

Deployment Process, Mechanization, and Testing for the Mars Exploration Rovers

Ted Iskenderian*

Abstract

NASA's Mars Exploration Rover (MER) robotic prospectors were produced in an environment of unusually challenging schedule, volume, and mass restrictions. The technical challenges pushed the system's design towards extensive integration of function, which resulted in complex system engineering issues. One example of the system's integrated complexity can be found in the deployment process for the rover. Part of this process, rover "standup", is outlined in this paper. Particular attention is given to the Rover Lift Mechanism's (RLM) role and its design. Analysis methods are presented and compared to test results. It is shown that because prudent design principles were followed, a robust mechanism was created that minimized the duration of integration and test, and enabled recovery without perturbing related systems when reasonably foreseeable problems did occur. Examples of avoidable, unnecessary difficulty are also presented.

Introduction

The highly integrated, collaboratively engineered mechanisms of the 2003 Mars Exploration Rover represent a significant achievement for the team that produced this spacecraft in record time. Two of these spacecraft were created as successors to the acclaimed Mars Pathfinder "Sojourner" rover that landed on Mars on July 4, 1997. The identical rovers landed on Mars on January 3rd and 24th, 2004.

Relative to most spacecraft, each Rover and the landing system that accompanies it feature a uniquely high bulk density. The roughly rectangular prism of the rover fits snugly within the stowed tetrahedron configuration of the lander. A complex deployment process results from this dense packaging. A total of 24 devices and mission-critical deployments must correctly function just to prepare the MER for driving off its base petal landing platform. The design of several of these mechanical systems depended on successfully implementing the concept for standup to which the rover was committed at an early date. Figure 1 shows an overall view of the MER before standup.

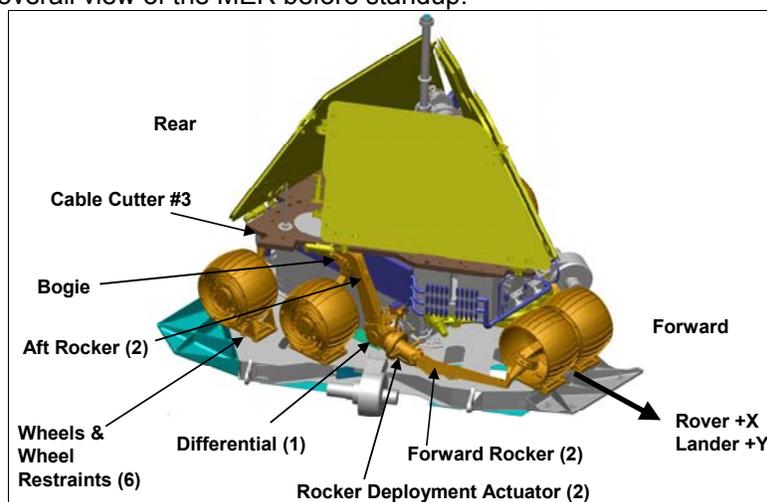


Figure 1. The stowed MER in its pre-standup configuration

* Jet Propulsion Laboratory, Pasadena, CA

The Martian environment also poses a severe challenge to mechanism designers. The design survival temperatures (which include margin) for the RLM's environment spanned from -120 to +35°C, and the qualification operating temperatures were between -60 to +35°C. The rover deployment system discussed herein must be able to survive these extremes over the 20 “sols” (Martian days) over which standup may occur immediately after landing. Some motors on the rover must survive many cycles from -105 to +8 °C over the mission's design lifetime, 90 sols. Additionally, beyond the anticipated high temperature during the mission, was the added necessity to reduce biologically active contamination to levels beneath those set by Planetary Protection engineers. Most mechanisms engineers on the MER met this requirement by baking their devices to +110°C for 50 hours prior to final integration with the spacecraft. Mechanisms had to survive not only routine launch vibration, but also the impact load up to approximately 49 g that would be experienced upon landing with the distinctive airbag cushioning system developed at JPL. As will be discussed, gales of dust particles were anticipated upon landing and during intermittent dust storms after landing.

Figure 2 shows the essentials of the standup sequence comprising the last of the mission-critical deployments and approximately when they occur. Figure 4 shows the rest of the sequence with a series of snapshots from the computer model. The functional requirements on the RLM throughout this process were as follows:

1. Lift the rover and suspension system far enough to allow the front wheels to clear the base petal (the plate to which the Rover was stowed) when they fully deployed. We designed to an arbitrary value, 75 mm, which proved to be adequate in all our tests. The time allocated to this function was 15 minutes.
2. Cause the rocker-bridge latch pivot, the juncture between forward and aft suspension, to rotate through a minimum angle of 42 degrees. (This angle was estimated to be about 39.5 degrees at the beginning of the project. It was identified as one of the poorly understood requirements, as it depended on the development of the rocker-bridge latch). The corresponding linear stroke requirement grew to about 244 mm (including 21 mm margin).

At this point, the front rockers and wheels rotate through their own deployment range of approximately 180°. This is accomplished through a series of motions shared between the wheel-steering actuators and the rocker deployment actuators. This sequence is detailed in the paper entitled “*The challenges of designing the rocker-bogie suspension for the Mars Exploration Rover*”.

3. Upon reversing, allow the loose pivot joint in requirement (2) to latch up at a fixed angle. Continue lowering the rover without interruption, while accommodating the fundamentally different kinematic relationships of the system.

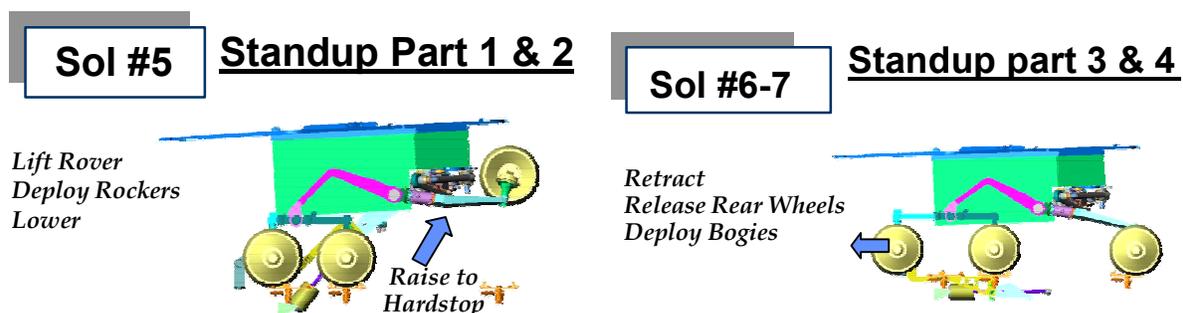


Figure 2. Sequence of Rover standup events

RLM Functional Requirement & Standup Sequence

The nomenclature of the RLM and its relationship to the rover is shown in Figure 3. This simplified computer model shows the rover's configuration just before the standup sequence is initiated.

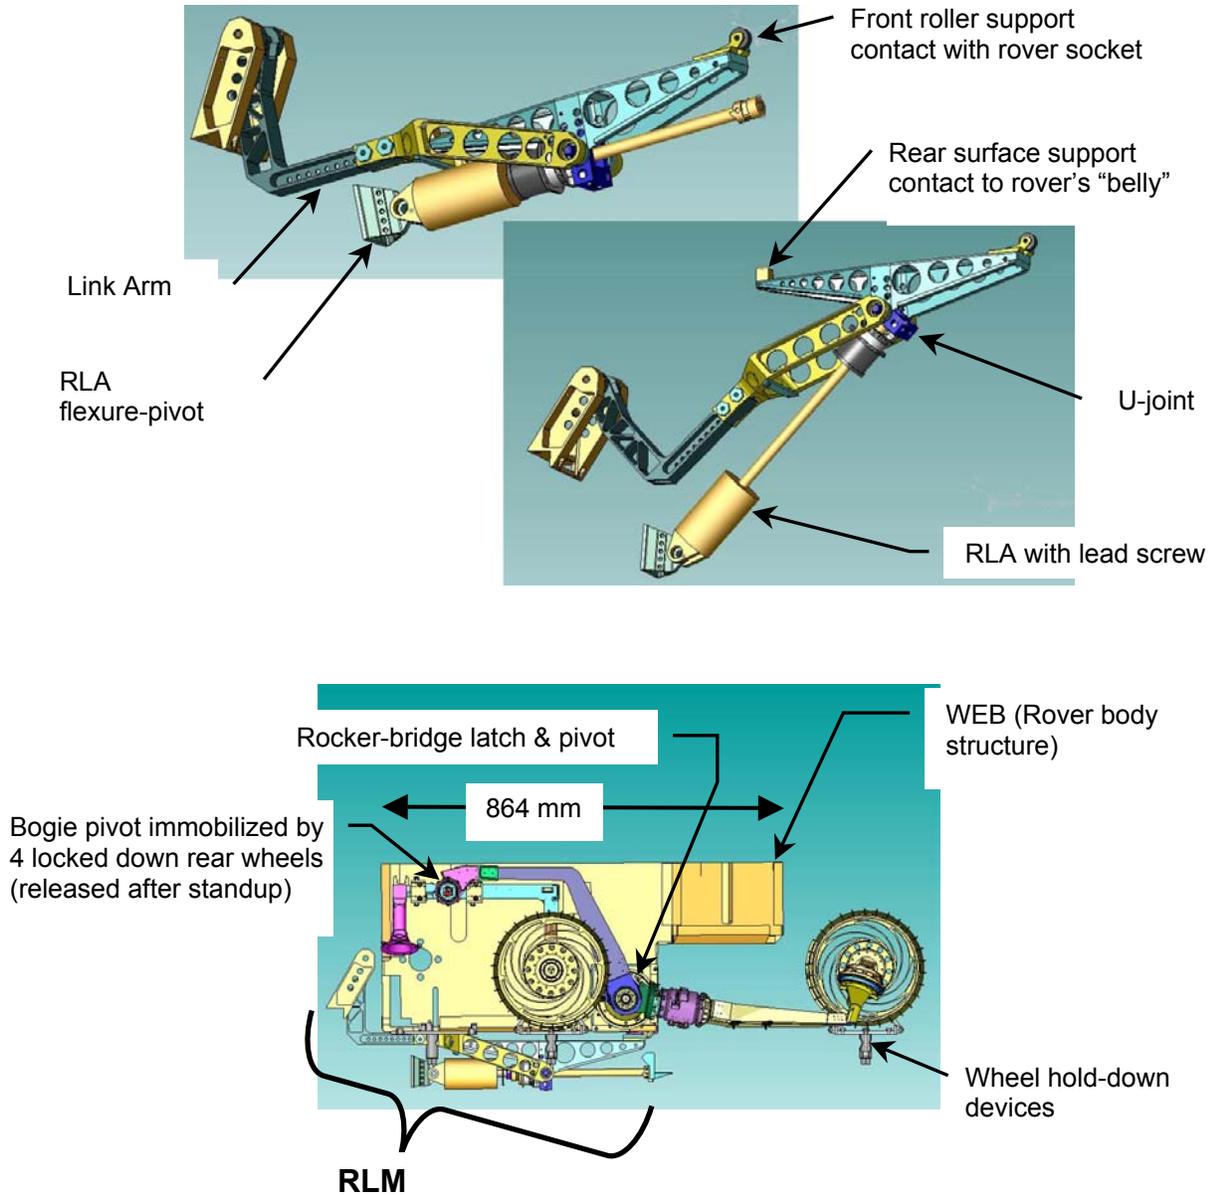
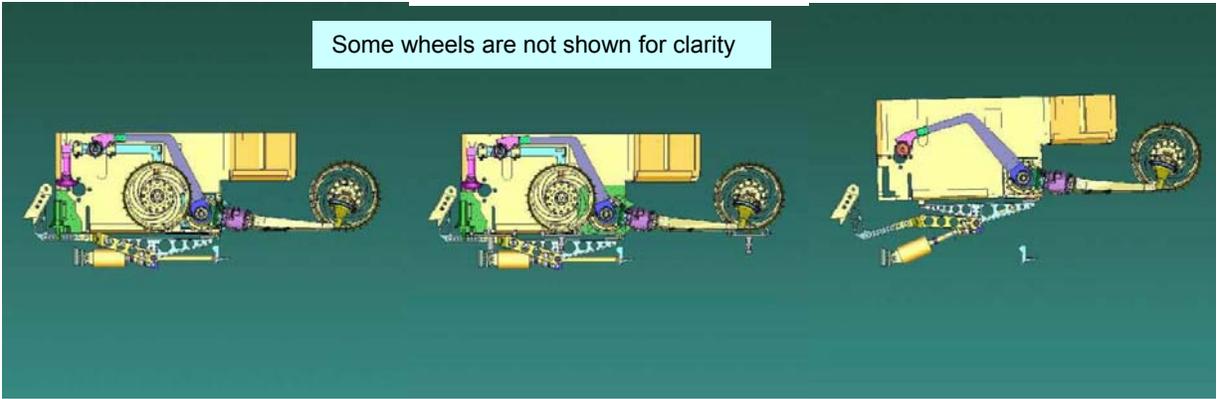


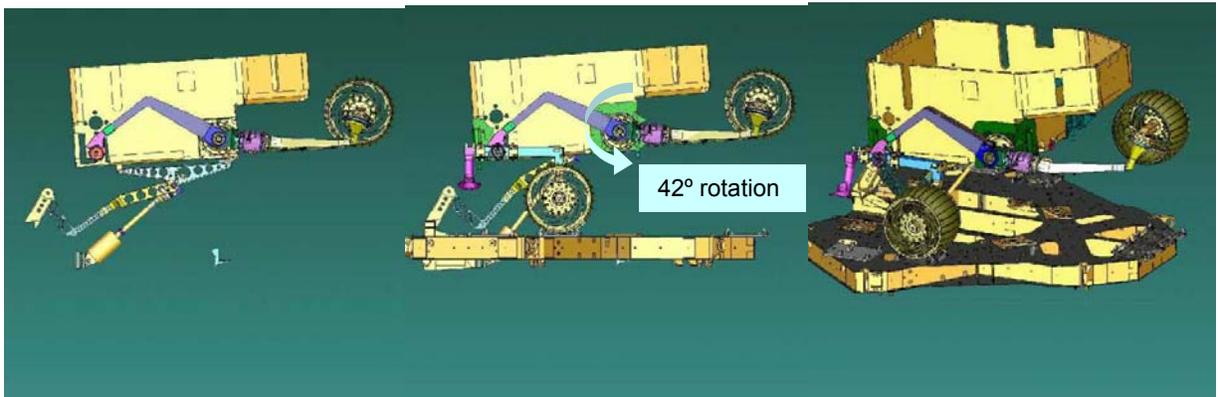
Figure 3

(Above) the RLM and its nomenclature, shown in stowed and deployed configurations. (Below) The rover before pyrotechnic separation devices release, showing the RLM nested within the thickness of the base petal. For clarity, only two of six wheels are shown.

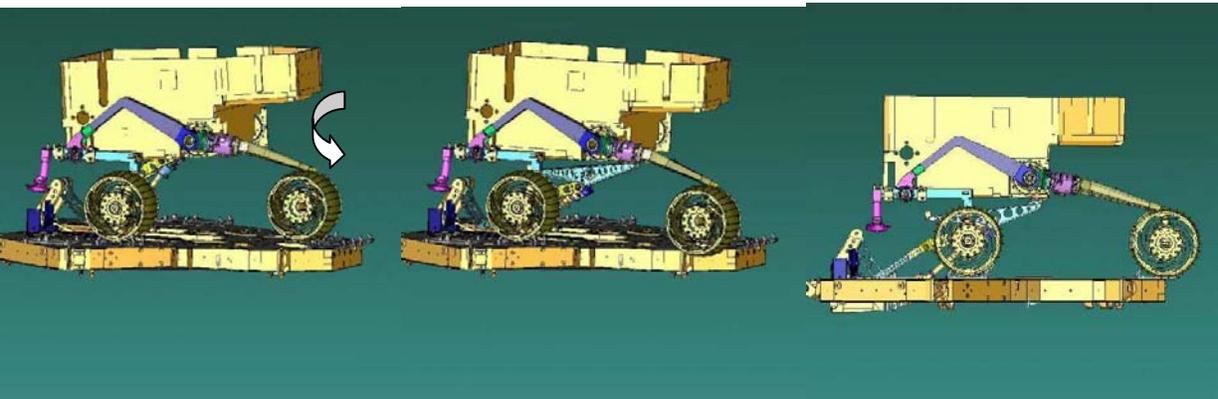


Separation nuts release — springs push off 5mm

RLM begins lifting



RLA nut travels about 244 mm until RLM stops engage. Configuration at maximum deployment



Front bogie & wheel deployment.

RLM lowers, rocker bridge latches engage. RLM continues to lower rover until wheel contacts base petal, then continues to retract and disengage

Figure 4
Sequence of deployments during standup

It was also desirable to accomplish the conversion between kinematic configurations in requirements 2 and 3 without resorting to pyrotechnic devices, as we anticipated testing the engineering unit many times while in a compressed schedule. Safety precautions and other issues peculiar to pyrotechnic devices posed a formidable threat to performing many tests in rapid succession.

The aggressive mass budget and restricted volume for this mission caused the engineering team to adopt an integrated, four-bar linkage approach to the stand up deployment scheme. The existing rover body structure (Warm Electronics Box, or WEB), lander, and the suspension system were utilized to create three of the four links and two of the four pivots, reducing the number & volume of the parts and active devices added to lift the rover.

The RLM functioned in two modes:

- First, it provided the motive force for lifting the rover in addition to functioning as the final link necessary for the four-bar system. That is, it controlled two degrees of freedom (vertical motion and rotation in pitch) when lifting in concert with the suspension system's linkage. In this lifting phase, the RLM supports the rover at two points: a shaped, Torlon pad that contacts the flat rover "belly", and a Torlon roller that is inserted into a specially shaped socket.
- Second, latches in the suspension system engaged in two steps as the deployment reached its apex and the RLM was reversed, thus altering the kinematics of the four-bar linkage by making the bridge joint between the front and rear rockers rigid. This second mode of operation then required the RLM to control only one degree of freedom, vertical motion. The disengagement of rotational (pitch) control is accomplished through passive means, by allowing the rear rover support to simply fall away while the front rover support roller rotates in its socket to permit the now-rigid rover and rocker system to rotate about a pivot on the rear bogie.

Actuator Selection

A trade study early in the design process identified a linear actuator as the preferred means of motivating the lifting process. The chief alternative, torque-actuated joints, was found to require very high torque and correspondingly high moments applied to beams and plate-type structural members. This resulted in an unacceptable mass and volume for a rotary actuator, along with the inefficient use of structural shapes. Only small angular joint excursion was required.

The actuator assembly was subcontracted. A leadscrew type of actuator was assumed (but not specified) in requests for proposal. When pure linear motion is desired from a leadscrew actuator, it is generally preferred that its torque be reacted internally to simplify the design of neighboring components. This would require extra mass, volume, and engineering effort however, so the mechanism's linkage was designed to react the leadscrew's torque. We ultimately selected a linear actuator from Ducommun-American Electronics (AEI); they offered an existing, qualified motor-gearhead with the addition of a leadscrew output module. Ducommun's actuator made use of their electronically commutated (brushless) motor design. JPL had already used another Ducommun motor of very nearly the same design in previous spacecraft, and it matched our drive electronics with its less-common center-tap winding configuration.

The leadscrew and thrust bearing stage represented a new design. The leadscrew, a two-start type with modified acme thread form, was custom made by Nook Industries. The completed Rover Lift Actuator (RLA) was qualification tested at the vendor on a linear-bearing dynamometer that was fitted with cables and weights to permit application of loads up to 4,087 N (920 lbf). The dynamometer was placed in a thermal chamber that allowed testing over the range of -60°C (with the aid of cold LN₂) up to +35°C.

Significant Design Features

When Ducommun proposed a leadscrew-based linear actuator, we required that it be provided with a dry, self-lubricating polymer nut. The reason was that the RLM was likely to be exposed to powerful jets of dust upon MER's landing and multiple bounces on the Martian surface. But with a stroke of about 244 mm (9.61 in) in a crowded space, we could find no viable design for a protective boot around the lead screw. We assumed that a conventional metal nut with wet lubrication would attract dangerous quantities of Martian dust to stick to the exposed leadscrew and seriously diminish the RLM's lifting capacity or jam it altogether. Instead of a flexible boot, dense brush wipers were added to sweep the screw ahead of and

behind the nut. We selected a nut made of Torlon® resin 4203 for its high strength and frictional properties, which we had characterized and found to be adequate in limited testing for other MER devices.

The RLM was required to have integrated hard stops to prevent the application of high force to the suspension system's elements when they meet their end of travel. The uncertainty in required stroke length was recognized early in the design process, so these stops were designed into the RLA's lead screw and made adjustable. The stop "dogs" were custom machined into adjustable steel nuts on each RLA and timed to engage with an appropriate overlap with corresponding dogs on the Torlon nut. Nook Industries machined their thread form into the nuts so that they could be adjusted by screwing back and forth on the leadscrew. One half-dog setscrew on each nut secured its final position and reacted the torque on each stop through shear force.

An early development test was designed to gain confidence in the leadscrew's ability to tolerate dirt. We purchased the closest commercially available item we could find that resembled our flight design: a single-start, rolled-thread acme screw with Polyethylene (PET) nut. We applied loads up to 1,156 N (260 lbf), which represented the highest load anticipated in Mars's gravity. Results were encouraging: torque required increased from 10% to 30%, with smaller increases noted at the highest load.

Most of the early part modeling and analysis was done in SolidWorks® and associated analysis software. However, due to limitations in SolidWorks at that time, the whole spacecraft database was maintained on Computervision® CADD5. When the engineer was satisfied with the general layout of the assembly, it was transferred by STEP or Parasolid file to the system to be integrated with the spacecraft. This translation presented an opportunity for error, and we did discover a significant error in the 4-bar pivot locations when the rover attitude at specific points along the lift trajectory did not match between the CADD5 model and the SolidWorks model. The discrepancy was discovered and remedied by a thorough comparison of the CAD models. We had made early plans to create an Interface Drawing to control swept volumes and critical interfaces, but the complexity of the design problem and the limited time available forced us to simply rely on the CADD5 database as the Interface control tool.

Lesson 1

Complex mechanisms comprising many parts with multiple motion trajectories are much easier to manage in CAD systems that show solid form in real time, that allow quick assignment of motion constraints, and allow real time movement of the system.

Lesson 2

Errors due to poor communication crept into the location of the pivots in the 4-bar link system. The errors resulted in frustration and wasted time when the two models gave incompatible trajectories for the WEB. This experience not only affirms the common wisdom that communication among team members is important, but especially points out that using two independent databases for design work requires careful planning to safeguard consistency of data. Much better still, use only one database.

Loads Analysis

The loads in the RLM were estimated with CAE software, and the effect of joint friction was estimated. Figure 5 shows the forces acting on the interface between the WEB and the RLM during the upward lift process. Two tools were used to evaluate forces: Dynamic Designer®, software that is integrated with SolidWorks, and ADAMS®. Several variable parameters were evaluated: rover/lander attitude with respect to gravity, different joint locations, and link lengths. The effect of increased friction at any one of the system's many pivoting joints could be studied this way as a means of quantifying functional margins. CAE simulation provided insight and understanding of system behavior, helping to design the system standup test plan. An example of analysis output for the linear actuator's force output is shown in Figure 6. Gary Ortiz of JPL performed the ADAMS modeling and documented the system's forces and kinematic behavior.

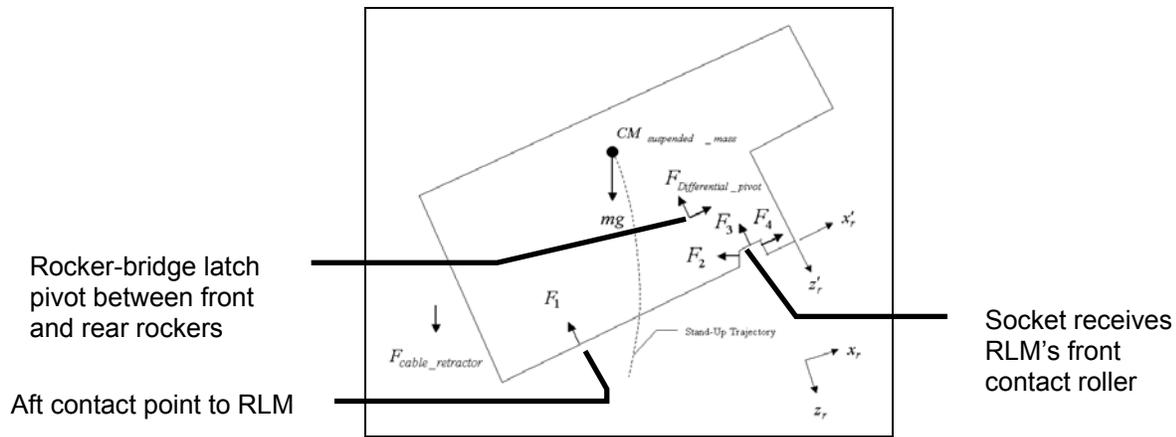


Figure 5

Free body diagram of the rover and its interfaces to the suspension system and RLM

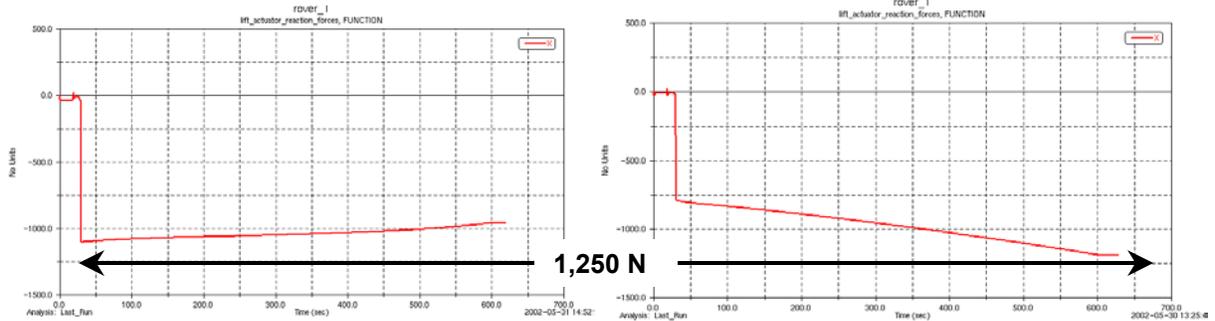


Figure 6

RLA force requirement vs. time for a nominal, flat lander attitude (left) and 20° pitch (right)

One gratifying indication from our analysis was that the maximum force required from the RLA would not be significantly different as a function of rover/lander tilt attitude; although the slope of the force curve would change, no one load case was clearly more difficult than another. We calculated forces at the actuator and joints over the whole envelope of possible deployment attitudes (MER is required to be deployable at any attitude up to 20° from horizontal).

Testing, Failure & Recovery

Our testing schedule demanded that parts must be fabricated before achieving complete confidence in the design. But CAE provided essential advance warning of problems when continuing analysis predicted that the rover's stability was threatened at certain tilt angles within the specification envelope. The RLM depends on a balance of forces due to gravity to hold it in firm contact with the rover, and it was feared that there was insufficient margin to prevent the front contact roller from pulling out of its socket at lander tilt attitudes greater than 20°. Anticipating that this joint would need refinement as a result of testing, the critical interface shape was designed as an easily machined and replaced titanium-alloy insert. The two plots in Figure 7 contrast the estimated contact force on the RLM's roller between the unstable configuration on the left with the remedied configuration on the right, that was designed long before the first test could be executed. This instability was confirmed in a (carefully controlled!) test. The replacement part was ready to be installed immediately, and the test program continued with little delay. Figure 8 illustrates the socket and replaceable insert.

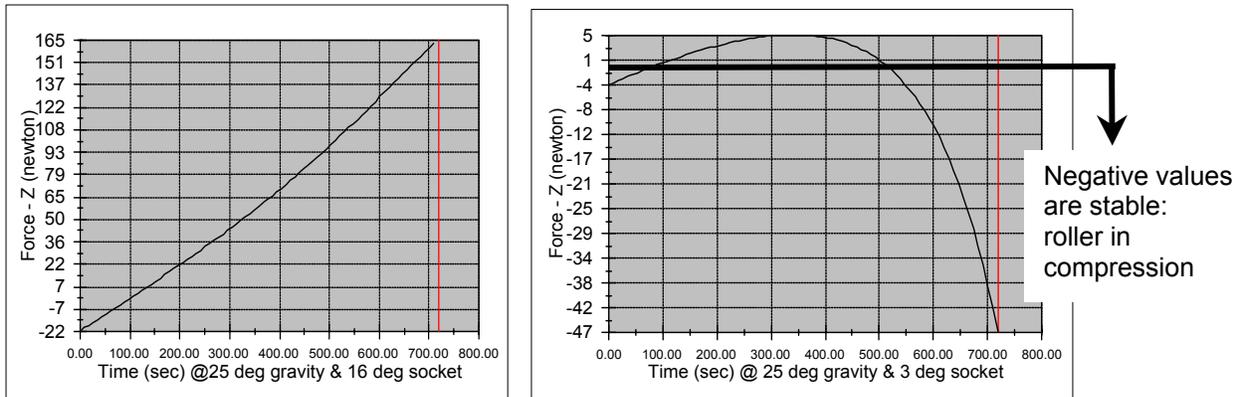


Figure 7

Contact force on the RLA's front contact roller vs. deployment time. Note the pronounced, monotonic transition through zero for the unstable joint configuration on the left. This indicates that the roller has lost compressive contact with its socket



Figure 8

Front roller to socket interface and the replaceable socket insert

Lesson 3

A modeling effort that continues concurrently with hardware assembly and test can prove beneficial, especially if uncertainties can be foreseen and modular components are utilized which can be easily modified or replaced as new information is made available.

Late in the test program, it was found that the real system strained so much in Earth gravity loads that the maximum travel of the RLA had to be extended by approximately 12 mm to cause adequate rotation at the rocker-bridge latch to approximately 42°. When modeling with ADAMS we had attempted to quantify the extra deployment made necessary by compliance and backlash, but the real system was too complex for an accurate estimate. The RLA's lead screw was designed to be long enough to provide adequate margin to allow this extension, although it was necessary to abandon the hardstop's adjustability feature and design a new hardstop with smaller volume to accomplish the stroke increase.

At about this same time, we noted another problem: a crack had opened up at the root of the hardstop that was machined into the Torlon nut. This occurred due to an oversight in the stress analysis; no stress concentrations had been considered in that location. The remedies for the cracked nut and to provide increased range of travel were accomplished in the same redesign; there was room at both extended and retracted positions to mount new steel parts (adding about 0.1 kg) that shunted the stop loads around the nut altogether. The broken Torlon nut is shown in Figure 9. Figure 10 illustrates the original extended stop design and the final version that used all the available screw length. The final design of the nut housing and retraction stop is shown in Figure 11.

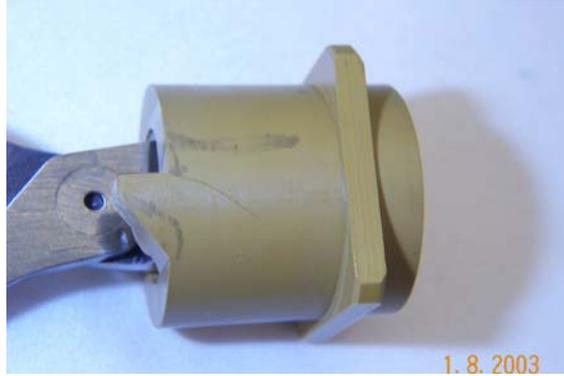


Figure 9

Torlon 4203 nut showing crack propagating from root of hardstop



Figure 10

Original extension stop design (left) vs. the redesigned version with longer stroke



Figure 11

Final nut housing and retraction stop design

Lesson 4

This reaffirms the previous lesson: modularity in design proved to be beneficial. Small, key components could be quickly unbolted and inexpensive replacement designs incorporated in a short time, rather than having to modify an expensive, complex machining with integral features. Try to foresee the need for extended capability and design for it.

Lesson 5

When significant uncertainties exist in the development of a design, try to provide room and features that permit different approaches to accomplishing the same task. There was adequate room at both ends of travel to modify the hardstops so that the original approach was discarded and new hardware fabricated at a late date.

The RLA was specified to work reliably at loads up to 3,685 N (830 lb_f). This would allow testing of the complete rover/ suspension system in Earth gravity without the assistance of off-loading devices. We considered this an important requirement, as it would support the difficult logistics of rapidly testing deployment in the many permutations of attitude and environment. It was also desirable to avoid off-loading devices to prevent possible damage to the delicate rover appendages. However, dynamometer tests on the actuator made it plain that it would not reliably exert the force required to test the lift process in Earth gravity. Instead, we found that, even at room temperature it would frequently stall at loads as low as 2,313 N (520 lb_f). Its capability was, of course, worse at -60°C, the lowest test temperature.

The force at which the RLA stalled was inconsistent from one test to another. In fact, one unit successfully (albeit slowly) lifted a rover in cold thermal-vacuum test, only to fail another test at ambient temperature in a supposedly identical RLM. We also noted that stall generally occurred at a motor speed of 26.0 rad/s (248 rpm) $\pm 5\%$ instead of nearly zero speed as we expected. Following the logic indicated in Figure 12, it was shown to high confidence that the poor and erratic performance was due to variable friction of the Torlon nut on the lead screw. The motor/gearhead put out consistent torque that indicated reasonable gearhead efficiency.

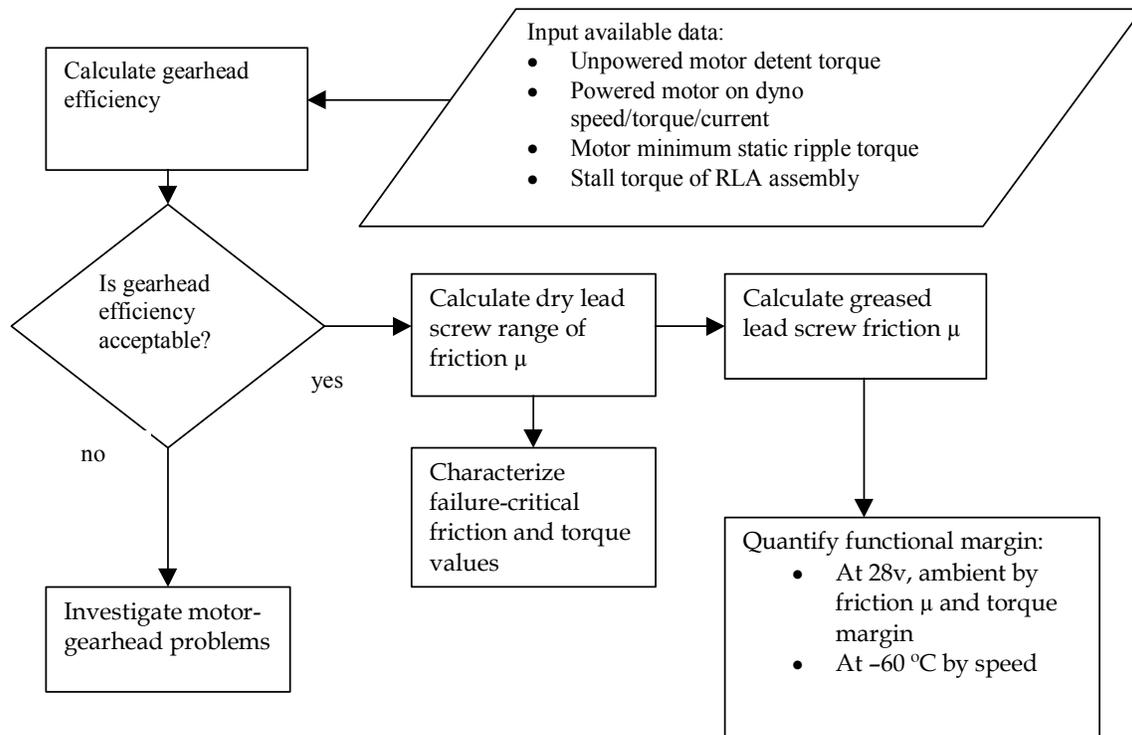


Figure 12
Diagnostic logic for RLA performance

We investigated several remedies. The first solution we attempted was to dry film lubricate the lead screw with a thin film of sputtered-on MoS₂. This process made the friction problem worse! Next, we removed the lubricant film and polished the screw with lapping compound. The Torlon nuts had a noticeably poor finish on their threads, so we tried polishing those too. Polishing improved the RLA's performance briefly, but the gains would not persist through all qualification tests; upon subsequent running-in, the level and variability of friction at the nut increased again.

With performance too poor for Earth gravity operation and schedule almost exhausted, we finally experimented with grease lubrication on the screw. Immediately after a failure to lift the rover in ambient conditions, we injected Braycote® 600 grease into the cavity just in front of the nut. As hoped, the rover deployed on the next try without the slightest trouble. Persuaded to favor grease lubrication by this dramatic improvement, now we had to prove that a blast of dirt would not jam the exposed leadscrew! The brush wipers, mentioned earlier, were made of electrically conductive brass-fibers to facilitate grounding across the interface between the steel lead screw to the u-joint that housed the Torlon plastic nut. They were intended to remove all but a fine layer of dust that might cling to the screw, but now we needed to show that a heavy load of mixed dirt would not pose significant risk to RLM operation. To make matters worse, we noted that some brush fibers came loose in some tests (Figure 13). It then became necessary to show that these fibers themselves could not jam the mechanism.

To prove that the RLM could tolerate stray brush fibers, we repeated a set of RLA dynamometer tests with a greased lead screw. Several brush seals were chopped up to prepare large amounts of brass fibers, and these fibers were applied aggressively to the lead screw before the RLA was activated. To our relief, no amount of effort to intentionally jam fibers of any length into the nut/screw interface succeeded in making a significant (>5%) increase in current required to lift any load up to the maximum test load. Photos of the lead screw and the Torlon nut covered with brush fibers are shown in Figure 14.



Figure 13

Photo showing brass brush fibers that were pulled out of their brush wiper assembly

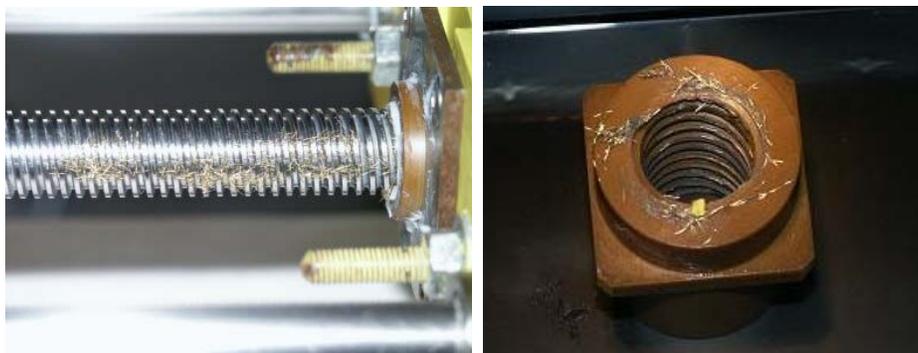


Figure 14

Left, the greased lead screw in the dynamometer covered with brush fibers. Right, the Torlon nut after the test. Performance was unaffected despite the trapped fibers

Finally, many system-level tests were performed over the whole standup sequence. The engineering model of the rover/ suspension/RLM package was tested for the RLM's ability to tolerate a heavy load of dirt. We formulated a mix of quartz particles of various sizes and basalt dirt and applied it to the full length of the RLM's lead screw in its pre-lift configuration. This material application was intended to simulate a worst case caking of dirt onto the leadscrew resulting from airbag impact and rolling. When the RLM achieved maximum lift, a coating of fine "Mars Dust Simulant" was applied to the lead screw in order to simulate the effect of airborne particles that might stick to the lead screw when it's exposed in the extended position. This measure addresses a contingency plan that anticipates the ground team is working on a standup or system anomaly for multiple sols.

Table 1 summarizes the material mixture that was applied prior to lift. Figure 15 shows a picture of this mixture and the appearance of the lead screw with this contaminant before lifting. The lead screw had residual Braycote 600 grease on it from a prior test. The material was forced onto and in between the threads of the screw from all sides through the entire length of the screw. Thus, both the total quantity applied and the level of penetration of the material into the thread depths was felt to be conservative relative to what could be naturally achieved during impact and rolling on Mars.

Table 1
Hard, sharp dirt contamination mixture applied to the RLA's leadscrew

Element	Volume Fraction
Nominal 420-1000 micron crushed quartz	1/3
MIL-E 5007 crushed quartz	1/3
ISIL (basalt) sand, $\leq 5\text{mm}$ particles	1/6
Athena "Mars Dust Simulant"	1/6

Figure 16 shows that a significant amount of material is gathered in the extension brush wipers of the RLM and that a small amount does escape past the brushes to the Torlon nut. In addition to the material entrained by the brushes, a significant fraction of the material is brushed off the lead screw and falls to the ground during operation.

No significant performance degradation was seen as a result of dirt testing, as measured by motor current necessary to raise the rover. Table 2 summarizes current and speed test data from system tests and the safety margins that correspond to those results.



Figure 15

The mixture of quartz and sand applied to the RLA's leadscrew for the dirt tolerance test



Figure 16

Views of the extension brush wiper after the dirt test and the Torlon nut with an accumulation of dirt

Table 2

Final MER 1 and MER 2 RLM/RLA margins as measured in system tests.

Test	Test Temp °C	Test Voltage	Measured stall current [mA]	Predicted stall Velocity, mm/s	Measured current [mA]	Measured Velocity, mm/s	Demonstrated Margin, Earth Weight		Predicted Margin, Mars Weight	
							current draw	velocity	current draw	velocity
MER 1 STT	-55	32.7	785	0.2	350	0.52	2.2	2.6	3.7	3.5
MER 2 STT2	-57	28.5	810	0.2	350	0.44	2.3	2.2	4.3	3.4

The performance margins noted here are robust, even in Earth gravity where wide margin is not required. It can be seen that the key to obtaining consistent performance was to control friction at the nut/screw interface of the RLA, and grease was the appropriate means to achieve this control. This affirms a motto—to quote Doug Packard, an esteemed, veteran engineer at JPL: “Grease is Good”!

Lesson 6

We responded to early indications that the self-lubricating Torlon® nut presented high and erratic friction with measures such as polishing and dry film lubrication. Unfortunately, the best solution we found could have been implemented much sooner if we had not been dogmatic about not flying a greased, “dust magnet” leadscrew. As it turned out, the greased screw was highly tolerant of large quantities of dirt, especially with its brush wiper-shields. If resources permit, it is wise to prepare and test a range of design solutions when significant uncertainties exist in all candidate solutions.

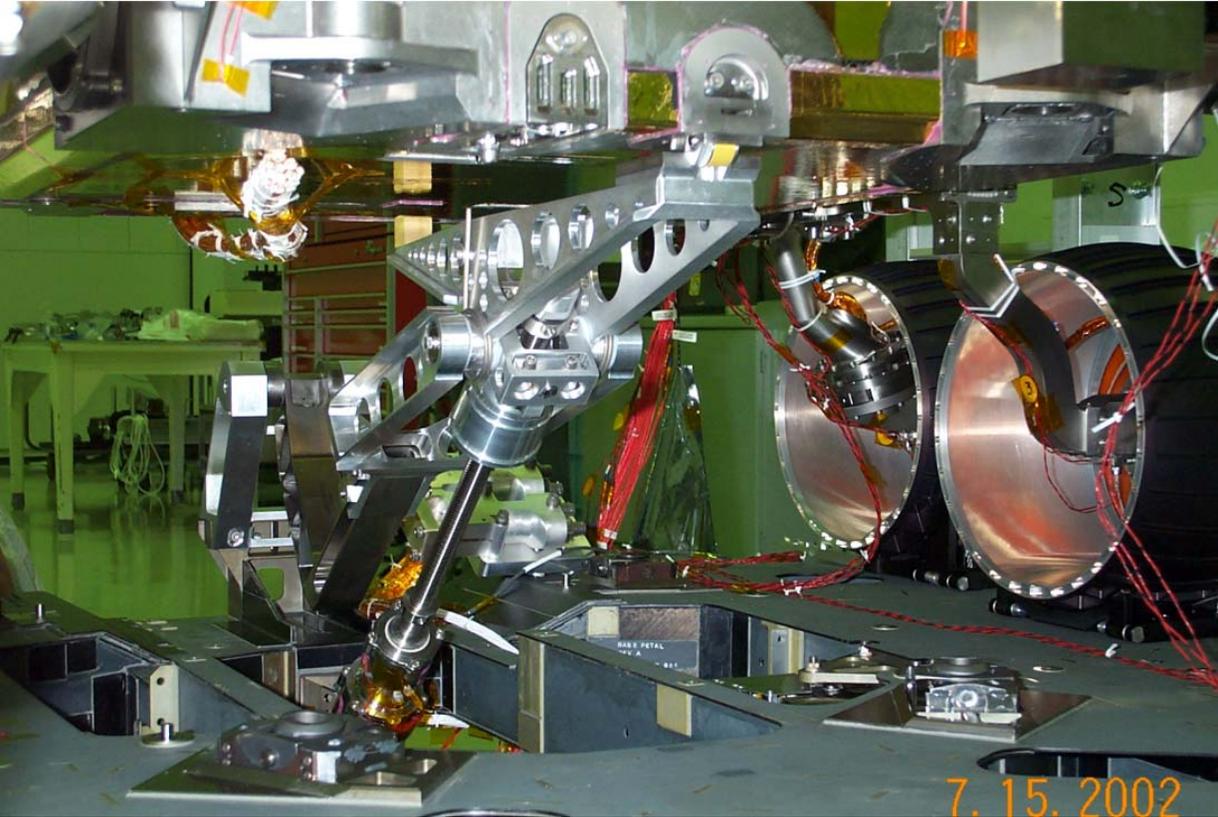


Figure 17
A view of the complete RLM underneath the rover

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.