STARDUST: IMPLEMENTING A NEW MANAGE-TO-BUDGET PARADIGM

Kenneth L. Atkins, Project Manager
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91011
kenneth.l.atkins@jpl.nasa.gov

Bredt D. Martin, Planning & Control Mgr.
Martin Consulting Group, LLC
25517 Paine Circle, Stevenson Ranch, CA 91381
bredt.d.martin@jpl.nasa.gov

Joseph M. Vellinga, Flight System Mgr.
Lockheed Martin Astronautics, Denver, CO
joseph.m.vellinga@lmco.com

Rick A. Price, Planning & Control
Lockheed Martin Astronautics, Denver, CO
rick.a.price@lmco.com

Abstract—STARDUST is the Discovery Program's fourth mission. It follows Lunar Prospector, Mars Pathfinder and the Near Earth Asteroid Rendezvous (NEAR) mission. Launched on February 7, 1999, the STARDUST flight system will collect comet samples during a 6 km/s flyby of Comet Wild 2 on New Year's day, 2004, and return the samples to Earth in January 2006. Enroute to the comet, STARDUST will also attempt to collect samples of interstellar dust.

The Jet Propulsion Laboratory (JPL) provides project and mission management with Lockheed Martin Astronautics (LMA) as the industrial partner for the flight and ground systems. LMA made strong use of developments in the Mars Global Surveyor (MGS) and the Mars Surveyor '98 projects preceding STARDUST. Under the stringent cost-caps of the Discovery Program, the STARDUST development was completed on time with nearly $2M remaining reserves to reprogramming into the flight phase.

The STARDUST management team aggressively worked to achieve this control through the use of Total Quality Management (TQM) and reengineering principles, and commercially available software tools. The approach was to develop project-to-project interfaces exploiting parts stores and common procurements, shared staffing and shared facilities. Inheriting parts, hardware, software and designs leveraged dollars efficiently to accelerate the development, while staying inside a constrained budget. Additionally, to achieve the required level of time efficiency and budget control, a new level of communications and data handling (read excellent Management Information System [MIS]) is mandatory. Finessing the rigidity of traditional Performance Management (or Measurement) Systems (PMS) and institutional/corporate cultures requires a new way thinking and a cheerleader aggressiveness. STARDUST organized toward the Integrated Product Development Team (IPDT) concept as a central feature. This was matrixed into a dedicated Product Development Office (PDO) at LMA. The PDO served the leveraging goal by serving two projects, STARDUST and Mars Surveyor '98. It avoided duplicate project-unique personnel structures and offered cost benefits to each project.

This paper provides details and example metrics characterizing the aggressive application of the *design-to-cost* paradigm and implementation by the STARDUST management team to achieve success under the Discovery Program budget constraints.

INTRODUCTION

The Discovery Program reflects a new way to continue the legacy of the Mariners, Voyager, Magellan, and Galileo in deep space exploration. Discovery is changing the way NASA does business. It is a central element in a *complete culture change* for planetary exploration and space science. Discovery's goal is to achieve results *faster*, *better*, and *cheaper*. It will be more effective, do more with less—specifically, carry out planetary flight missions with highly- constrained total cost.

STARDUST was selected from a pool of 28 proposals in 1994. It becomes the fourth mission in the series, preceded by: Near Earth Asteroid Rendezvous, Mars Pathfinder, and Lunar Prospector.

Historically, planetary missions evolved to large, complex platforms with up to 14 scientific experiments and price tags of up to $2 billion. These missions endeavored to do remote-sensing and in-situ investigations on extremely stringent diets of power, mass, and volume. The struggles in the scientific community to be one of their cramped passengers were difficult and frustrating. With their high price tags, such missions are difficult to afford.

STARDUST is in the process of *reversing* the paradigm. It is a *sample return* mission whose fundamental premise is to bring the essence of the solar system, material from a comet, home! With samples back on Earth, literally *hundreds* of experimenters can participate in analyzing the thousands of particles returned to Earth. They can apply existing instruments—with relatively unlimited power, mass, and volume constraints—which are operational in the finest labs and universities. This will allow participation in solar-system exploration by a broad community. And the opportunity is offered at a Discovery price, less than 10 % of the traditional approach!

STARDUST is the *first* program approved for return of material from a solar-system body since the Apollo and Luna sample-return missions of the 1970s and, more importantly, the *first ever* program for return of material from a comet. As such, it becomes a model for planning follow-on sample-return missions to other planetary bodies. The simplicity and compactness of the Sample Return Capsule (SRC) should be very attractive to follow-on applications. Figure 1 shows the STARDUST spacecraft in its sampling configuration.

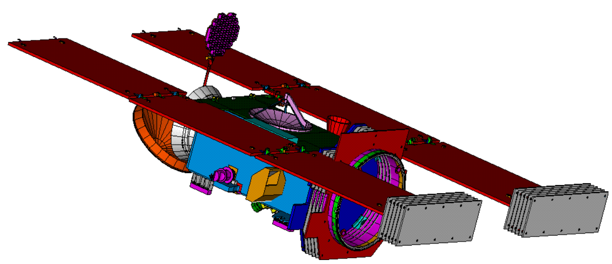

Figure 1. STARDUST Spacecraft

The major features of the STARDUST flight system are: the Sample Return Capsule (SRC), about a meter in diameter, shown open like a clamshell at the rear of the spacecraft, with the dust-collector grid deployed into the dust stream; the Whipple shields, consisting of two plates with Nexel™ curtains between to stop the high-speed particles from impacting sensitive spacecraft elements, shown at the front of the spacecraft and solar-arrays; solar-arrays, shown along each side of the spacecraft; and the Cometary and Interstellar Dust Analyzer (CIDA), to be provided by Germany. The flight system also carries an upgraded Voyager camera to provide optical navigation capability. The plan is to also use this camera for imaging the nucleus of the comet to a resolution an order-of-magnitude better than *Giotto* imaged *Halley*.

CONTINUOUS IMPROVEMENT AND REENGINEERING PRINCIPLES

To operate within NASA's better, faster, cheaper paradigm, and to meet Discovery Program requirements, the STARDUST team was challenged to achieve a very efficient program. Meeting this challenge entailed implementation of many new ways of doing business, which required a combined approach of adopting, changing, and inventing business processes. The paramount goal was, and continues to be, the implementation of *best business practices* throughout the project. How would this goal be met with a distributed team?

VIRTUAL COLOCATION: HOW IT WORKS

Doing business *globally* is becoming a necessity in today's business environment. The JPL teams in Pasadena had to be functionally intertwined with the LMA teams a thousand miles away in Denver. With the new role of the PI being in charge, it was essential for the PI, who resides in Seattle, to be able to fulfill his role on a very frequent basis from a location removed from JPL and LMA. Co-Investigator (Co-I) team members located around the country and in Germany must interact often with the other teams.

Through the use of commercially available software tools, and some not-so-commercially-available software tools, the team was linked via an Information and Communications System (ICS) detailed below. The ICS facilitates easy access and frequent communications among all team members, which has significantly contributed to the success of the project as a whole.

A primary benefit of the ICS is to enable the distributed team members to work together and share information as if separated by an office down the hall rather than a thousand miles away. The structure promotes team cohesiveness and open communications—there are no secrets across institution boundaries. Project budgets were defined and worked as one integrated team, and not as a customer-contractor relationship. This relationship proved very beneficial when initial baseline budget plans exceeded a funds available profile. A solution was jointly worked by the teams, and not merely thrown over the fence to the other party.

A second benefit of the ICS is savings on travel costs. In addition to the dollars spent, there is a substantial lost effective time factor, and additional stress on personnel in being away from home. STARDUST, as a matter of course, conducts its Monthly Management Reviews (MMRs) and other recurring reviews co-located "virtually", with no personnel travel required.

Software Tools

A critical decision made early in the program was to decide upon a set of multi-platform, commercial-off-the-shelf (COTS) software which would be uniformly used by all project personnel, regardless of location or affiliation, for the duration of the program. Prior experience had repeatedly shown that purported software translators never quite did the job 100 % of the time. Moreover, a requirement of translation prior to using the information stored in the files would inhibit communication among the team members— an unwanted result.

While word processing, spreadsheet and presentation software easily defaulted to Microsoft Office™ for various reasons including cost, availability, and ease of use, decisions affecting the program-control aspects of the project were more complex. LMA Flight Systems, on previous and other concurrent programs, used Microsoft Project™ for network scheduling, Microframe Program Manager™ (MPM) for financial data processing, including earned-value, and FastTrack™ scheduler for presentation-quality top-level schedules. The similarity of these programs and the availability of data made the decision to use the same software tools the logical choice, regardless of the potential superiority of other software products.

While Microsoft Project was one of the more popular network schedulers of JPL personnel, it was not the only product used. The advantage of using the same software as LMA far outweighed the advantages of any other project-planning application. The use of MPM for earned-value performance management was another, more difficult problem. JPL, as an institution, did not support any commercial earned-value management (EVM) software product, as a sponsor had never previously required the earned-value process. However, by fortuity, the STARDUST planning and control (P&C) manager had previously applied MPM to another project. Thus, the conclusion of using LMA's institutional-standard program-control software was a relatively simple solution. The team at JPL also adapted FastTrack as its high-level scheduler for convenience and continuity with LMA's company practice.

It is noteworthy that starting from a process framework did not drive the team to have COTS software entirely meet intended requirements. Rather, adjustments were made to the ideal processes to accommodate the functionality of the commercially available software, and adaptations were made as necessary to make the process whole, thereby meeting the *best business practices* goal. *"Good Enough"* became the master of *"better."*

File and Server Design

At the center of virtual co-location are one or more file servers. The STARDUST file servers hold an electronic library of all documents produced during the life of the project. A carefully planned structure of the files on the server is essential to ease of use and subsequent retrieval of information. Files must reside in folders that group naturally and allow simple navigation to reach a needed file. The objective is to avoid confusion and the feeling of being "lost in a maze." The design has the first-level folders align along functional and work-breakdown-structure (WBS) lines. For example, top-level folders exist for Business Management, Reviews, Project Engineering and Integration Team (PEIT), and NASA HQ. In general, all folders are fully accessible by project team members to facilitate flow of information between personnel. To maintain some confidentiality of information, certain folders are provided to limit access to members external to the project, e.g. NASA Headquarters, Outreach affiliates, and foreign scientists. Figure 2 presents a view of the STARDUST server directory-folders.

The STARDUST servers are configured to provide local access at both JPL in Pasadena and at LMA in Denver. Every 30 minutes, the servers mirror locally generated information to the server at the remote location via a dedicated T1 line. The purpose of this configuration is twofold. File transfer time, from the server to a desktop for local users, is relatively instantaneous compared to some expected delay over the Internet. In cases where megabytes of information are transferred, this efficiency is essential to the smooth operation of the server and provides incentives for personnel to use the system. A mirrored approach provides each generator of a large amount of information to have a complete set of its information at all times. In the event of a network failure between JPL and LMA, LMA and JPL personnel still have current information. The local users will have a copy of the remote information most recent before the network failure. The local access server also provides double-click file-execution capability. All team members removed from the JPL–LMA mirrored sites may access information via File Transfer Protocol (FTP).

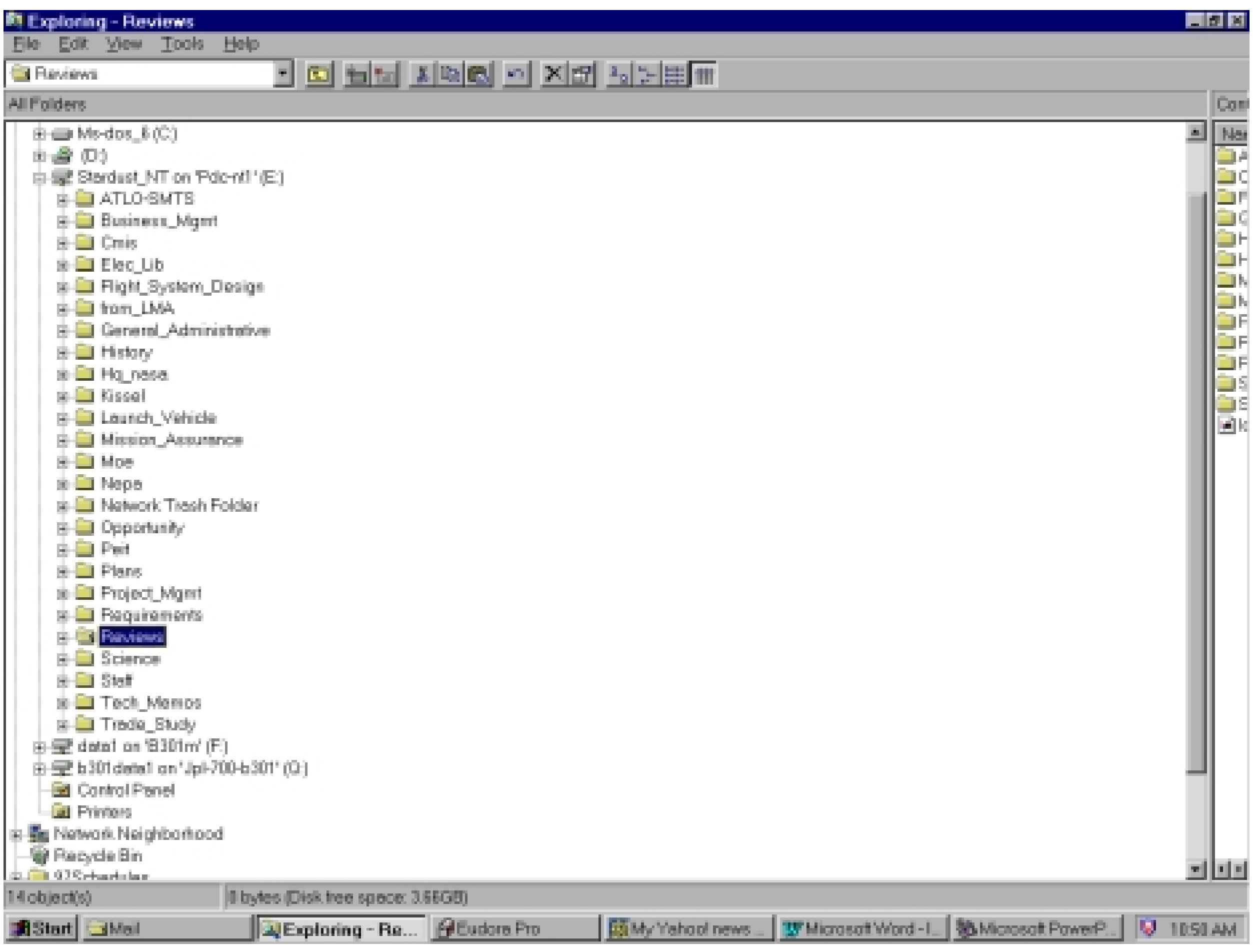

Figure 2. STARDUST Mirrored Server Directory Structure

Firewalls and Data Networks

At the beginning of the program, commercially available software did not exist that could provide secured mirroring of information between two remote sites. To compensate, LMA developed its own internal “store” program. It initiated pushing of data to the remote site each half hour, and pulling data from the remote site each half-hour, staggered 15 minutes from each push. LMA was required to initiate all communications to and from outside its “firewall” to maintain security for its information networks. This firewall system precluded other solutions, for example, operating Windows NT™ (NT) mirrored-servers over a wide area network (WAN).

An issue in providing server access was the ability to provide multi-platform access to information. The STARDUST project initially faced a situation where the UNIX servers at the local and remote locations ran an Appletalk™-only emulator shell, a “universityware” program. While the Appletalk-only access was sufficient to meet LMA’s need, JPL’s requirements were for multi-platform access, including Apple™, Windows 3.1™, Windows 95™ and Windows NT workstations. Initial attempts to use Windows directly with the UNIX operating system resulted in many scrambled files and frustration among the users. Going to a 100 % Apple-compatible user-set was not an option at JPL.

To solve this problem, a Windows NT server was placed in series with the UNIX server at JPL. The NT server communicated with the UNIX machine via Appletalk-emulators on both ends, which resulted in high-quality data being available to all platforms from the NT. But this arrangement caused additional problems. Because an automatic read/write of files to the NT from the UNIX machine was not possible within the operating systems, the ICS manager was required to manually push and pull information to/from the UNIX from the NT at least once a day. This operation lasted about 18 months. Local users were happy with high-quality information, but the practice resulted in additional workload for the ICS manager, and the arrangement was certainly not optimal. Recently, commercial software became available which emulates an NT server and which runs on a UNIX platform. LMA and JPL migrated to this new system and it works well.

Your Place or Mine (The Meet-Me Line)

In conjunction with the mirrored servers, the STARDUST project implemented a dedicated “meet-me” teleconferencing line. It allows up to approximately 33 concurrent calls. The teleconferencing line provides the audio portion of the virtual meetings, while files on the server provide the visual portion of the meeting.

Typically, Microsoft PowerPoint™ slides are prepared prior to a meeting, and the files are shared so each participant has ready access. The conferencing facility at JPL comprises a meeting room with a screen, a computer with video output connected to a ceiling-mounted projector, and a tabletop conference phone. Participants in meetings have dialed in from Russia, an airline in flight, and a Hawaiian phone booth. With a laptop computer and a modem, the system is globally accessible for participation *equivalent to being in-situ*.

Early in the project, a decision was made not to invest in videoconferencing facilities for three main reasons. First, the cost of each facility is in the neighborhood of about $50,000 and a minimum of two facilities would have to be installed at least one at JPL and one at LMA. Second, videoconferencing does not easily provide the operating flexibility from many locations worldwide, with each attendee participating on an equal basis. Third, having participated in numerous videoconferences, the value added by “virtually looking” at remote participants during the discussion was determined to be minimal.

IMPLEMENTING CHANGE – CULTURE CLASH

For many years numerous aerospace companies and JPL operated under the Apollo paradigm; that is, design and schedule generally were concerns paramount to cost. As a result, tools were not developed and used in tracking cost at JPL to the extent that a for- profit company would. For cost-account holders, penalties were severe for over-running planned cost, and near heroism was bestowed for under-running cost plans. Thus, a mindset developed (for self-preservation’s sake) to intentionally budget such that it was highly unlikely that a cost overrun would occur. A further mindset developed in which budgets were considered grant-like, and recipients would hoard them, whether likely to spend them or not.

Budget not spent in one year was insisted to be rolled-over into the next, rather than be recaptured in reserves. These conditions resulted in fantasy budgets that in many cases could not be related to real people and real tasks.

Given these operating conditions, program-control teams had great difficulty in planning for contingencies or understanding what reserves actually existed, because so much money was typically hidden in the planned budgets. Thus, projects operated under smaller fiscal reserves than necessary and certain decisions regarding expenditures on risk-reduction items could not be intelligently made.

In tracking planned cost and actual cost only, without the integration of schedule, it was difficult to tell whether planned work in fact had been accomplished or whether it was merely deferred. Thus, during a budget revision, it was more likely that re-planned prospective budgets would rise without prior warning. The real problem was that there was no *clear* correlation between work planned and work accomplished or between work accomplished and actual cost.

Others had implemented earned-value management (EVM) programs with limited success. Existing earned-value cultures included those who left the room at the mention of earned-value and those who staunchly advocated tri-service-certified methodologies. Our challenge was an opportunity to gain the benefits of having performance-measurement information to facilitate management of the program in a cost-capped environment, without enduring the distress and cost of adhering to the strict discipline of most traditional performance-measurement systems. How to do it *faster, better, cheaper*?

Our *tailored* approach has enough discipline to maintain baseline and data-integrity without the unnecessary restrictions typical of a tri-service–validated system. The approach was successfully tested during Phase B, and fully implemented at the start of Phase C/D development. Operationally, all significant variances were initially investigated by the program-control team, and then were addressed by the cost-account owners as necessary to explain the variations. This *tailored* approach to EVM helped enable the cultural change toward accepting it as a legitimate tool to help manage the program. The tailored EVM data proved to be a valuable indicator of true performance against baseline plans.

Integrating the Work Breakdown Structure (WBS)

One of the unique features of STARDUST's WBS was its full integration between JPL and LMA. That is, there are no overlapping WBS elements between the two enterprises, and they fit together as an integrated whole. Such a structure benefits the project in a number of ways. From an organization standpoint, the team members located remotely from each other are seen as part of a whole, rather than segregated by a particular affiliation or identity (ID) badge.

From a program-control viewpoint, there are numerous advantages to a unified WBS. Budgets and cost-accounts are uniformly identified at the same level of the WBS, regardless of the origin of the work. The WBS uniformly identifies schedules, associated with each WBS element. Project documents, regardless of origin, are uniformly identified with a WBS from whence it came. And, at the end of the cost-accounting period, earned-value can be rolled-up within the WBS structure without concern for misidentification of costs incurred.

Schedules

1. Detail vs. Intermediate or Top Level—A Communication Challenge—

A challenge in the maintenance of multi-level schedules is to insure that the information is internally consistent among the schedules, and effort is not duplicated in maintaining the schedules. The desired result is that the master information is contained in a single location, and other compilations of the schedule information are derivative, rather than duplicative.

Network-schedule software such as Microsoft Project, while capable of showing rollup information, is not geared toward the display of Level 1 (project level) information in a clean fashion. As described above, a natural selection was FastTrack Scheduler, already used by LMA institutionally to fulfill this need. While FastTrack provides additional flexibility not inherent in MS Project, it does require manual input. Clearly, a more optimum solution was needed, but wasn't available. A "make-shift" worked but not efficiently.

With MS Project, intermediate-schedules derived from detailed networks can be created by dynamically linking desired information from the detailed schedules into a separate schedule file. The information in the intermediate-schedule file is then automatically updated when the detailed schedule information is updated, and thus internal consistency is automatically maintained. But this method has a serious drawback, which impedes dissemination of the information. The dynamic linking process includes specific file directory information. This aspect prevents the transport of files from one directory to another without breaking the dynamic link, which requires manual input to repair. The virtual co-location aspect of STARDUST required that the files be easily transferred within and without each institution, thus the dynamic linking of the files was not a workable answer.

A solution to the problem of maintaining intermediate-schedules is found by using the multi-project capability of MS Project. Desired-to-be-displayed information in an intermediate-schedule was tagged in a common-text field of the detailed schedules. All detailed schedules may then be simultaneously loaded into MS Project, and all tagged tasks from the detailed schedules may be selected for display. In this manner, a one time only effort allows identification of desired tasks. The information is fully transportable by file name and independent of directory name.

2. Network Schedules—Critical-paths—

While MS Project is probably not the most flexible scheduling tool on the market today, it is relatively easy to use and operates equivalently on Windows and Apple computer platforms. A serious limitation in the product itself is its cumbersome and limited ability to link a task in one network to another task in a second network. These links are critical when numerous products are being fabricated in shops not necessarily under the control of the end-product holder. The key to the success of any networked schedule is the accurate modeling and control of hand-off points between task/ budget owners.

A semi-manual approach was developed to identify and constrain the known links (or hand-off points) within the separate network files. This approach included duplicating the hand-off points within the delivering and receiving networks. For example, when a delivery was agreed upon by both sides of an interface, it appeared in the delivering network as “Deliver XYZ Box to ALTO 68320.” It also appeared in the receiving network as “Receive XYZ Box from 64400.” In the delivering network, it is constrained as a “Finish No Later Than” activity type, and as a “Start No Earlier Than” activity type in the receiving network. These activities are also indicated as receivables or deliverables in a text field, and can be sorted on to provide additional management attention, if necessary. Managing interfaces this way is a bit cumbersome, but it proved effective. An advantage of identifying and constraining activities in this manner is that it provides early indications (via critical-path networking) when a hand-off point is in jeopardy, as if all of the individual networks were contained in a single database.

KEY METRICS

The STARDUST program-control paradigm was to not to implement a single tool to accomplish all goals—this approach typically results in many compromises which impede implementation of *best business practices*. Rather, the program-control teams developed a *suite* approach that included a number of programs and metrics working together to plan and analyze performance. This *suite* is detailed below.

<u>Event-Driven Performance Assessment Metric (PAM)—</u>

The program-control team residing at LMA generated and maintained approximately 30 detailed network schedules. One of the metrics used by the program-control team at JPL to monitor LMA schedule performance was a Performance Assessment Metric, or PAM, created from the detailed network schedules. Depending on the length of the task in the network, the PAM assigns a number of events to the task. For tasks of one month or less, start and finish events were assigned. For tasks with more than one month, events were assigned for each month during the task. This method avoided a one-month task being assigned the same weight as a six-month task. The events were then graphed cumulatively over the time-span of the project, and indicate a planned, late-finish, and actual completion of the events. The actual line should fall between the other two, and if the actual line falls at or below the late-finish, the metric indicates that schedule reserve is lost and critical-paths are likely affected. Figure 3 presents an example PAM.

Each month, the number of actual events is compared against the plan to yield a schedule-only–driven schedule performance index, which is then compared against the earned-value Schedule Performance Index (SPI) for a crosscheck. While the methodology appears to be somewhat arbitrary, it is interesting to note that the resulting schedule performance observed to date with this system has very closely mirrored the performance indices output from the PMS system.

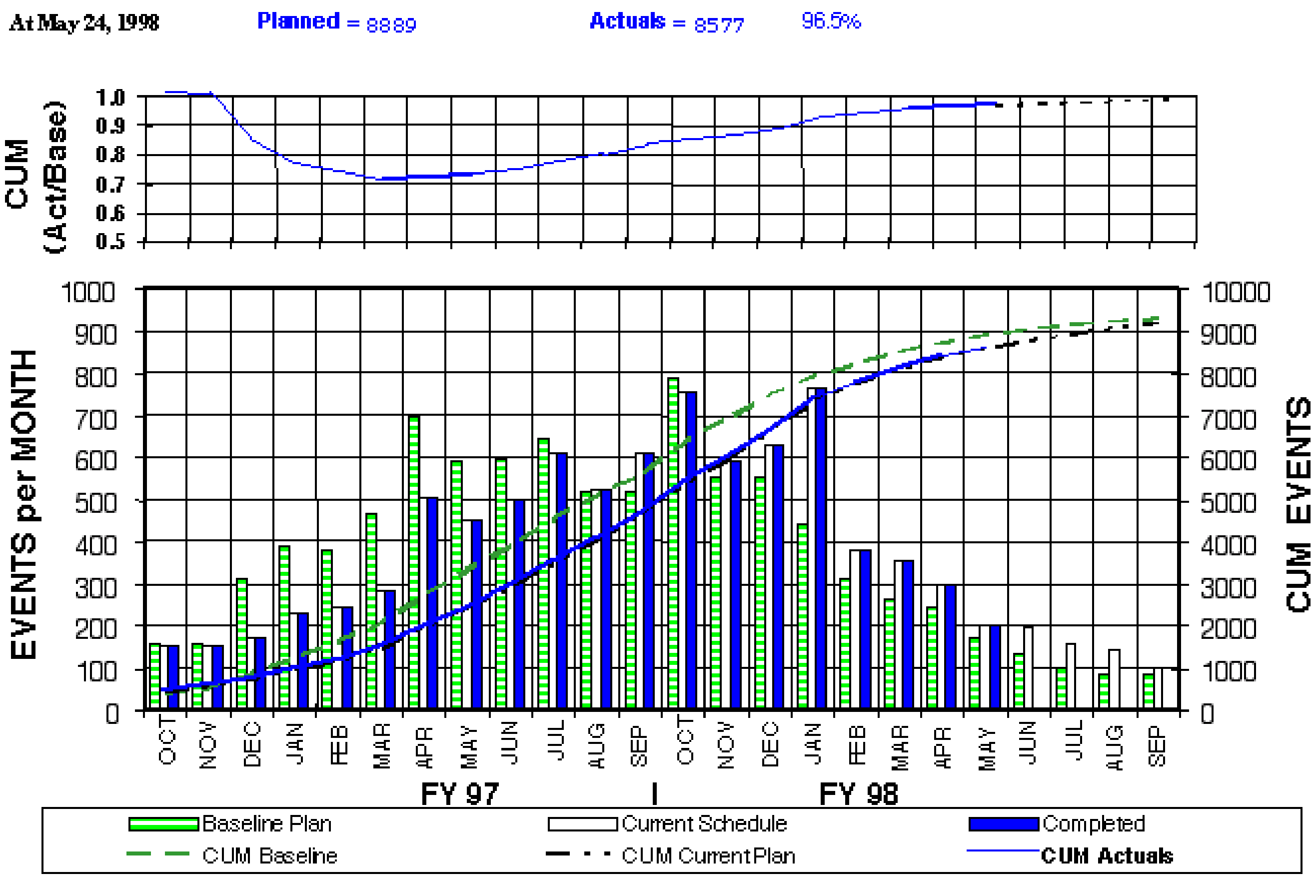

Figure 3. STARDUST Performance Assessment Metric

Earned-value—

The earned-value metrics are central to the understanding of the pulse of the program over time, i.e. true performance. In previous programs, a contractor would provide only a NASA form 533 to the customer on a monthly basis, which contained data that was at least a month old by the time it was received. With the common software tools used between JPL and LMA, evaluation of earned-value results were greatly simplified. In accord with the new openness and partnership between the two organizations, LMA provided their internal earned-value information in the form of an electronic file to the program-control team at JPL, usually within one week after earned-value sessions had been completed. The ability of the program-control team to identify potential problems early is thereby greatly enhanced.

Figure 4 presents a sample earned-value metric.

Financial Risk Management

1. Managing Reserves—

When faced with an essentially fixed-price program, effective management of financial and schedule reserves is vital to the survival of the project. The penalty of a cost-overrun of 15 % in today's NASA environment is cancellation. For Discovery Programs the 15 % will not be tolerated; the program starts with an agreement with NASA that the program will be completed for the commitment in the proposal, *and no more*. The result of a schedule stretch-out is a cost overrun, which effectively is cancellation.

For missions that are dependent on critical planetary trajectories in order to fit on the small launch vehicles permitted for Discovery missions, schedule slip is not an option.

The implementation of detailed budget plans is central to the understanding of the reserves posture. Detailed planning necessarily stimulates the planner to better understand the nature of the task, which generally results in fewer omissions that need to be covered by reserves at a later time.

STARDUST

COST PERFORMANCE MEASUREMENT

August 1997

Program Earned Value

FISCAL YEAR '9

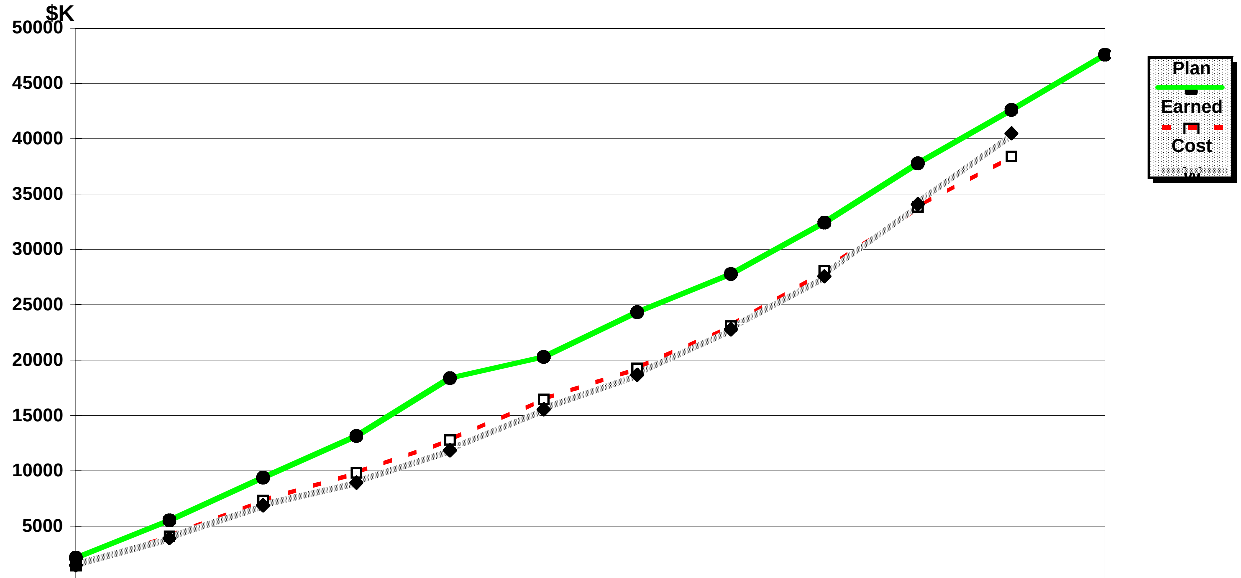

		Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
C	Plan	2155	5542	9386	13153	18378	20298	24340	27799	32435	37799	42628	47588
U	Earned	1448	4104	7337	9831	12793	16479	19273	23091	28079	33854	38396	
M	Cost	1470	3919	6881	8942	11869	15572	18683	22783	27574	34097	40479	

	Oct	Nov	Dec	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep
Schd Var	-707	-1438	-2049	-3322	-5585	-3819	-5067	-4708	-4356	-3945	-4232	
Cost Var	-22	185	456	889	924	907	590	308	505	-243	-2083	
Plan- Act	685	1623	2505	4211	6509	4726	5657	5016	4861	3702	2149	

Figure 4. STARDUST Earned-Value Metric

After assessment of the general design-maturity of the flight subsystems, many of which were essentially build-to-print of prior or concurrent programs, the team settled on a 10 % reserves floor guideline based on the program's cost-to-complete. This floor was intended to provide a bounded required-reserves level, and would not be violated except under extreme circumstances. Any sustained move toward the 10 % floor will result in heightened focus on reserves maintenance. Figure 5 presents an example percent-reserves vs. Cost to Go graph against a 10 % reserves floor.

Encumbrances against the reserves include hard liens, or those that were accepted by the teams and management, and soft liens, which were relatively more uncertain threats anticipated by team members. The total value of the encumbrances of the soft liens are reduced by a factor of probability of occurrence to yield an effective encumbrance against reserves.

A benefit to understanding with relative accuracy the fiscal reserve picture is the ability to periodically release reserves to reduce technical risk while still maintaining adequate reserves to allow for unknown unknowns. During the first year of Phase C/D, STARDUST purchased approximately $900K in risk-reduction items, including additional testing equipment, parts, and spare electronic boards. The additional equipment is intended to be a preemptive strike to avoid contention for test equipment during board and box testing, and against problems, which typically occur downstream in the program, usually during ATLO. The additional equipment would facilitate the addition of parallel operations, should they be necessary to recover the schedule.

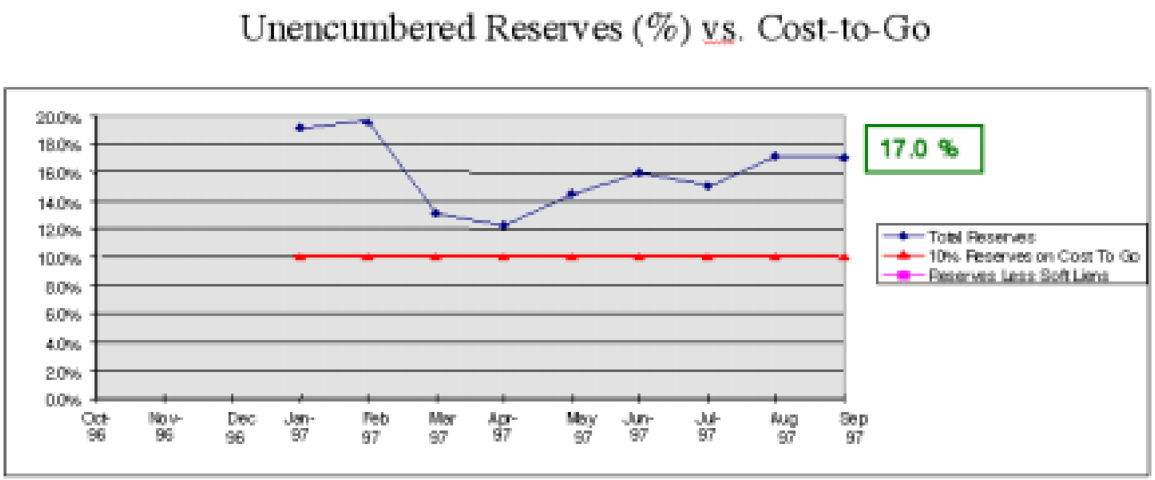

Figure 5. STARDUST Reserves on Cost To Go Metric

SUMMARY AND CONCLUSIONS

The STARDUST project has been a pioneer in implementing the *faster, better, cheaper* paradigm. It has successfully implemented an integrated information- and communication-system infrastructure that virtually co-locates project teams around the country and overseas. The team employed a suite of program-control tools and processes brings critical information about project progress and managing-to-budget to personnel in a clear informative manner. Travel time and expense was reduced over traditional implementations. Over one million dollars were allocated from reserves to risk-reduction items during the development as preemptive strikes against future anticipated problems.

In sum, STARDUST made the transition to a new paradigm of managing-to-budget. It works and we recommend others consider the merits.

ACKNOWLEDGEMENTS

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, in Pasadena California; Lockheed Martin Astronautics, in Denver, Colorado; and the University of Washington, in Seattle, Washington. The Jet Propulsion Laboratory is under contract with the National Aeronautics and Space Administration to manage the STARDUST project. Lockheed Martin Astronautics is under contract with the Jet Propulsion Laboratory, California Institute of Technology, to provide the STARDUST spacecraft. Dr. Don Brownlee of the University of Washington is the STARDUST project Principal Investigator.

***Bredt Martin** joined STARDUST as Manager of Business Operations and Information Systems in January 1996 following its selection as the Discovery 4 mission. He first worked at JPL in 1986 as a thermal environmental engineer in the Mission Assurance division, supporting various projects including the Microwave Limb Sounder for the Upper Atmosphere Research Satellite.*

In 1990, he became Planning Manager for the corrective-optics Wide-Field/Planetary Camera (WF/PC) II, which was installed into the Hubble Space Telescope during its First Servicing Mission in 1993. Under the program management and control techniques he developed, WF/PC II was delivered to the NASA customer ahead of schedule and within budget. Later, he was Program Control Manager for the SeaWinds scatterometer project. There, he co-developed the innovative, low cost earned-value system, which is now applied to the STARDUST project. He holds a Bachelor of Science degree in Mechanical Engineering from the University of Washington and a Master of Business Administration degree in Corporate Finance from the University of Southern California.

Ken Atkins *joined STARDUST as Project Manager in June 1995 during its Phase A competition for selection as the Discovery 4 mission. He first worked at JPL as a member of the technical staff in propulsion and mission analysis focusing on small body missions, then managed U.S. options on missions to Halley's comet in the late '70s. He next managed the JPL Power Systems Section, which, under his leadership, successfully delivered the Galileo power subsystem.*

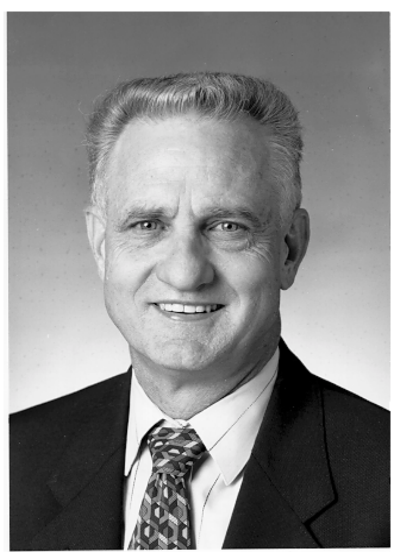

He later managed the Flight Command and Data Systems section focusing on deliveries of the Cassini command and data subsystem, development of flight software for the low-gain Galileo mission, and flight operations for Voyager, Galileo, and Mars Observer. He also managed the integrated avionics development for the Mars Pathfinder Discovery *mission. He has a doctorate in aeronautical and astronautical engineering from the University of Illinois.*

***Joe Vellinga** was LMA Program Manager for the STARDUST spacecraft and sample return capsule development. He was the proposal manager for the seven spacecraft proposals submitted by LMA in October 1994 for the first Discovery competition and was Program Manager for both the Suess-Urey and STARDUST competitive Phase A studies. He was program manager for the VIRSR seven color imager for the polar orbiting operational weather satellites Phase B study for NASA/GSFC. He also managed the program definition efforts for high resolution imaging spectrometers. He was program manager of the Manned Maneuvering Unit program bringing the flight units out of storage, testing them and preparing plans for re-certification and re-flight. He did the Faint Object Spectrograph system engineering, managed system test and managed the program for integration into the Hubble Space Telescope. He was payload manager for the SCATHA spacecraft, a research satellite carrying over 20 separate payload instruments. He supervised experiment integration compatibility activities for the Skylab corollary experiments.*

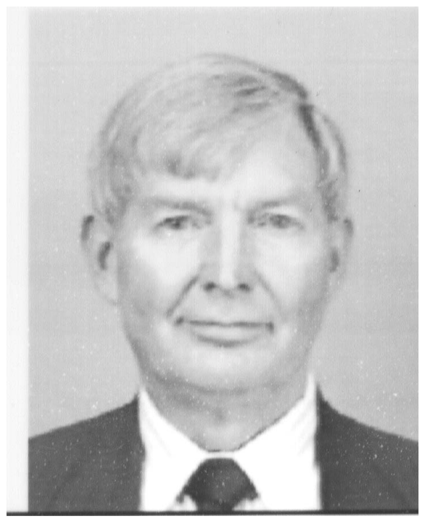

Earlier activities included propulsion research and development, bi-propellant reciprocating engine development for lunar power and aerodynamic analyses of supersonic inlet control systems. He has a Bachelor of Science degree from the University of California at Los Angeles.

***Rick Price** joined the STARDUST Flight System team as Chief of Program Planning and Control in April 1996. He has worked at LMA since 1980, beginning with the Titan 34D program. He has served as Chief of Program Planning and Control for the Commercial Titan, Titan II, and Titan IV launch vehicle programs. His experience includes extensive critical path schedule development and analysis, and Cost/Schedule Control System Criteria (C/SCSC) applications in both commercial and DOD environments. Rick was also instrumental in LMA's corporate development and implementation of the Integrated Product Development (IPD) management approach. He has a Bachelor of Science degree in Business Management from California State University Long Beach.*