interaction is essential when it comes to making decisions and taking actions based on decision support analysis tools. CMMD technologies have the capability to provide mixed-initiative planning and scheduling where the system can explain to the user how it arrived at the current context, what intervention needs to be decided upon, and provide the user the capability to drill down and look at all information under its span of coordination and reconciliation.

D. Implicit maintenance management

One of differences in domains between DoD's JSF and NASA's CTS is that there is a clear demarcation between the roles of a military pilot and a maintenance engineer, whereas the crews in NASA's CTS will need to be more jack-of-all-trades in that they will be directly responsible for vehicle maintenance while on a mission. The reason is that there is little if any time for a jet fighter pilot to work vehicle servicing issues while engaged, for example, in combat and the aircraft is always meant to return to a servicing depot in a short time. In contrast, the environment for a CEV crew is more akin to that of a submarine that is hundreds of thousands—in case of Mars, millions—of miles away. In this situation, the crew has to be capable of administering to all of the maintenance needs in situ because there is no nearby service depot. In this situation where on-board, on-site mission supportability is of paramount importance because of remoteness there will be a tendency to maintain margins and redundancy through the generalization of use of equipment types. In addition, crew resourcefulness may result in equipment types being used in unexpected ways. Such re-purposing of equipment, if not tracked and accounted for adequately, may result in strained, if not broken, logistical supply lines that could impact mission feasibility. Coordination of general purpose equipment as already noted will be a technical challenge but one that must be done in order to deal with the unforgiving environment of space. CMMD's integrated planning and scheduling and federated reasoning about location technologies will be brought to bear to solve the problem of coordinating operations and maintenance by using the maximum amount of knowledge about the state of the equipment as well as mission goals. The coordinated operations/maintenance cycles will result in missions that will have maintained equipment supporting them, and no operations will be scheduled that cannot be supported by equipment.

VI. Conclusion

Our hypothesis is that program affordability can be impacted dramatically through the use of system coordination and reconciliation technologies such as CMMD. We have discussed how through the re-targeting and refinement of our technologies of 1) integrated planning and scheduling, 2) coordination with external scheduling systems, 3) handling plan disruptions, 4) probabilistic planning and scheduling, and 5) federated reasoning about location, we are poised to repeat and exceed the DoD-sponsored CACE results found in its domains for NASA's domains. Fundamentally, the kinds of scheduling and planning problems found in both of these have a great deal in common and we are confident that we can readily leverage prior experience and expertise for NASA's benefit. As a corollary to our hypothesis, it will benefit program affordability for NASA to take on an acquisition strategy for system health and logistics information infrastructure that is as far reaching as what DoD has done for its programs such as Joint Strike Fighter. The results of such acquisitions, where systemic features as discussed are embedded in the Constellation Transportation System, will give CMMD more with which to optimize maintenance and operations cycles for the continual attainment of program affordability across multiple spirals of space exploration campaigns while ensuring day-to-day practical *supportability*.

References

---

[1] Code T Office of Exploration Systems Enterprise Project Constellation Work Breakdown Structure, retrieved January 14, 2005, http://exploration.nasa.gov/documents/documents.html

[2] Human and Robotic Technology(H&RT) Formulation Plan, retrieved January 14, 2005, http://exploration.nasa.gov/documents/documents.html

[3] Dixon, H., and Ginsberg, M. : "Inference methods for a pseudo-Boolean satisfiability solver", The Eighteenth National Conference on Artificial Intelligence (AAAI-2002).

[4] Selman, B., Kautz, H., and Cohen, B., "Local Search Strategies for Satisfiability Testing.", Final version appears in Cliques, Coloring, and Satisfiability: Second DIMACS Implementation Challenge, October 11-13, 1993. David S. Johnson and Michael A. Trick, ed. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 26, AMS, 1996

[5] Sirkis, J., Childers, B., Melvin, L. et al, "Integrated Vehicle Health Management(IVHM) on Space Vehicles: A Space Shuttle Flight Experiment", *Key Engineering Materials Vols. 167-168(1999)* pp.273-280,Trans Tech Publications, Switzerland

[6] Hess, A., Calvello, G., Dabney, T., (2004) "PHM a Key Enabler for the JSF Autonomic Logistics Support Concept", *2004 IEEE Aerospace Conference Proceedings*[CD-ROM] IEEE Catalog Number:04TH872C ISBN:0-803-8156-4

[7] Byington, C., Kalgren,. P.W., Dunkin,. B.K., Donovan, B.P., "Advanced Diagnostic/Prognostic Reasoning and Evidence Transformation Techniques for Improved Avionics Maintenance", *2004 IEEE Aerospace Conference Proceedings*[CD-ROM] IEEE Catalog Number:04TH872C ISBN:0-803-8156-4

[8] Atlas, L.; Bloor, G.; Brotherton, T.; Howard, L.; Jaw, L.; Kacprzynski, G.; Karsai, G.; Mackey, R.; Mesick, J.; Reuter, R.; Roemer, M.; "An evolvable tri-reasoner IVHM system", Aerospace Conference, 2001, IEEE Proceedings. ,Volume: 6, 10-17 March 2001, Pages:3023 - 3037 vol.6

[9] Provan, G., "An Open Systems Architecture for Prognostic Inference during Condition-Based Monitoring", *2003 IEEE Aerospace Conference Proceedings*[CD-ROM] IEEE Catalog Number:03TH8652C ISBN:0-7803-7652-8

[10] Kirkland, L., Pombo, T., Nelson, K., "Avionics Health Management: Searching for the Prognostics Grail", *2003 IEEE Aerospace Conference Proceedings*[CD-ROM] IEEE Catalog Number:03TH8652C ISBN:0-7803-7652-8

[11] Keller, K., Simon, K., Stevens, E., Jensen, C., Smith, R., Hooks, D.,"A Process And Tool For Determining The Cost/Benefit Of Prognostic Applications", *AUTOTESTCON Proceedings, 2001, IEEE Systems Readiness Technology Conference,*: 532-544.