DIFFERENCED DOPPLER APPLICATIONS FOR MARS EXPLORATION ROVER NAVIGATION

Timothy P. McElrath[(1)], Neil A. Mottinger[(1)], Brian M. Portock[(1)], Darren T. Baird[(1)], Geoffrey G. Wawrzyniak[(1)], Eric J. Graat[(1)]

[(1)]*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109, USA, firstname.lastname@jpl.nasa.gov*

ABSTRACT

One of the navigation benefits of the twin-spacecraft Mars Exploration Rover (MER) project was the opportunity to learn from the first vehicle in order to improve the results for the second. In particular, differenced-Doppler data were collected and analyzed after the operational support of launch and Mars atmospheric entry for the MER-A spacecraft, with the results incorporated into realtime support of launch and atmospheric entry for the MER-B spacecraft. In the case of launch, differenced Doppler from two stations within the same tracking complex proved to be as useful as angle data in improving the post-launch single-complex orbit solutions. For atmospheric entry, differenced Doppler with intercontinental baselines proved useful in measuring the atmospheric accelerations, such that the landing site location could be significantly improved over the a priori results.

1. INTRODUCTION

The Mars Exploration Rover project launched two nearly identical spacecraft to Mars on June 10 and July 8, 2003. While many more details of the spacecraft (and their navigation) are given elsewhere [1-4], for the purpose of this paper it is only important to know that the spacecraft were spin-stabilized with an antenna that was offset from the spin axis, that the telecommunications system operated only at X-band (8400-8450 MHz) downlink frequencies, and that the spacecraft did not carry an ultra-stable oscillator (which would have been suitable for navigational use). The first and second spacecraft launched will be referred to as MER-A and MER-B, respectively. The two spacecraft arrived at Mars on January 4 (for MER-A) and January 25 (for MER-B), 2004. During both launch and atmospheric entry the MER spacecraft used low-gain antennas (LGAs), which were mounted on the cruise stage or the back of the aeroshell and provided reasonable coverage over all expected spacecraft attitudes. The spacecraft transmitter operated essentially continuously from launch vehicle separation through landing on Mars, before going into an intermittent duty cycle on the surface to conserve power.

2. LAUNCH

2.1 Requirements

The primary goal of navigation for interplanetary spacecraft immediately following launch is to improve the orbit knowledge enough to allow ground stations to acquire the spacecraft signal within the main beam of the antenna. Although the trajectory of the spacecraft will certainly need to be corrected, this can typically wait for several days in all but a tiny fraction of cases which hover just above launch vehicle failure. For X-band radio systems, the pointing requirement is to be within 0.032 degrees of the boresight (which amounts to a 3 db loss) for the 34-meter diameter antennas of the NASA/Deep Space Network (DSN). This requirement is never met by the expected launch vehicle injection covariance, so the first pass necessarily uses an acquisition aid in the form of a ~1-meter X-band antenna attached to a 26-meter S-band antenna (used primarily for Earth orbiter support). The acquisition aid system provides angle observables which can be used to correct the pointing of the 34-meter antennas as long as the radio link margin allows (typically only a few hours).

2.2 Tracking Data

Once the 34-meter antenna has acquired the signal, Doppler data is available, initially one-way, and then two-way or three-way (different transmitting and receiving stations) when an uplink is transmitted for the spacecraft to transpond coherently. Range data can also be obtained after the uplink is established, but only by the transmitting station. In addition to Doppler and range, the angle data from the acquisition aid is available to navigation, in the form of the angular position of the two antenna axes, which for the 26-meter antennas are aligned with the local North and East directions when the antenna is pointing to zenith.

The Doppler data collected by the DSN antennas at X-band is accurate to at least 0.1 mm/sec of one-way range rate, but the navigationally-poor quality of the spacecraft oscillator made this available only when the spacecraft was in a coherent mode. Of course, differencing Doppler data between tracking stations removes the effects of the spacecraft oscillator, as well as any other spacecraft-generated data signatures. One such effect is the spin signature, which at launch was very close to 12 rpm for both vehicles. The Doppler data were compressed to 10 seconds to remove most of the spin effect, but there were still small spin and nutation signatures in the resulting data. As a result, the coherent Doppler data weight was conservatively increased to 2 to 3 mm/sec, to cover the envelope of the remaining signature, and one-way Doppler was not used. After initial acquisition of both vehicles, the spacecraft team spent a navigationally unpleasant amount of time (30-40 minutes) pondering telemetry before any uplink was established, thus limiting the amount of early Doppler and range data available. Due the circularly polarized antennas on the spacecraft, the Doppler data were biased proportionally to the spacecraft spin rate, with a sign that has always defied prediction (but been easy to determine post facto).

Range data is acquired by transmitting a sequential code at a bandwidth of up to ± 1 MHz. The ranging noise is strongly correlated to link performance, which results in very small noise values at launch. The range data weight was 1 meter (one-way), with a 1-2 meter range bias assumed for each tracking station. Range points were acquired at a rate of about once per minute.

Operational robustness dictates that launches are supported by at least two tracking stations at each complex, which provides the opportunity to difference the Doppler data between them. While the baseline formed by the station location difference can be as small as about 250 meters, the advantage of such data is that all spacecraft signatures (due to spin and oscillator instability if one-way) and most transmission media signatures cancel, leaving almost completely thermal noise. This data can be obtained as soon as the downlink signal is tracked at two stations, regardless of the coherency state of the spacecraft transponder, and thus much sooner than typical two-way Doppler and range. The differenced-Doppler data within a complex was weighted at 1 mHz (over 60 seconds), or 0.036 mm/sec of one-way velocity along the station baseline. All DSN stations at the same complex use the same frequency reference, so no bias correction or estimation was necessary at this data noise level.

The angle data system on the 26-meter antennas was originally designed to operate at S-band, and mostly still does so. The X-band system suffers from both fewer targets to update the angle calibrations, and lower signal strength on the much-smaller antenna. Consequently, angle data are treated by estimating a bias in each axis with a 0.1 degree a priori sigma, and weighting the data at 0.1 degrees for a sample rate of one per 10 seconds. Although portions of angle data may display better characteristics, over a whole pass there are typically higher-order signatures present at levels approaching the bias and data weight, which do not do a good job of representing such signatures in the estimation process. Consequently, solutions including angle data often have larger errors than would be considered statistically likely (i.e. more than 1 sigma).

2.3 Estimation Process

The uncertainties in the spacecraft state are the dominant errors, so a simple least-squares filter was used for all launch estimates, with the state and any data biases as parameters. In addition, the effect of errors in Earth orientation, station locations, and transmission media calibrations was added without allowing these parameters to change (or in JPL parlance, these parameters were "considered"). Typically the state was assumed to have an effectively infinite a priori uncertainty, since more realistic uncertainties were both hard to develop and still too large to be useful. The important data biases included those for angles, Doppler (due to spin), and range (station delay calibration), but none for the differenced Doppler. The resulting estimates were mapped to a variety of epochs and coordinate systems, most importantly including the rise time of the second station in spherical coordinates. Since this gives the geocentric angular uncertainty, it closely approximates the expected pointing error for a station tracking at that time due to the significant

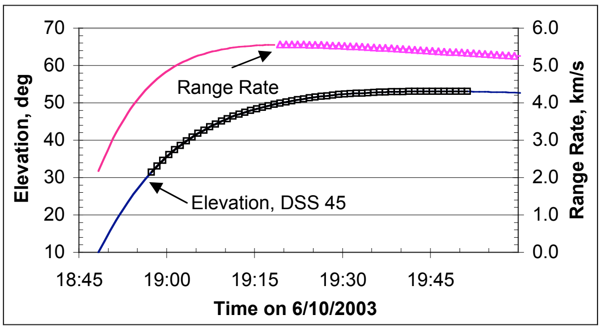

Fig 1. MER-A Launch Canberra Range Rate, Elevation

geocentric range growth by the end of the first pass.

2.4 MER-A Results and Analysis

The initial acquisition complex for MER-A was Canberra, which consists of two 34-meter antennas (DSS 34 and 45) and a 26-meter antenna (DSS 46), as well as a 70-meter antenna that isn't used for launch. In this case the two 34-meter antennas acquired almost simultaneously at about 10 degrees elevation, while

DSS 46 didn't produce useful angles until 9 minutes later (at 30 degrees elevation). Figure 1 shows the range rate and elevation profiles for this pass. The symbols indicate when each data type was available, which shows that the uplink was not established until about 30 minutes after rise. By this time much of the range-rate signature had vanished, making the Doppler data less useful. The angle data also became unusable after 54 minutes (before the end of the plot in Figure 1) due to non-physical jumps.

Although only the first 72 minutes of tracking data (and the first 40 minutes of 2-way data) is shown, several hours of Canberra range and Doppler, along with the available angles, were used to update the pointing predicts for the second station (Madrid in this case). Madrid acquired successfully, and the geometric strength of the two-station tracking data quickly removed the remaining orbit estimate errors, to the degree that pointing predicts were never an issue again (all of which is typical after launch).

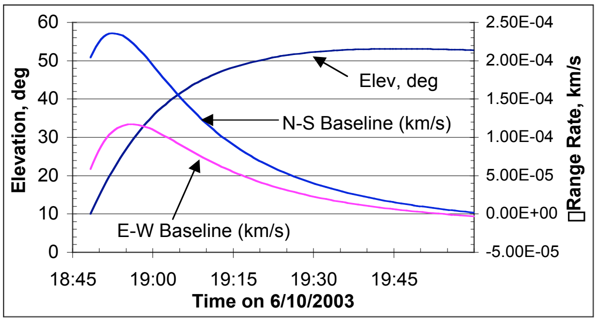

Fig. 2. MER-A Launch Canberra □-Range Rate

Differenced-Doppler data were not used operationally for MER-A launch, but the relatively poor performance of the angle data caused additional data types to be considered, of which differenced-Doppler data was the most obvious (as well as being the only one available). Before MER-B launch, procedures for using differenced-Doppler data were demonstrated on the MER-A data, with results that were better than those using angle data (both in combination with range and Doppler data). For this paper additional analysis was performed to study the actual and potential usefulness of differenced Doppler combinations. The results are given for the data span shown in Figures 1 and 2, since even though longer data arcs are typically available, the usefulness of angles and differenced Doppler falls off quickly with increasing geocentric range.

The actual 34-meter station baseline at Canberra (between DSS 34 and 45) lies almost exactly in an East-West (E-W) direction, with a length of 389 meters. A hypothetical equivalent North-South (N-S) baseline was simulated for the same data span. As can be seen from Figure 2, the N-S baseline has almost twice the signature of the actual E-W baseline, due to the MER-A groundtrack, which headed eastward as it passed well north of Canberra before turning through North on the way to the outbound asymptote trace (close to the equator for a low-declination departure).

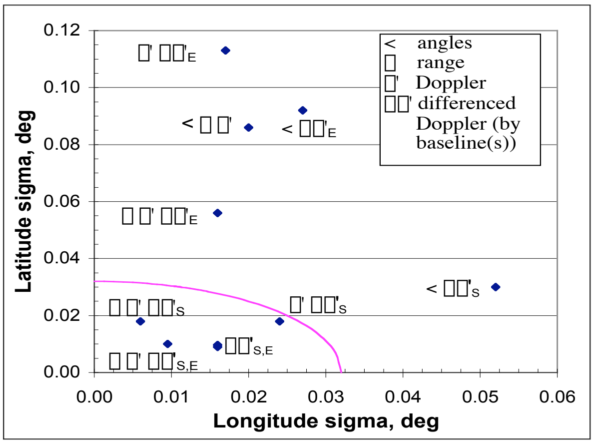

Fig. 3. MER-A latitude/longitude sigmas by data type, at second station (Madrid) rise

The resulting latitude and longitude sigmas at Madrid rise for various data type combinations are shown in Figure 3. Range and Doppler alone for this short arc has a combined uncertainty of nearly a degree (and so is off the chart). Note that the antenna pointing requirement is shown as a circle at 0.032 degrees about the origin. Although neither combination meets the requirement, actual differenced-Doppler (E-W baseline, denoted $\square\square'_{E}$) with range and Doppler is much better than angles (<) with range and Doppler. The N-S differenced Doppler ($\square\square'_{S}$) with range and Doppler meets the requirement, as does differenced Doppler alone from both baselines simultaneously ($\square\square'_{S,E}$). Angles with either baseline are similar to angles with range and Doppler in total error. It is interesting to note that removing range from solutions with range, Doppler, and one baseline show that range appears to be measuring the angular component not provided as well by the remaining differenced-Doppler baseline.

2.5 MER-B Results and Analysis

MER-B launch was the opposite of MER-A in many ways: the launch took place at night on the second opportunity for that day, and following a long coast, the initial acquisition was at Goldstone for a pass that went much more overhead, at a time closer to spacecraft separation (and hence angularly faster). This resulted in a delay in signal acquisition of about 10 minutes from the Goldstone rise.[1]

[1] The first (by a few seconds) post-separation acquisition of MER-B was obtained visually (aided with binoculars) from JPL.

Although the navigation team used differenced Doppler operationally for MER-B, the station visibility fortuitously included a short Madrid pass in the middle of the Goldstone pass. A Madrid 34-meter station (DSS 54) was scheduled fairly late in the planning process, and was able to acquire without any help from a 26-

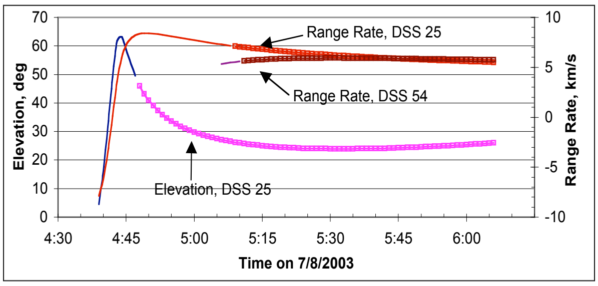

Fig. 4. MER-B Launch Goldstone Range Rate, Elevation

meter antenna (since the 26-meter antenna at Madrid has no X-band acquisition aid) and provided 3-way Doppler data. The strong geometry of the Madrid Doppler data combined with Goldstone Doppler and range was all that was necessary to meet the second station (Canberra) pointing requirements with significant margin. However, the other data type combinations were used to help validate the results, and thus were still useful.

Figure 4 shows the elevation and range rate profile for Goldstone and range rate profile for Madrid, with symbols marking the start of the tracking data (angles and differenced Doppler started at the same time, shown with the elevation symbols). The uplink started 20 minutes after the initial 1-way tracking, which was sooner than MER-A relative to signal acquisition and separation from the launcher. Although all of the early data were used in operations, for the results in this paper the start of the two-way data was delayed until the same separation-relative time as MER-A, or about 16 minutes to 05:25, to improve the basis for comparison between the two.

The DSN complex at Goldstone (in California) consists of four 34-meter antennas, as well as a single 26-meter and 70-meter antenna. The two antennas scheduled for MER-B launch were DSS-25 and 24, which are only 258 meters apart, with a mostly N-S separation, so this was not optimum for using differenced-Doppler data, as will be reflected in the results later. One of the 34-meter antennas (DSS 15) is 10 kilometers (mostly north) from the other three, making it the ideal station for use with any other for obtaining the best possible differenced Doppler, as long as media and spacecraft effects don't start to interfere (which they probably would not, but which would need to be confirmed).

As with MER-A, simulated stations 389 meters north and east of DSS-25 were simulated, to evaluate the usefulness of each baseline. Figure 5 shows the differenced Doppler signature for each simulated baseline. Note that most of the signature is before the start of the actual differenced-Doppler data, although the magnitude from that point on is similar to MER-A. The large early signatures (which were hard to acquire

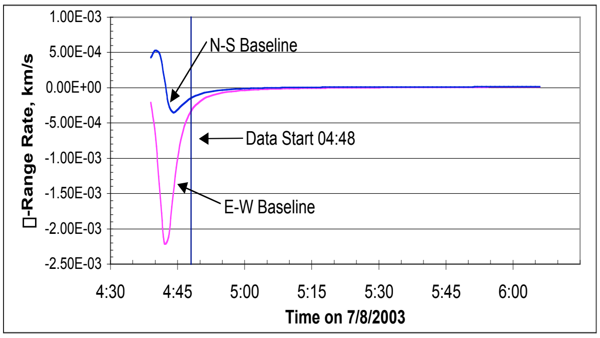

Fig 5. MER-B Launch Goldstone Simulated □-Range Rate on 389-meter Baselines

due to the high coincident angle rates) are probably characteristic of nearly-overhead passes, but the sensitivities to such passes bears more investigation.

The results for various real and simulated data types are mapped to the second station rise (Canberra in this case, since Madrid did not get any improved predicts) latitude and longitude uncertainties in Figure 6. Several differences are apparent compared to MER-A. The angles, range, and Doppler solutions are quite effective,

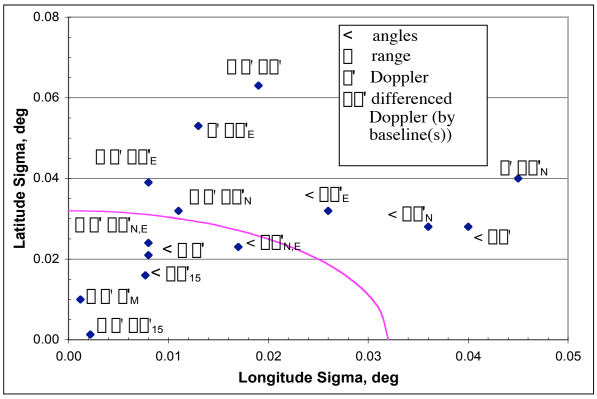

Fig. 6. MER-B latitude/longitude sigmas by data type, at second station (Canberra) rise

and more so than any combination of a single short baseline of differenced Doppler with range and Doppler. Differenced Doppler alone from both 389 meter baselines is off the chart at 0.171 degrees, and range and Doppler alone is ten times worse. Angles and both differenced-Doppler baselines would have met the pointing requirement. The strength of the operational solution using Madrid three-way ($\square'_M$) with Goldstone Doppler and range is clear. The long baseline (DSS15) differenced Doppler ($\square\square'_{15}$) combined with range and

Doppler provides the best combination on the chart, which shows that even a few kilometers of baseline can provide powerful geometry.

2.6 Launch Conclusions

Clearly differenced-Doppler data from the existing stations of the DSN is useful during initial acquisition passes, although the geometries where it is less effective (as seen for MER-B) need to be understood better. In addition, the DSN is moving towards arrays of smaller antennas, and so understanding the benefits of differenced Doppler as a function of baseline orientation and size is more than a theoretical exercise. Until arrays are implemented, project navigation teams should request the longer baselines at the initial acquisition complex where possible (primarily Goldstone) and not neglect to schedule any additional DSN complexes or available non-DSN tracking stations that have viewperiods during the initial acquisition pass.

3. ARRIVAL

The rotational phase of Mars at the arrival time of the MER spacecraft fortuitously aligned the selected landing sites with the project's desire to have multiple DSN complexes in view during the critical entry, descent, and landing (EDL) events. Consequently, both EDLs were visible from the Goldstone and Canberra complexes, which trained all of their antennas on Mars

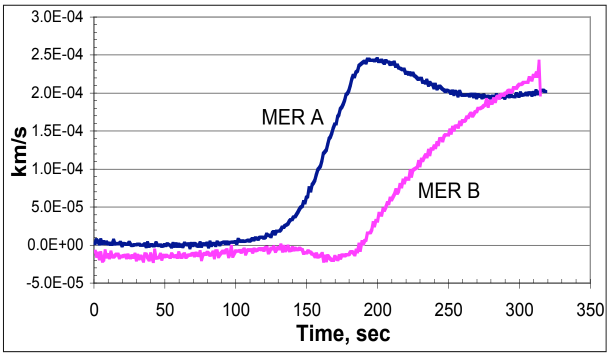

Fig. 7. Differenced-Doppler data for MER-A/B EDL

for the event. However, the only stations expected to have a chance of maintaining carrier lock throughout most of EDL were each complex's 70-meter antenna. The spacecraft signal switched to 1-way over an hour before entry, to avoid mode transitions during uplink interruptions caused by staging events on the spacecraft. At the entry attitude, the spacecraft spin axes (and consequently the LGA boresights) were 22 and 39 degrees off Earth point, respectively, for MER-A and MER-B.

The relatively weak signal and unpredictable dynamics prompted predictions that even the 70-meter antennas would not be able to track in a closed-loop mode. However, both antennas tracked all the way down to parachute deploy for both vehicles, at which point the frequency predict error (mostly from the expected variations in the parachute deploy time) caused both stations to drop lock, providing a redundant indication that the parachute had deployed. This outstanding DSN performance prompted glee on the part of the navigation team during both EDLs, as it provided another set of differenced-Doppler data, the only data type capable of improving the estimate of the landing site without receiving telemetry data from the surface. Although the one-way data oscillator uncertainties precluded more than a qualitative detection of atmospheric entry, the differenced-Doppler data, though unable to improve the spacecraft state at entry, were able to constrain the error growth in at least one component of the velocity due to uncertainties in the Martian atmosphere, resulting in a much improved position estimate on the surface. Figure 7 shows the differenced-Doppler signature for MER-A and MER-B. The data have had a linear trend removed (based on the first 100 seconds plotted), and so the result can be taken as the effect of encountering the atmosphere. The signature for MER-A is clearly more dynamic than that of MER-B, which is reflected in the results below.

3.1 MER-A Landing

Although the potential of differenced-Doppler data was known before MER-A EDL, it was assumed that significant post-processing would be required to obtain an estimate. However, a few hours after the successful landing (and related celebration), one of the authors attempted a fairly simple estimate and obtained useful results. Starting from converged trajectory from the last

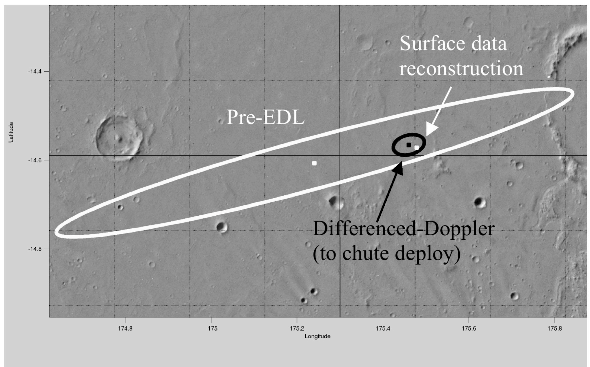

Fig. 8. MER-A landing site estimates

the days, the spacecraft was modelled as a sphere travelling through the MarsGRAM atmosphere model. Preliminary work had established an approximate constant value of the drag coefficient (C_D) to match the landing locations obtained from higher-fidelity models. C_D was then estimated as both a constant and a stochastic parameter, to account for average and time-varying differences from the nominal model, respectively. In addition to atmospheric parameters, it was necessary to estimate a bias in the differenced-Doppler data to account for differences in the frequency standards between the Goldstone and Canberra complexes. The differenced-Doppler data was weighted close to the RMS, which was about 0.045 Hz (for 1 second data).

The resulting estimate was mapped to the time of the end of the data arc, corresponding to the time of parachute deploy. The formal errors were under 1 kilometer, but of course did not include any effect of winds or surface bouncing, and assumed that the motion after parachute deploy was straight down (whereas actually the spacecraft continued downtrack for at least several hundred meters). All of these effects were judged to add up to a surface location and 3σ uncertainty of –14.563° ±1.5 km latitude, 175.459° ±2.4 km longitude. The locations and error ellipses for the pre-EDL estimate, the differenced-Doppler estimate, and the final reconstruction estimate (based on surface tracking from the Earth and the Mars Odyssey orbiter) are shown in Figure 8. Note that the final surface estimate and the differenced-Doppler estimate differed by about two sigma, which is not too surprising given the relative lack of study into the actual error contributions for this scenario.

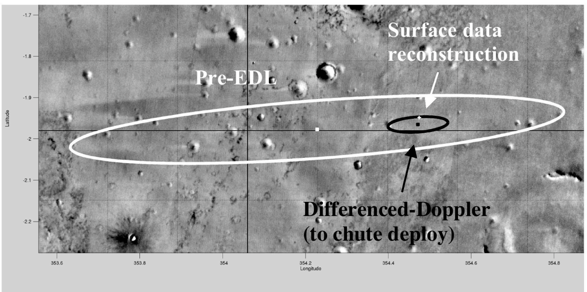

Fig 9. MER-B landing site estimates

3.2 MER-B Landing

Following the MER-A experience, the process used was analyzed to try to make a priori use of information from higher-fidelity models, and to try to improve the timeliness of the estimate. The adopted process was to use C_D values as an a priori model, both before and after parachute deploy, and to map the result to the altitude of the local surface. This accounted at least approximately for motion after the parachute deployed, as well as C_D changes during supersonic flight, and allowed tighter constraints on the earlier drag profile. The weaker differenced-Doppler signature, while not recognized at the time, probably made better modelling more important than for MER-A. The differenced-Doppler data also were slightly noisier than MER-A's, and a weight of 0.073 Hz (at 1 sec) was used.

These preparations allowed a landing site estimate and 3σ error ellipse of –1.965° ±1.2 km latitude, 354.471° ±4.3 km longitude, to be generated well within one hour of the time of the MER-B landing, and then displayed for the rest of the JPL flight team. Figure 9 shows the pre-EDL, differenced-Doppler, and final reconstruction estimates. Even with the larger uncertainties in the differenced-Doppler estimate (compared to MER-A), the final reconstruction differes by almost 3 sigma. Clearly more understanding of the errors affecting these sorts of estimates is needed to improve their statistical consistency.

4. CONCLUSIONS

In retrospect, it is remarkable that the beginning and end of the flight of the MER spacecraft was significantly aided by using differenced-Doppler data. The beginnings of an understanding of the effects of differenced-Doppler data during initial acquisition have been developed in this paper, and the possibilities of this data type's use for EDL and other events with high velocity uncertainty have been demonstrated.

5. ACKNOWLEDGEMENTS

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

References

1. D'Amario L. A. "Mars Exploration Rovers Navigation Results," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA-2004-4980, AIAA, Washington, DC, 2004.

2. Portock B. M., Graat E. J., McElrath T. P., Watkins M. M., and Wawrzyniak G. G., "Mars Exploration Rovers Cruise Orbit Determination," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA-2004-4981, AIAA, Washington, DC, 2004.

3. Knocke P. C., Wawrzyniak G. G., Kennedy B. M., Parker T. J., Golombek M. P., Duxbury T. C., Kass D. M., and Desai P. N., "Mars Exploration Rovers Landing Dispersion Analysis," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA-2004-5093, AIAA, Washington, DC, 2004

4. Baird D. T., McElrath T. P., Watkins M. M., D'Amario L. A., Guinn J. R., Portock B. M., Graat E. J., and Wawrzyniak G. G., "Mars Exploration Rovers Entry, Descent, and Landing Navigation," *AIAA/AAS Astrodynamics Specialist Conference*, AIAA-2004-5091, AIAA, Washington, DC, 2004.