Management and systems engineering of the Kepler mission

James Fanson [*] [a], Leslie Livesay [a], Margaret Frerking [a], Brian Cooke [a]
[a] Jet Propulsion Laboratory, California Institute of Technology,
4800 Oak Grove Drive, Pasadena, CA 91109

ABSTRACT

Kepler is the National Aeronautics and Space Administration's (NASA's) first mission capable of detecting Earth-size planets orbiting in the habitable zones around stars other than the sun. Selected for implementation in 2001 and launched in 2009, Kepler seeks to determine whether Earth-like planets are common or rare in the galaxy. The investigation requires a large, space-based photometer capable of simultaneously measuring the brightnesses of 100,000 stars at part-per-million level of precision. This paper traces the development of the mission from the perspective of project management and systems engineering and describes various methodologies and tools that were found to be effective. The experience of the Kepler development is used to illuminate lessons that can be applied to future missions.

Keywords: Kepler, spacecraft, exoplanet, transit, telescope, management, systems engineering

1. INTRODUCTION

The Kepler mission provides a rich history from which useful lessons can be drawn. Competitively selected in 2001 as the 10th mission in NASA's Discovery program, Kepler was successfully launched into an Earth-trailing solar orbit in March 2009. Kepler's formulation phase was unusual in that the composition of the project team was fundamentally altered following selection. Some of the choices made during this period resulted in unintended implementation difficulties downstream. This paper describes programmatic and technical challenges encountered during the development phase, and management and systems engineering approaches that were found to be particularly effective in addressing them. The paper is organized around five broad themes that the authors believe are important to successful mission implementation and that are illustrated by specific examples from the Kepler experience.

The paper's perspective is that of the core project management team at the Jet Propulsion Laboratory (JPL), most of whom joined the project in 2006 during a period when things were not going well and firm management action was deemed necessary. The authors wish to express their appreciation to all of the many individuals who contributed to the Kepler project, especially those at the NASA Ames Research Center (ARC) and Ball Aerospace & Technologies Corp. (Ball).

2. A BRIEF PROGRAMMATIC HISTORY OF THE KEPLER MISSION

NASA's interest in detecting extrasolar planets spans decades. Efforts to produce technology and mission roadmaps to achieve this goal gained momentum in the 1980s [1,2]. Further impetus resulted from the first detection of an extrasolar planet orbiting a normal star (51 Pegasi) in 1995 [3]. One method for detecting extrasolar planets is to measure the dimming of stars when planets transit across their disks. William Borucki at ARC championed the idea of a space-based photometer to measure a large number of stars over several years to search for Earth-size planets orbiting in the habitable zone; that is, the region around a star where the temperature of a planet could support liquid water on its surface. Borucki proposed such a mission at several opportunities beginning in 1992, but years of effort were required to demonstrate the feasibility of making measurements at the required level of precision [4]. Detecting an Earth-like planet transiting a Sun-like star requires photometric precision at the level of one part in one hundred thousand, considered beyond the state of the art until the turn of the millennium.

[*] James.L.Fanson@jpl.nasa.gov; phone 1 818 354-7014; fax 1 818 354-5461; jpl.nasa.gov

The fourth and winning proposal from ARC envisioned a streamlined implementation, very much in the spirit of the *faster-better-cheaper* paradigm in vogue at NASA during the 1990s. Ball would design and build the spacecraft and photometer (a 1.4-meter class telescope with a 95-megapixel focal plane array) and operate the mission with the support of the Laboratory for Atmospheric and Space Physics (LASP). ARC would manage the development, build the science data analysis pipeline, and lead science operations. By the time the mission was selected in December 2001, several in-flight failures had shifted NASA's implementation paradigm from *faster-better-cheaper* to *mission success first*, with an increased emphasis on rigor and oversight by experienced mission developers. Three months after Kepler's selection NASA established a new policy requiring all robotic missions to be managed by JPL or the Goddard Space Flight Center (GSFC), NASA's two field centers having the most experience with such work. The Kepler project was subsequently directed to add to the team either JPL or GSFC to provide project management, systems engineering, and mission assurance leadership during development. JPL was chosen, and Kepler mission formulation commenced in 2002.

By 2006, due to a combination of factors, it became evident that Kepler could likely not be completed within the agreed-upon budget. A series of replans were undertaken during the spring of 2006, resulting in a revised forward plan and budget and a new launch date of February 2009. A new management team was appointed to implement this plan; however, within a year it became clear that continued difficulties—some technical, others organizational—again threatened successful completion within the agreed budget. By this time, a new Associate Administrator (AA) had assumed the leadership of NASA's Science Mission Directorate. The new AA was determined to stem mission cost growth. Therefore, in the summer of 2007, the Kepler project, NASA, JPL, ARC, and Ball collectively committed to a set of agreements to enable the mission development to be completed at no increase in cost.

These agreements involved further re-organization of the project, which was converted from a Principal Investigator (PI)-led mission to a directed mission led by the Project Manager. Ball reorganized, streamlining a number of internal processes. Some planned tests were descoped and the associated risks accepted. Under these new rules of engagement, the project successfully completed development, launch, and commissioning, and began science operations within the budget agreement established in 2006.

As of this writing, Kepler has discovered its first exoplanets [5] and is well on its way to answering one of the most profound questions of the ages: are there other Earth-like planets in the Universe?

3. THEME 1: ORGANIZE BOTH THE TEAM AND THE PRODUCT ALONG SIMPLE LINES WITH CLEAR RESPONSIBILITY

Adding JPL to the team posed unique challenges. The existing team of ARC and Ball had developed the mission concept together over several years and had established a rapport and mutual understanding; they were understandably unenthusiastic about the change in management. In response, JPL elected to implement an atypical, more "hands-off" strategy. The social context of this transition, coming as it did after selection, had far-reaching implications for the project's development.

The resulting project organization complicated the governance relationships between JPL, ARC, and Ball in important ways. For example, while day-to-day project management and spacecraft development was the responsibility of JPL, contractual responsibility for the photometer—clearly the most technically challenging portion of the flight vehicle—remained with ARC. Direction on the photometer contract flowed from the ARC Contracting Officer's Technical Representative (COTR) through the Contracting Officer to Ball. It was not clear to Ball whether JPL had authority to direct work that was funded under the ARC contract. In fact, authority was divided, which resulted in confusion and a lack of crisp, timely decision making.

The situation was again complicated when a cost-saving decision was taken early on to consolidate to a single flight computer for both the spacecraft bus and photometer instrument. The flight computer and flight software were added to the photometer contract at ARC, in part because it would be needed early for testing with the photometer and, in part, because the different Center overhead structures made subcontracts less expensive if executed through ARC. This resulted in a situation where scope changes made to fault protection, for example, required changes to be negotiated with two Contracting Officers at two NASA centers for two contracts having different terms and conditions and different fee structures.

Beginning in 2006, new management arrangements were implemented to clarify lines of authority and strengthen decision making. ARC and JPL agreed that while CoTR responsibility for the ARC contract with Ball would remain at ARC, JPL would provide overall technical leadership. At JPL, two key team members were added: a Deputy Project Manager with responsibility for the entire flight system and an Instrument Manager with responsibility for the photometer. This decision allowed work to proceed in both contracts with clearer lines of responsibility and, thus, quicker decision making. It was not an optimal solution and continued to cause problems in contract management and change control; however, a significant streamlining of the team management was achieved. The decision to change the management structure arose from a multi-day retreat focused on improving management of the project. The impact of this retreat was significant: everyone involved had a greater understanding of the issues from different perspectives and everyone agreed on a common forward path.

Systems engineering was less affected by the complex governance structure between the various players. A strong and continuous systems engineering team was in place throughout the development and was instrumental in anchoring the team through the various replans and descopes. Project-level systems engineering management and issue resolution was led by the Project Systems Engineer (PSE) at JPL. This included top-level requirement verification, interaction with the science team, and execution of the top-level system validation campaign. Flight system-level systems engineering management and anomaly resolution was coordinated by the Flight System Systems Engineer (FSSE) at Ball. Areas of responsibility included flight system-level requirements verification, board- and box-level anomaly investigation and resolution (with project approval), fault protection testing, mission scenario test development and execution, and flight system resource management. Ground system systems engineering and technical management was led by the Ground System Engineer (GSE) at ARC, who was responsible for ground system requirements verification, ground system anomaly resolution, ground system thread testing, and operations readiness tests.

Relatively late in the project implementation cycle, the PSE shifted to the role of Chief Engineer (CE) and the Deputy PSE moved up to PSE. This permitted the CE to focus on investigating and resolving specific technical issues that arose and, thus, protect the project schedule. In the final phase of development, the PSE's responsibility was broad and the CE's was deep. Interactions between the PSE, FSSE, GSE, and CE were dominated by discussions of interface agreements, cross-boundary deliverables, and ways in which modifications in one area could relieve difficulties in another. This division of labor allowed clear lines of responsibility and clear areas of interaction and co-dependence. Moreover, the addition of the CE created a path for off-loading the investigation of major issues and allowing the system leads to continue making progress in unaffected areas and focus on system validation.

4. THEME 2: MAINTAIN A STABLE IMPLEMENTATION ENVIRONMENT

It is well understood in the world of acquisition that changes to contract scope cost money. What is equally true, but less well appreciated, is that changes to any aspect of the implementation environment—requirements, processes, rigor of policy enforcement, personnel, funding profile, schedule—all have cost (and risk) ramifications. Change is also fertile ground for unintended consequences.

Kepler experienced several perturbations to its development timeline that contributed to cost and schedule growth. Shortly after Preliminary Design Review (PDR), mid-way through government fiscal year 2005, NASA found itself with limited available funding and sharply reduced the funds to Kepler for the remainder of the fiscal year. This had the effect of slipping the launch date and triggering another replan. Had this funding profile been part of the project's original plan, it would have been possible to effectively use the added time to buy down risk before ramping up the team in preparation for the PDR. The timing and magnitude of the reduction resulted in a steep roll-off of contractor personnel and a subsequent loss of corporate- and Kepler-specific knowledge. In addition, the schedule change put Kepler in contention for resources with a defense program at Ball, which was an ongoing cause of difficulty through delivery. The impact of this funding profile change was clearly underestimated by the project; also, in important respects, the momentum lost during this period was not recovered until much later in the development. A stable, appropriately funded Phase B is critical to successful project development. Changes in Phase C funding might have implications beyond a specific project (i.e., within an implementing company) that have profound impacts on the project's ability to deliver.

5. THEME 3: DON'T UNDERESTIMATE COMPLEXITY OR THE DIFFICULTY OF SCALING EXISTING TECHNOLOGY

It was a point of particular note that the Kepler flight vehicle involved no new technologies (Figure 1). While this was technically true, it was also the case that the photometer's focal plane array assembly (FPAA) featured a focal plane of unprecedented scale: 95 megapixels, or nearly one square foot of silicon. The FPAA comprises more than 20,000 components and represents the most complex cubic foot of space flight hardware ever developed by Ball (Figure 2). Located at the prime focus of the telescope, the FPAA supports 42 science and four fine guidance sensor charge coupled devices (CCDs) mounted to a convex surface, together with ten readout electronics boards necessary to achieve part-per-million photometric precision.

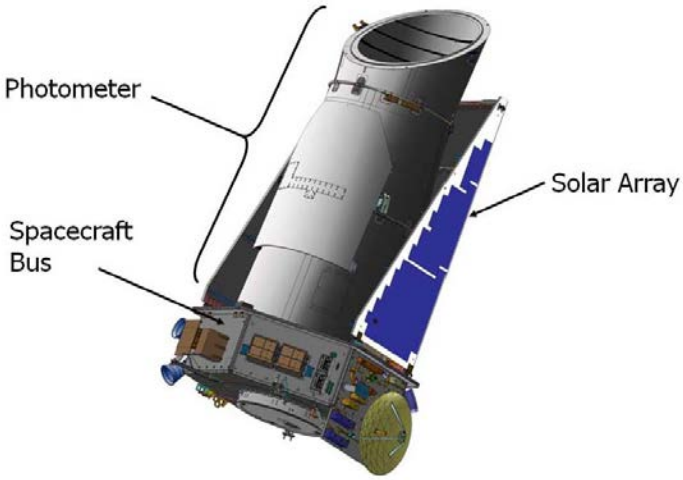

Figure 1. *Kepler Flight Vehicle*. Kepler is a space telescope having a wide field of view photometer capable of detecting Earth-size planets orbiting other stars in the galaxy.

Figure 2. *Focal Plane Array Assembly*. The Focal Plane Array Assembly (FPAA) contains 20,000 components, including 42 science CCDs that form a 95-megapixel focal plane.

While simple in concept, the design and packaging of the FPAA was anything but simple in practice. A sophisticated thermal isolation system was needed to maintain the CCDs at -95º C while in close proximity to readout electronics generating 110 W of heat. Heat pipes, as well as power and signal cable harnesses, were needed to span between the telescope prime focus and the exterior of the vehicle. They had to be nestled along the legs of a spider support structure, which also served to position the focal plane array at the optical focus of the telescope. Achieving low noise performance while simultaneously achieving high readout speed and minimum power dissipation was a daunting challenge.

Several decisions resulted in unanticipated difficulty downstream. For example, to reduce readout electronics power dissipation, a decision was made to power off signal chain amplifiers when not in use during CCD readout. This resulted in degraded performance from the amplifiers because they were never in thermal equilibrium. It also resulted in low-level oscillations from some of the amplifiers, which introduced systematic pattern noise into roughly one-third of the focal plane.

One of the many cost-cutting measures involved a decision not to build a full engineering model readout electronics unit but, instead, to perform development tests on only a pair of boards. This resulted in late discovery of cross talk between the fine-guidance sensor readouts and the science CCDs, which introduced additional systematic noise into the focal plane. By the time these various noise affects were fully appreciated it was too late to correct them in the hardware. Costly and time consuming ground software had to be developed, much of it post-launch, to mitigate the noise via sophisticated calibration techniques.

The degree of engineering effort required to design and test the FPAA was significantly and chronically underestimated. While the technology was not new, the scale and complexity, together with the demanding performance requirements, represented a technical and management challenge of the first order.

Several important systems engineering measures were brought to bear on these challenges. Ball brought to the team very strong requirements verification discipline. More than 23,000 individual requirements were verified across the flight segment. JPL complemented this with strong validation discipline, including stress testing and in-depth data review with participation from the science team. The approach was extended to the flight system-level through carefully constructed mission scenario tests, fault protection tests, end-to-end information system tests, and operations readiness tests. The PSE, CE, and FSSE formed a closely coordinated team to ensure a solid intellectual basis for the verification and validation campaigns. This was particularly important in view of the fact that end-to-end performance testing of part-per-million photometric precision was not practicable; a combination of tests and analyses were required to establish confidence in system-level performance.

6. THEME 4. HAVE EFFECTIVE METRICS AND MANAGEMENT CONTROLS

Following the replan in 2006, the team was again working at full steam. It became necessary to develop a set of metrics and reporting tools to assist in assessing performance to plan. Ball had developed a detailed integrated master schedule (IMS) of approximately 16,000 lines, rigorously updated every week. Given the large number of parallel activities underway, it was difficult to assess schedule performance. Traditional earned value metrics are lagging indicators: cost data is typically available only monthly; trend information, when it emerges, is often several weeks old, too late to enable timely proactive measures to be taken.

The management team, therefore, undertook to develop metrics that could provide an approximation to earned value, but on a weekly cadence. A set of several hundred tracking milestones were selected from the IMS, each being significant to the delivery of a subsystem, test, or integration activity. Tools were built to enable the IMS to be processed each week and to produce graphs of the cumulative number of tracking milestones accomplished per week compared to plan for each of the many parallel activities (Figure 3). In addition, tables were automatically produced that showed the status of every tracking milestone planned for completion during each week (Figure 4). The tables included the baseline completion dates, projected completion dates from the current IMS, schedule slack, and percent complete, all color coded to show at a glance which milestones were late, which were on time, and which were early. Together with weekly workforce charts for each area, these data provided a simplified earned value assessment each week. We found this to be very effective in quickly detecting variances from plan since trend information could be established on two-week intervals. We also developed a tool to track and trend schedule margin on the critical path and to compare the actual depletion of schedule margin with the maximum desired depletion rate. As the critical path through the network logic

changed from week to week, we could still determine the overall trend in margin depletion and use this information to modify the plan and restore margin or manage the release of funded schedule reserve (Figure 5).

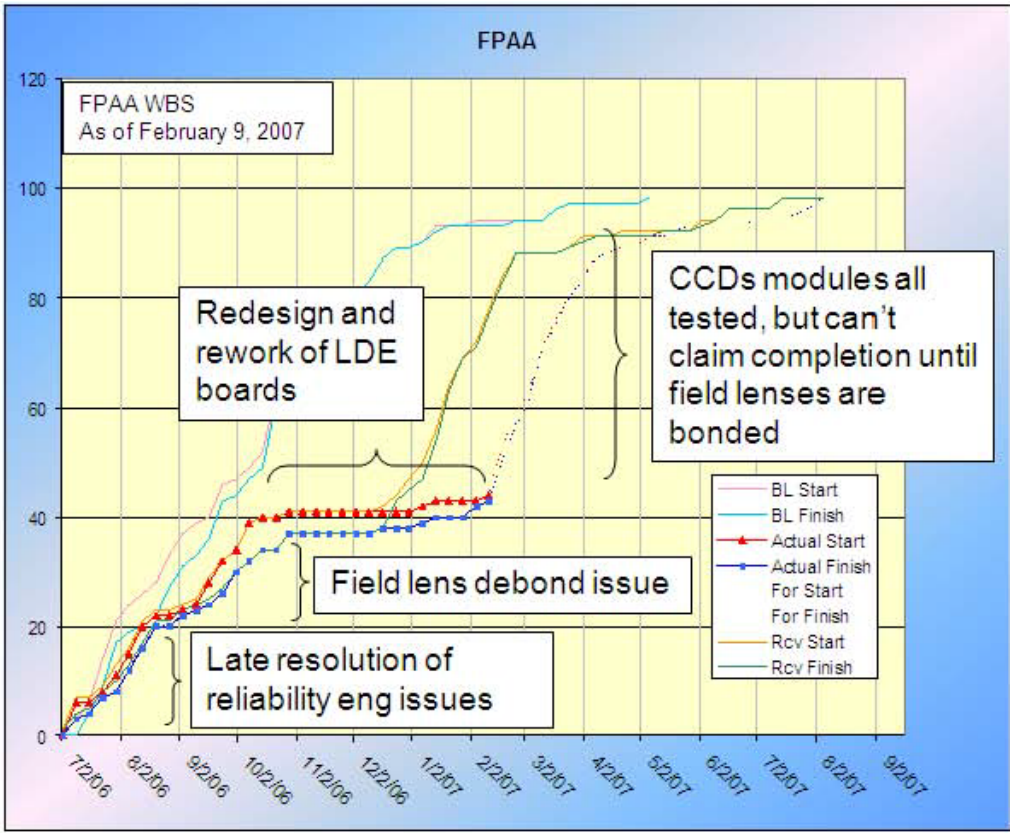

Figure 3. *Weekly Milestone Tracking Metric Example*. This graph illustrates the milestone tracking metric for the FPAA as of February 2007. The curves plot cumulative milestone counts as a function of time. The leftmost pair of curves (labeled "BL" in the legend) represents the baseline plan activity start and finish counts. Actual accomplished starts and finishes are plotted in red and blue data points, and show the affect on schedule of several technical problems being worked. The tool also had the ability to plot recovery plan curves (one such pair is labeled "Rcv" in the legend). The dotted lines (labeled "For" in the legend) are extrapolations of milestone counts based on the network logic contained in the IMS. The tool permitted close tracking of schedule variances on a weekly cadence. This example shows the worst schedule variance encountered after the 2006 replan.

										2/18/07	2/18/07	2/18/07	2/18/07	2/18/07	2/18/07
IPT	ubsyste	Resp	Float	UID	stone Gr	Name	Dur	BL Start	BL Finish	Start	Finish	Start Var	Finish Var	PC	Float
FPAA	FPA Mo		0	2626	ptm	Module 2 assy complete	0	8/25/2006	8/25/2006	2/22/2007	2/22/2007	124	124	0%	33
FPAA	FPA Mo		0	23996	ptm	Module 3 assy complete	0	8/25/2006	8/25/2006	2/26/2007	2/26/2007	126	126	0%	31
FPAA	FPA Mo		0	2642	ptm	Module 4 assy complete	0	8/25/2006	8/25/2006	2/27/2007	2/27/2007	127	127	0%	34
FPAA	FPAA LI		2	14936	ptm	Flight Bd Pair Test #1 - PH0060	5	9/14/2006	9/20/2006	2/20/2007	2/23/2007	109	108	0%	-11
FPAA	FPAA LI		2	14937	ptm	Flight Bd Pair Test #2 - PH0060	5	9/21/2006	9/27/2006	2/28/2007	3/5/2007	110	109	0%	-13
FPAA	FPA Mo		0	2651	ptm	Module 5 assy complete	0	9/21/2006	9/21/2006	3/5/2007	3/5/2007	113	113	0%	30
FPAA	FPA Mo		47	2659	ptm	Module 6 assy complete	0	9/21/2006	9/21/2006	3/7/2007	3/7/2007	115	115	0%	28
FPAA	FPA Mo		44	2667	ptm	Module 7 assy complete	0	9/21/2006	9/21/2006	3/9/2007	3/9/2007	117	117	0%	26
FPAA	FPA Mo		37	2675	ptm	Module 8 assy complete	0	9/21/2006	9/21/2006	3/13/2007	3/13/2007	119	119	0%	24
FPAA	FPAA LI		2	14938	ptm	Flight Bd Pair Test #3 - PH0060	5	9/28/2006	10/4/2006	3/6/2007	3/9/2007	109	108	0%	-13
FPAA	FPAA LI		2	14939	ptm	Flight Bd Pair Test #4 - PH0060	5	10/5/2006	10/11/2006	3/12/2007	3/15/2007	108	107	0%	-13
FPAA	FPAA I&		4	21210	ptm	Certify focal plane test set	9	10/9/2006	10/19/2006	1/2/2007	2/20/2007	57	84	94%	24
FPAA	FPAA LI		2	14940	ptm	Flight Bd Pair Test #5 - PH0060	5	10/12/2006	10/18/2006	3/16/2007	3/21/2007	107	106	0%	-13
FPAA	FPAA I&		20	16901	ptm	Proof load DAA vibe plate w/ load cells	10	10/12/2006	10/25/2006	2/15/2007	2/15/2007	86	77	100%	0
FPAA	FPA Mo		37	2701	ptm	Module 11 assy complete	0	10/19/2006	10/19/2006	3/1/2007	3/1/2007	91	91	0%	36
FPAA	FPA Mo		30	2709	ptm	Module 12 assy complete	0	10/19/2006	10/19/2006	3/20/2007	3/20/2007	104	104	0%	23

Figure 4. *Tracking Milestone Variance Table*. Tables showing current milestone status compared to the baseline plan (or recovery plan) facilitated fortnightly variance discussions with subsystem managers.

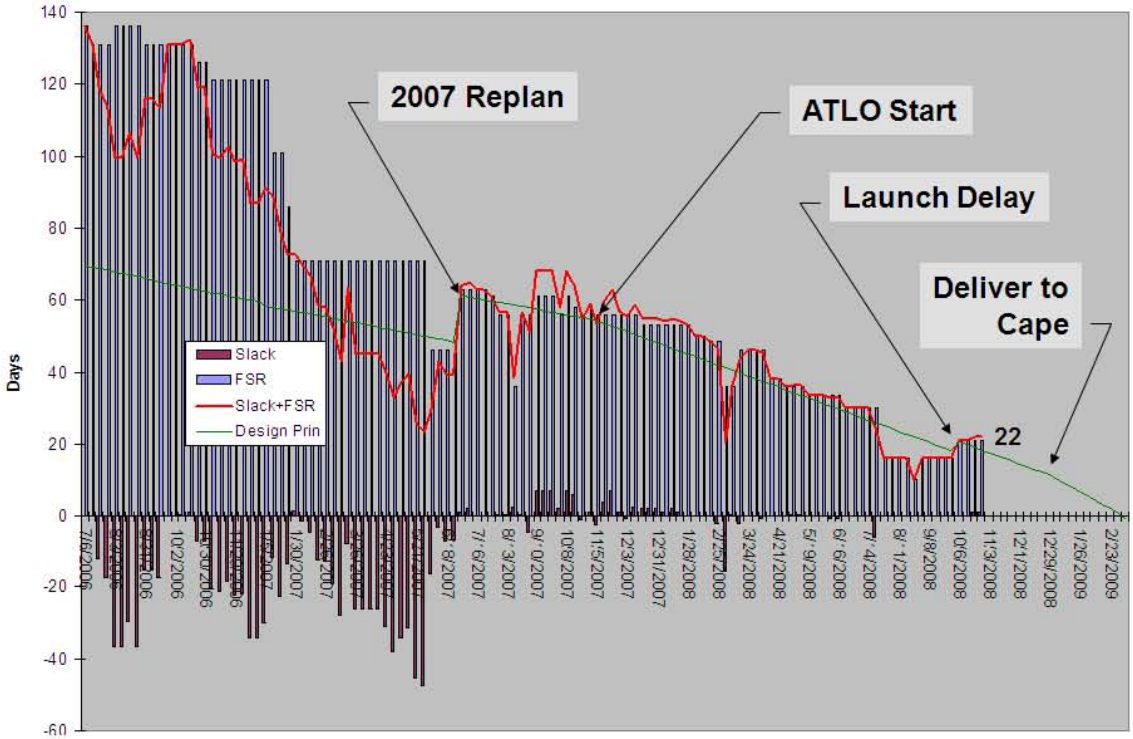

Figure 5. *Critical Path Schedule Reserve Depletion Trend*. Critical path schedule slack and funded schedule reserve (FSR) were tracked each week and compared to the JPL Design Principles depletion rate guideline (green line). Reserve management following the 2007 replan was excellent.

The weekly metrics enabled much more focused discussions with managers on exactly what was causing variances from plan and what mitigation actions might be taken. We also elected to move from large Monthly Management Meetings to smaller, more interactive meetings on a fortnightly cadence. Thus, the interval required to identify a problem, understand its cause, and work solutions with the appropriate individuals was considerably shortened.

To keep the entire team aware of the bigger picture, a decision log was maintained for every decision taken by the management team. Every major problem, concern, or issue being worked by the management team (actual problems as distinct from risk items in the risk database) was tracked on a spreadsheet and ranked in order of criticality (Figure 6). The problems and concerns were scored to arrive at a composite "technical stress index" that was used to communicate the overall technical burden being carried by the team (Figure 7). From 2006 through in-orbit checkout in 2009, a total of 202 individual items were tracked.

Risk Management was established early in the project development cycle, but by early 2007 had lost significant focus and discipline. The addition of several new team members provided impetus to revitalize the risk management approach and enforce better risk management discipline on project leadership. The first part of the revitalization effort re-evaluated several hundred existing project risks for relevance to the current phase of implementation. Almost one hundred risks were retired as a result of this re-evaluation, mainly due to the fact that the milestone for risk realization had passed. Other risks were retired or rejected because they took the form of "if team A does not do their job on time then team B will be impacted." While these risk statements were acknowledged as true, the revitalized risk process treated receivable/deliverable commitment tracking and schedule optimization as a project management task, not as a risk. Finally, the remaining risks were clearly described in the form of an "If.. Then... Resulting in..." statement, ranked for likelihood and consequence, and given clear assignees. Assignees were then tasked with developing a mitigation plan and identifying several evaluation milestones. On a weekly basis, the project risk management team reviewed new risks, reviewed and approved mitigation plans, and monitored the progress of active mitigation efforts.

Decisions involving the acceptance of significant risk were arrived at through a formal process involving the full management team. Following a briefing and discussion of available options, the decision was taken and individual approvals recorded on a Risk Decision Form, together with a statement of the risk and any dissenting opinions.

Kepler Flight Segment			2/9/2007	2006						2007			
			Technical Stress Index →	12	16	20	28	26	22	22	18		
	PROBLEM/CONCERN/ISSUE LIST			Rebaseline			SYS CDR			GND CDR		Trend	Project Office Point of Contact
	Item	Lead Org	First Identified	J	A	S	O	N	D	J	F		
1	ATK Magna MDRs Delaying Metering Structure	BATC	Oct-06				C	P	P	P	P		Peg Frerking
2	ATK Corona Bonding Anomalies (now at Magna)	BATC	Jul-06	P	P	P	P	C	C	C	C		Jim Fanson
3	Field Flattener Lens Debond	BATC	Sep-06			C	P	P	P	P	C		Jim Fanson
4	Compl. of Reliability Anal. Prior to Build Starts	BATC	May-06	C	C	C	C	C	C	C	C		Jim Newell
5	ATK Damage to Deployable Cover Panel	BATC	Nov-06					C	C	C	C		Jim Fanson
6	Detection of Transits on SSTVT	BATC	Nov-06					I	I	C	C		Riley Duren
7	"Undershoot" Corruption of PSF (LDE)	BATC	Oct-06				P	P	P	C	C		Peg Frerking
8	"Ringing, Ramp & Slope" Image Artifacts (LDE)	BATC	Oct-06				C	C	C	C	C		Peg Frerking
9	Digital Card Failure (SRAM) on WISE Unit	BATC	Oct-06				I	C	C	C	C		Jim Newell
10	LN200S IMU WDM Failure	BATC	Jun-06	I	C	C	C	C	C	C	C		Jim Fanson
11	Ka-Band TWTAs	BATC	Sep-06			C	C	C	C	C	C		Leslie Livesay
12	SDST K-band Failure on MRO	JPL	May-06	P	P	P	P	C	C	C	C		Leslie Livesay

Figure 6. *Problem/Concern Chart*. Problems (P), concerns (C), and lower level issues (I) were identified and ranked, with trend information provided for each. A "technical stress index" was derived by summing the number of problems (each given a score of three) and the number of concerns (each given a score of one).

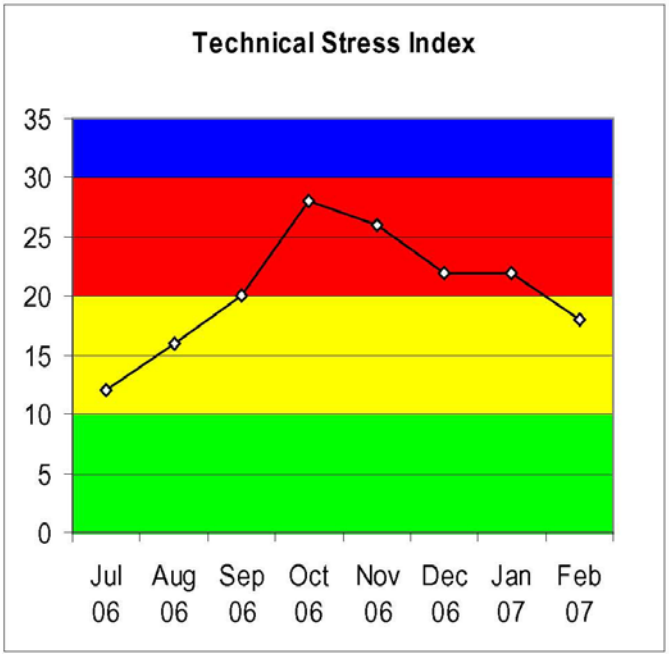

Figure 7. *Technical Stress Index*. The technical stress index provided a rough measure of the overall technical burden being carried by the team. It was calculated as a weighted sum of the number of problems and concerns being worked at any given time (see Figure 6). The trend curve was plotted on a color coded scale to indicate low (green), moderate (yellow), high (red), or extreme (blue) risk to meeting project-level commitments. The excursion into the red zone resulted in the need for a replan in 2007.

7. THEME 5: HAVE AN EFFECTIVE MEANS TO IMPLEMENT REPLANS, AND THINK OUT OF THE BOX

Replaning is an inevitable feature of project development and often requires rapid trading of risk, cost, and performance in support of decision making. The ability to make trades quickly can be very important to maintaining programmatic credibility in times of crisis.

Kepler's scientific capability depends on a combination of many technical factors, including the number of stars measured, the completeness of the data for each star, the contiguity of that data, the sensitivity of the measurements, and the duration of the search. These factors combine in complex and subtle ways. To help the team maximize scientific capability during replans, an analysis tool was developed to quantify the ability of Kepler to find planets of various types given a specific set of engineering parameters. This "Science Merit Function" tool was central to the project's ability to make design trades, manage risks, and, ultimately, to decide what capabilities could be sacrificed when cost growth required descopes (Figure 8). It was also used to develop confidence in the scientific integrity of the mission at key programmatic decision points.

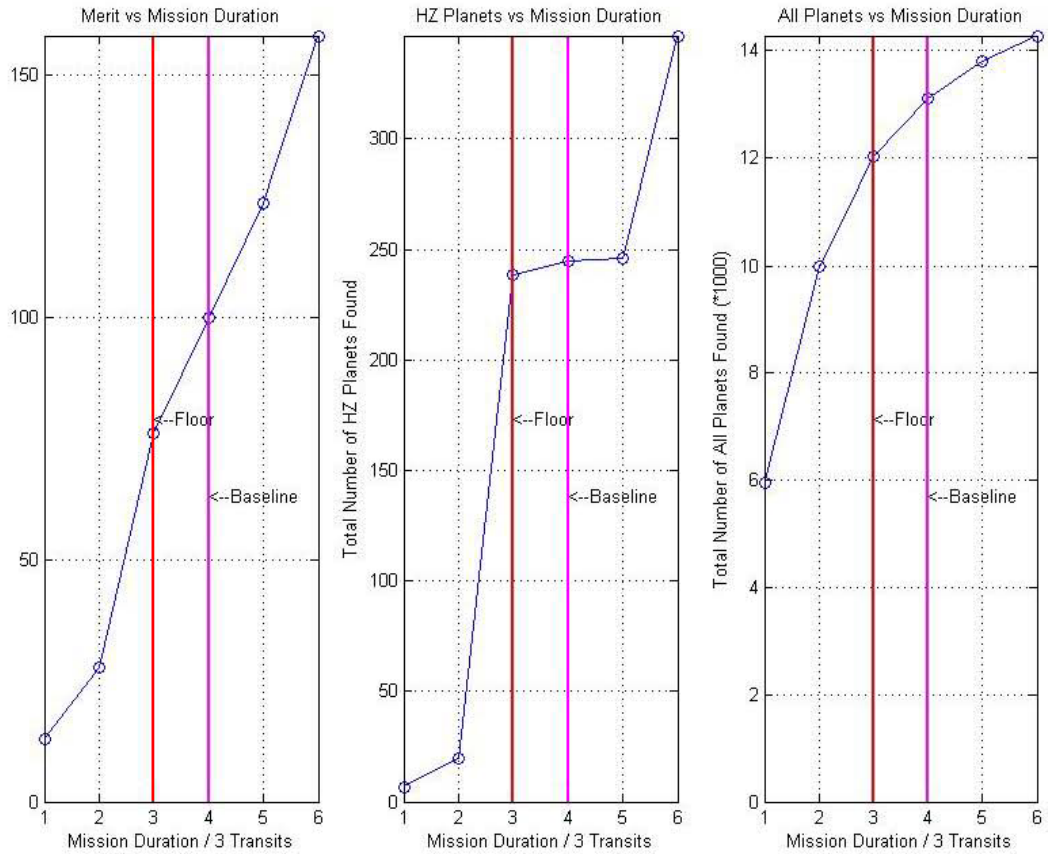

Figure 8. *Science Merit Function Tool*. A Science Merit Function tool was developed to assist with trades and descopes during development and to establish confidence in the science capability at key gate reviews. Merit was defined as the planet detection rate, scaled to 100 points for the baseline mission. This example shows the relative number of planets of various types that can be discovered as a function of mission duration in years. Notice the strongly nonlinear dependence for habitable zone (HZ) planets.

Kepler executed several replans during its development. The 2006 replan involved the rapid assessment of about 100 descope options, of which 18 major and about 50 minor were implemented. The most eventful replan occurred in the summer of 2007, when the project was required to develop a forward plan $42 M less expensive than what was believed necessary to achieve launch in early 2009. Approximately 50 descopes were evaluated, with 25 implemented. During

both of these descope phases, the project elected to have a Red Team and the NASA Program Office participate in the process. This facilitated the independent risk assessment and timely approval of the resulting plan. Descopes alone were not sufficient, however. Faced with such a daunting challenge, the project team, supported by the various implementing institutions, pulled together a set of far-reaching agreements requiring each stakeholder to put additional "skin in the game," but which made possible a credible plan. There is nothing like the threat of cancellation to steel the will of those committed to mission success.

8. CONCLUSIONS

In spite of the many challenges faced by the Kepler project, the team and its supporting institutions, including NASA Headquarters and the Program Office at the Marshall Space Flight Center, collectively overcame every obstacle, and delivered a mission that promises to revolutionize our view of our place in the Universe (Figure 9). Many difficulties could have been avoided by different choices early in formulation, when the team was being assembled and the project organized. Some challenges were programmatically driven; some were beyond the control of the team and NASA. The Kepler experience illuminates several important aspects of sound implementation practice: 1) organize the team and the product along simple lines with clear responsibility; 2) maintain a stable implementation environment; 3) don't underestimate complexity or the difficulty of scaling existing technologies; 4) have effective metrics and management controls; and 5) have an effective means to implement replans and think out of the box.

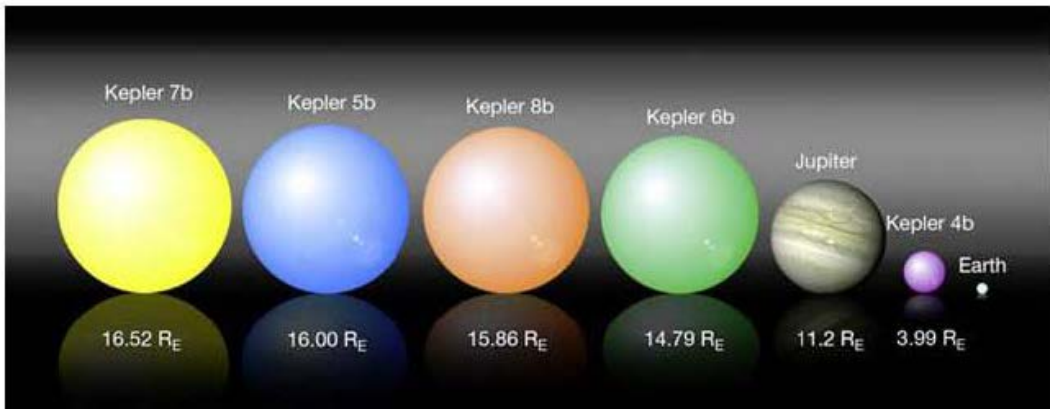

Figure 9. *Kepler's First Exoplanet Discoveries*. Kepler discovered five new planets in its first 10 days of operation. They are shown here to scale alongside Jupiter and Earth. Discovery of Earth-size planets orbiting in the habitable zones of stars will require several years of observations.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the helpful suggestions of Riley Duren and Gentry Lee in the preparation of this manuscript. The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1] Burke, B. et al., Tops: Toward Other Planetary Systems, NASA (1990).

[2] Beichman, C. et al., A Road Map for the Exploration of Neighboring Planetary Systems (ExNPS), JPL Publication 96-22 (1996).

[3] Mayor, M. and Queloz, D., *Nature* 378, DOI:10.1038/378355A0, Nature Publishing Group, 355 - 359 (1995).

[4] Borucki, W. et al., Kepler: Search for Earth-Size Planets in the Habitable Zone, Proceedings IAU Symposium No. IAUS253, International Astronomical Union, 289-299 (2008).

[5] Borucki, W. et al., *Science* 19: Vol. 327. No. 5968, DOI: 10.1126/Science.1185402, American Association for the Advancement of Science, 977 – 980 (February 2010).