The Effects of Clock Drift on the Mars Exploration Rovers

Khaled S. Ali[1] and C. Anthony Vanelli[2]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA,91109

All clocks drift by some amount, and the mission clock on the Mars Exploration Rovers (MER) is no exception. The mission clock on both MER rovers drifted significantly since the rovers were launched, and it is still drifting on the Opportunity rover. The drift rate is temperature dependent. Clock drift causes problems for onboard behaviors and spacecraft operations, such as attitude estimation, driving, operation of the robotic arm, pointing for imaging, power analysis, and telecom analysis. The MER operations team has techniques to deal with some of these problems. There are a few techniques for reducing and eliminating the clock drift, but each has drawbacks. This paper presents an explanation of what is meant by clock drift on the rovers, its relationship to temperature, how we measure it, what problems it causes, how we deal with those problems, and techniques for reducing the drift.

Nomenclature

ET = Ephemeris Time
FSW = Flight Software
HGA = High Gain Antenna
IDD = Instrument Deployment Device
IMU = Inertial Measurement Unit
MER = Mars Exploration Rover
REM = Rover Electronics Module
SCLK = Spacecraft Clock
sol = a Martian day
UHF = Ultra-high Frequency

I. Introduction and Background

The Mars Exploration Rovers (MER) are a pair of mobile robots, named Spirit and Opportunity, equipped with instruments to examine the Martian geology (see Fig. 1). Spirit launched on June 10th, 2003 and landed on Mars on January 4th, 2004. Opportunity launched on July 7th, 2003 and landed on January 25th, 2004. The prime mission for each rover was to last for 90 sols, or Martian days, but they have far outlived that duration. Spirit roamed the Martian surface until the mission officially ended in May of 2011. Opportunity has operated on Mars for over eight years and continues to do so. Both rovers have made important discoveries about the history of water on Mars and the Martian geology[1,2].

Like many machines, the MER rovers need to know the time. They need to know the absolute time, to conduct certain activities, such as communication windows, at the time the mission operators on Earth expect. They also need to know about the relative passage of time, to determine when to do things relative to some other activity, such as when to stop an activity that might be running too long or how long to allow the rover to "sleep" before waking it up.

The MER rovers keep track of the absolute time in Spacecraft Clock (SCLK) time. SCLK time is the rover's estimate of the current Ephemeris Time (ET) (a.k.a. Barycentric Dynamical Time (TDB))[3,4], expressed as ephemeris seconds past the J2000.0 epoch.

---

[1] Robotics Software Engineer, Mobility and Robotics Section, 4800 Oak Grove Drive/Mail Stop 264-528.
[2] Senior Engineer, Guidance and Control Section, 4800 Oak Grove Drive/Mail Stop 264-828, AIAA Member.

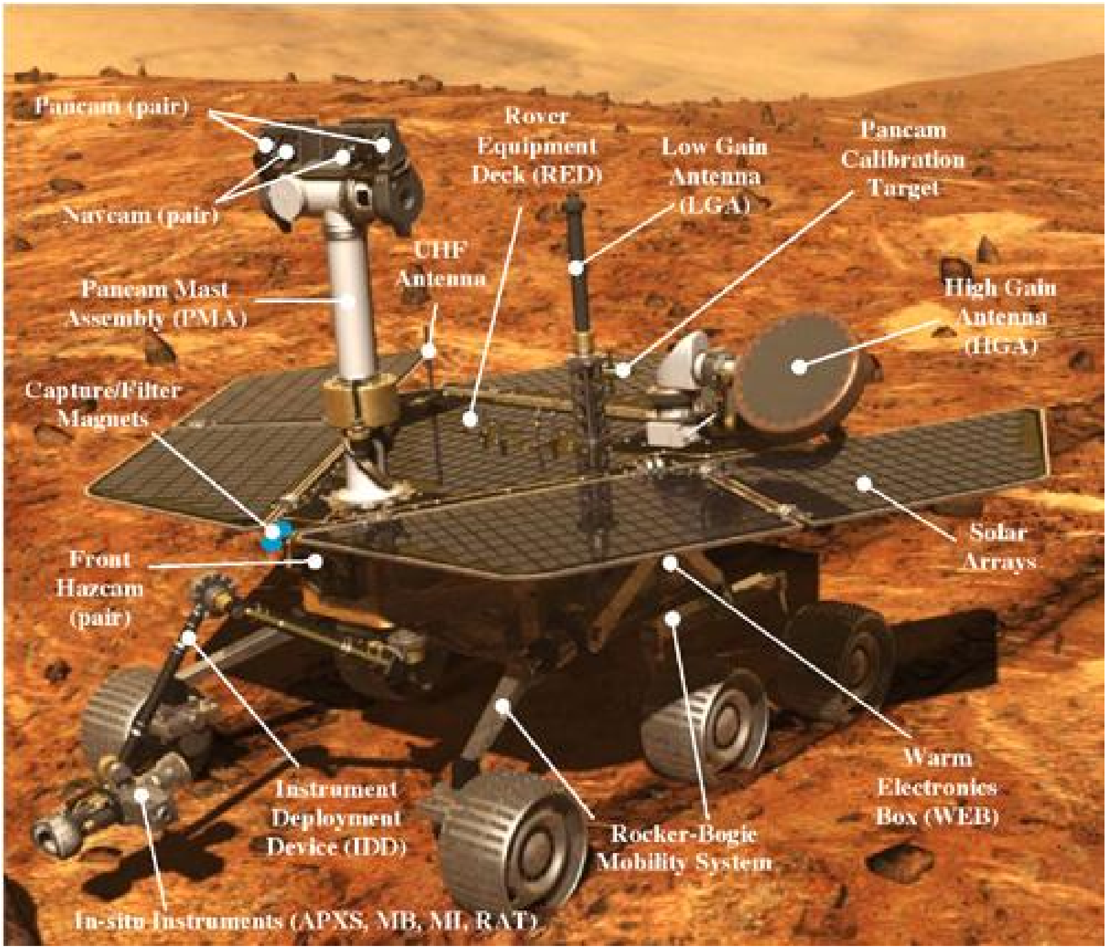

Figure 1. An artist's rendering of a Mars Exploration Rover, with the external components labeled.

The MER rovers each have three clocks to track the passage of time. The clocks are called the Mission Clock, the Spacecraft Clock, and the Alarm Clock. The Mission Clock and the Spacecraft Clock keep track of the time in SCLK Time, monotonically counting upwards from the time they were set before launch. The Mission Clock is powered by the rover batteries. As long as the batteries provide enough power, the Mission Clock will countinue counting the seconds. The Spacecraft Clock is volatile and is seeded from the Mission Clock each time the Flight Software (FSW)[5] boots up from a cold state. Nominally, the rovers perform a cold boot each Martian morning, after remaining in a "sleep" state throughout the Martian night. The FSW uses the Mission Clock only to seed the Spacecraft Clock after each cold boot, and thereafter uses the Spacecraft Clock when it needs to know the time.

The Alarm Clock is a countdown timer. The FSW sets it with a number of seconds before going to sleep, overwriting the current value. The Alarm Clock counts down until it reaches zero seconds, at which point, it "rings". Then the bus is powered up and the FSW wakes up again. After reaching zero seconds, the Alarm Clock rolls over to 27 hours again and continues counting down.

II. Clock Drift

The SCLK time on each rover is the *rover's* estimate of how many seconds have passed since the start of the J2000.0 epoch. Since the Mission Clocks on the two rovers were set before lauch, and have never been reset, the rovers' SCLK times have drifted significantly. The SCLK times on Spirit and Opportunity differed from each other, and they differed from our clocks here on Earth. We define the rover's clock drift as the difference between its Mission Clock's SCLK time and the J2000.0 Ephemeris Time, so positive clock drift indicates that the Mission Clock is running faster than our clocks on Earth. The amount of clock drift is determined by the temperature of the clock, as described further in Section IV.

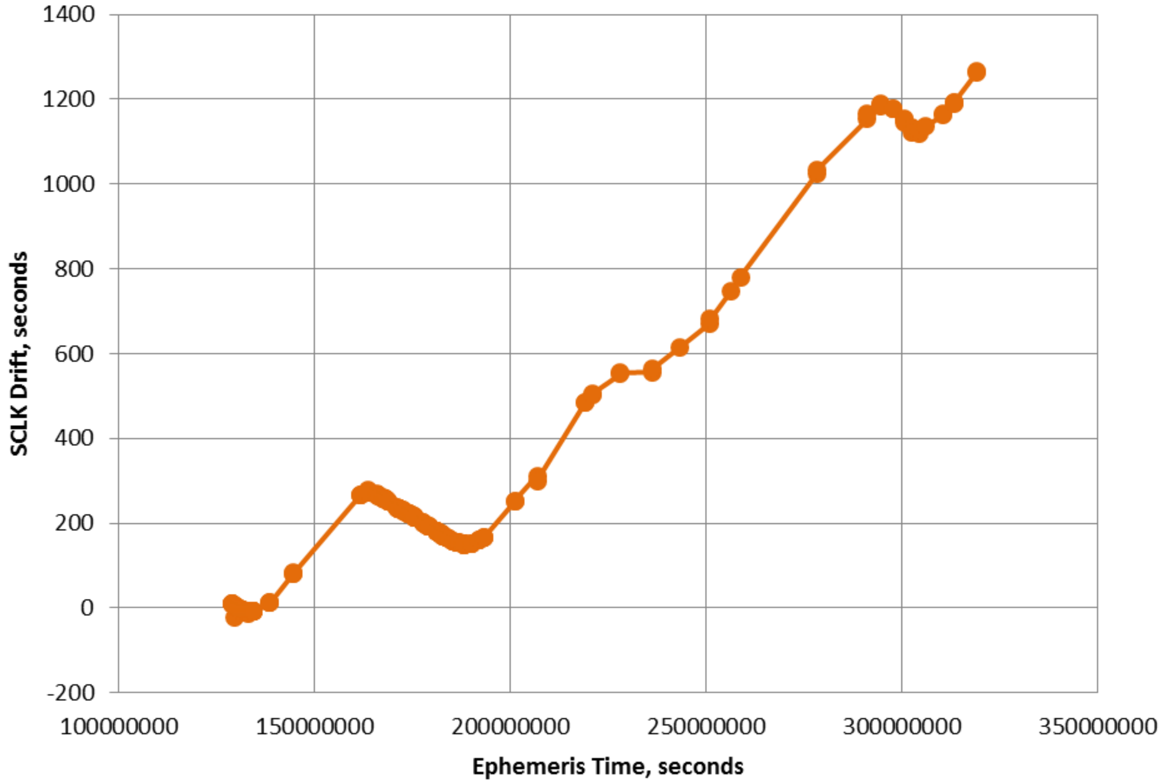

Figure 2. Spirit's clock drift from landing to the end of mission. The circular markers represent instances when the clock drift was measured. The lines connecting the markers are an interpolation of the drift.

Figure 2 shows the progression of Spirit's clock drift from the time of landing until UTC February 11, 2010, which was the last time it was measured before the end of the Spirit mission. When Spirit landed on Mars, its clock drift was about 10 seconds. The drift rose to 21 minutes and 5 seconds as of the last measurement.

Figure 3 shows the progression of Opportunity's clock drift from the time of landing until UTC July 9, 2012. When Opportunity landed on Mars, its clock drift was about 8 seconds. As of July 9, 2012, the drift was 19 minutes and 19 seconds.

The primary method we use to measure the rover's clock drift are Time Correlation Packets. These are telemetry from the rover containing information which can be used to determine the SCLK time at which the packet was sent. Using knowledge of the one-way light time and the time the packet was received, we can determine the Ephemeris Time at which the packet was sent, and thus we can determine the clock drift. We can only use this technique when we communicate directly to Earth, which is not commonly done. More typically, the rover communicates to one of the orbiting spacecraft around Mars, and that satellite later relays the data to Earth, which does not allow for determining the clock drift from Time Correlation Packets.

There are a few other less accurate methods we sometimes use to estimate the clock drift. One technique involves comparing the timing of the start and end of communications windows to when that event should have happened. Another technique involves examining the magnitude of tilt-based attitude estimation errors, which are primarily caused by the clock drift.

III. Problems Caused by MER Clock Drift

If the mission had lasted 90 sols, as planned, then the magnitude of the clock drift would not have been problematic. However, more than 3000 sols later, the magnitude of the clock drift causes several problems for the FSW and mission operations.

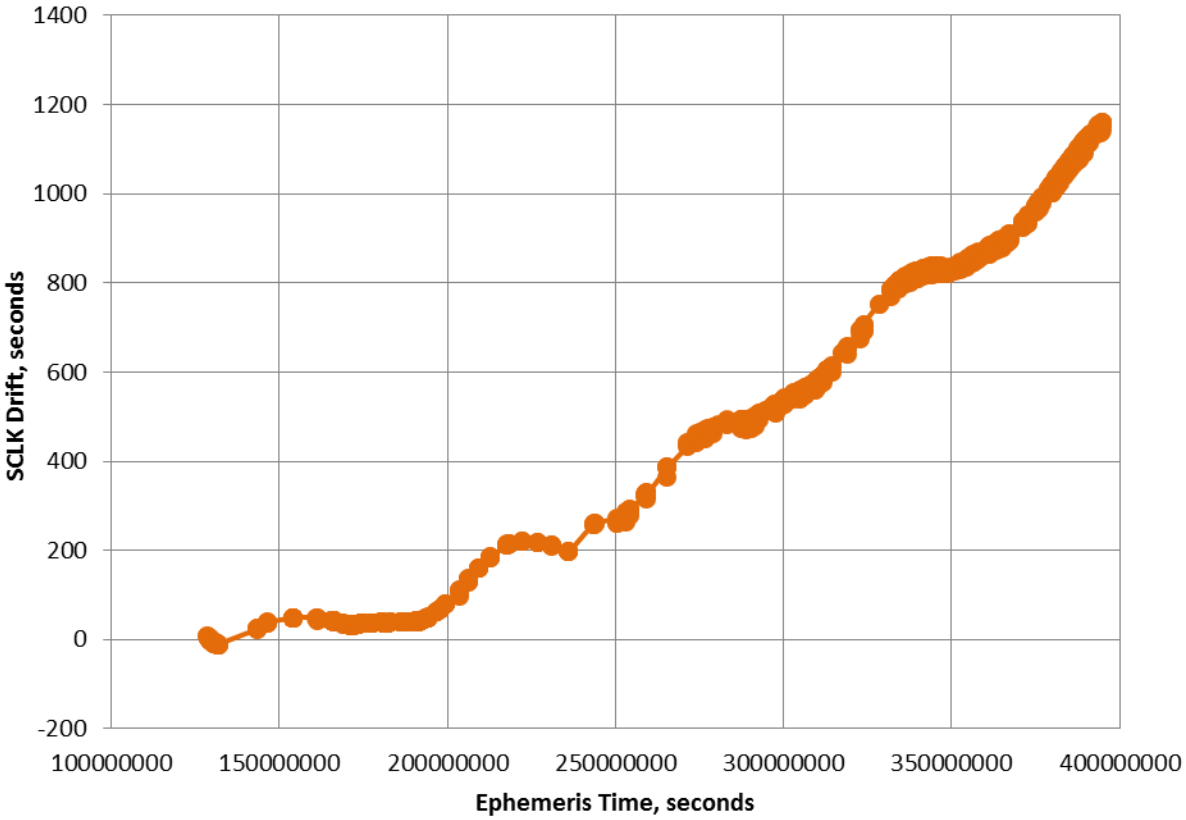

Figure 3. Opportunity's clock drift from landing until July 9, 2012. The circular markers represent instances when the clock drift was measured. The lines connecting the markers are an interpolation of the drift.

Since we can measure the clock drift, we can alleviate some of the problems it causes for mission operations. The ground tools use the measured relationship between SCLK time and Spacecraft Event Time (SCET), where SCET is the UTC time events happen on the spacecraft. This relationship is stored in a set of infrequently updated files, which we call Spacecraft Clock Coefficient Files, or SCLKvSCET files for short. The ground tools read in the latest SCLKvSCET file. Between updates, the relationship is linearly extrapolated from the latest data points. The rover, however, does not have access to these files, and the FSW is not designed to use this information. So, the FSW uses the uncorrected SCLK time when scheduling events in absolute time.

Some activities continue to be problematic, due to the clock drift, including attitude estimation, driving, robotic arm activity, imaging, power analysis, and telecom analysis.

A. Attitude Estimation

The rover's primary means of determining its attitude uses observations of the sun with one of its Pancam cameras[5] (see Fig. 1). These sun-based attitude updates find the sun's location in the sky relative to the rover and compare that observation with where the sun *should* be based on the rover's belief of the current time and its position on Mars. If the clock is incorrect, then the attitude estimate will be incorrect. The error in the attitude estimate due to the clock drift shows up as tilt error, meaning a difference between the estimated and real roll and pitch components of the attitude. On Opportunity, as of July 9, 2012, the tilt error caused by clock drift was about 4.8 degrees.

If the clock is running fast, as Spirit and Opportunity's are, the rover will believe it is tilted more to the West than it actually is. This is because the sun appears to be further behind in its path through the Martian sky than it should be at the believed time. The rover determines an attitude for itself which would put the sun at this location at the believed time. Likewise, if the clock was running slow, the rover would believe it was tilted more to the East than it actually was. The incorrect time knowledge has the same effect as if the rover had incorrect longitudinal knowledge of its position on Mars. A clock running fast affects the rover attitude estimate in the same way as a position estimate indicating the rover is further East on Mars than it really is. The incorrect time knowledge is the root cause of the attitude estimate error, which causes many of the other problems described in later sections.

As an example, assume a rover that is actually level with no tilt, facing due North, and a clock running 20 minutes fast. At the time the sun should have just risen over the horizon, let us say 7:30 LST for the sake of example, the rover believes it is 7:50 LST. So, the rover believes the sun should be about 5 degrees higher in the sky than the Eastern horizon. But, since the actual time is the time of sunrise, the sun is observed on the Eastern horizon, which is just over the right side of the rover's deck. The rover "knows" (incorrectly, due to the clock drift) that the sun is 5 degrees higher over the right side of its deck than what it observes. To account for the difference, the rover adjusts its attitude estimate to have a 5 degree Westerly tilt. This reconciles, for the rover, why it observed the sun just over the right side of the deck, rather than 5 degrees higher.

The tilt error could be removed using the rover's accelerometers, which the FSW can use to update only the roll and pitch information of the attitude estimate. So, the rover could be instructed to determine its attitude using sun observations and then correct the resulting tilt error using the accelerometers. This would produce an accurate attitude estimate. However, as explained further below, the current correct HGA pointing is due to the cancellation of the errors in the time estimate and the attitude estimate. If only one of these errors is eliminated, then the other will negatively affect the HGA pointing. Since MER operations relies on the HGA for commanding the rover, we must preserve accurate HGA pointing, and thus using the accelerometers to correct the tilt error is not a viable strategy.

The FSW has multiple methods to acquire and update the rover's attitude estimate[6]. The clock drift affects one of these techniques for updating the rover's attitude. The method we typically use, dubbed QuickFineAttitude, relies on the fact that the attitude estimate error is small enough that the sun will be in the image when the camera is commanded to point at the sun. So long as the attitude error is not too great and the clock has not drifted too much since the last attitude update, this technique works.

Due to the clock drift, the other method for commanding an attitude update, called GetFineAttitude[6], can no longer be used in the same manner as when the rovers landed. This technique first uses the accelerometers to determine the rover's tilt. This step is independent of time and not affected by the clock drift. Then, using the observed tilt and the erroneous time from the Spacecraft Clock, the FSW determines the elevation above the horizon which it expects to find the sun. This derived elevation will be incorrect, due to the error in the time that was used as an input. The rover searches for the sun with its cameras along a swath of the sky at the derived elevation. Since the cameras have a 16 degree field of view, the sun will be entirely in one of the images if the error in the derived elevation is less than about 7 degrees. However, if the observed elevation of the sun is too far off from the expected elevation, one of the sanity checks for this technique will fail. Then the observed sun is rejected and the attitude update fails. To continue using the GetFineAttitude method, the sanity check must be disabled or loosened. If the clock drift is allowed to grow to 28 minutes Mars Local Solar Time (about 28.77 minutes Earth time), the derived sun elevation error will be greater than about 7 degrees. Then the sun will not be fully present in the images that are used to search for the sun, and the attitude update will also fail.

B. HGA Pointing and Telecom Analysis

The clock drift can cause problems when the Earth is close to the horizon. The High-Gain Antenna (HGA) FSW[7] uses the SCLK time to determine when the Earth is above the horizon and to predict when the Earth will rise or set over the horizon. The HGA FSW will prevent an HGA communication pass from starting if it believes the Earth is below the horizon, and it will stop an ongoing HGA communication pass if it believes the Earth has set over the horizon. Since the SCLK time is incorrect, the HGA FSW may unnecessarily stop or prevent an HGA communication pass.

Additionally, the rover's attitude estimate and the time of the HGA passes are used as inputs to ground tools to predict when the Earth will be below the horizon or the deck. Whereas the rover's attitude estimate input comes from telemetry, and has error due to the clock drift, initially, the time used as an input to the ground tool had no error. Therefore, the output generated by the ground tools could have errors when used to predict when the Earth will set over the rover deck. Recently, we updated our telecom/ACS ground tools to use the latest estimates of the clock drift SCLKvSCET file, which eliminated this problem.

Possibly counter-intuitively, the attitude estimate error caused by the clock drift does not cause problems with pointing the High Gain Antenna (HGA) at the Earth. While the FSW uses the rover's attitude estimate as one input when computing where to point the antenna, it also uses the time to determine the Earth's position[7]. The estimate of the Earth's position in Local-Level Frame (LLF), a frame affixed to the Martian surface at the current location, is inaccurate due to the clock drift. However, the error is the same as the error in the rover's knowledge of the sun's position. Any additional lock drift which occurs between the determination of the attitude and the pointing of the

HGA will still contribute to pointing error, but attitude estimate updates are done frequently enough to keep this error small. Therefore, the attitude error caused by the incorrect sun position knowledge, which in turn was caused by the incorrect time knowledge, allows the rover to point the HGA at the Earth, even though it has incorrect Earth position knowledge.

C. Driving and Instrument Deployment Device Operation

The FSW used for driving has various safety checks to stop the drive if conditions become dangerous[8]. These checks run each time the rover drives an arc, and some of them monitor the rover's tilt. Limits on the tilt can be set in the sequences of commands that the Rover Planners, who plan the drives and motion of the robotic arm, create to drive the vehicle[9,10]. In addition to monitoring the tilt, the rover can be told to monitor sub-components of the tilt, such as roll, pitch, or northerly tilt. The FSW that performs these checks monitors the tilt derived from the accelerometers, which has no error due to the clock drift. The Rover Planners, however, determine the appropriate tilt limits for each drive by looking at images of the terrain they intend to drive through, with maps of the predicted tilt overlayed onto the images, as well as digital elevation models (DEMs) generated from the images. Those overlays and DEMS are generated from the rover attitude estimate, with the error due to clock drift, embedded in their metadata.

The inconsistency between the tilt observation used by the safety checks and the tilt information used to set the check limit forces the Rover Planners to add extra margin into the tilt limits they set. This extra margin must have a magnitude equal to the tilt error, because a tight tilt limit might cause a false positive, stopping a drive that was safe. Recently, the ground operators who produce the tilt maps and DEMS have begun producing additional image products, generated using the tilt information from an average of the accelerometer readings taken after the latest drive, but before the images were taken, eliminating the tilt error from the overlay.

Whereas the tilt error is consistent in a particular compass direction, the error in the roll and pitch are based on the current rover yaw. The Rover Planners typically cannot use their knowledge of the compass direction of the tilt error to apply different amounts of margin for roll and pitch limit. If the yaw changes during the drive, the amount of margin required for the roll and pitch components will change as well. This can cause confusion, which has been responsible for a prematurely halted drive on at least one occasion.

Additionally, the FSW's rover position estimate[6] suffers from the incorrect attitude estimate resulting from the clock drift. The position estimation is based on wheel odometry, with the direction of each position update based on the latest attitude estimate. Since the tilt estimate is incorrect, the position updates are added in an incorrect direction. Effectively, the rover believes it is driving on a greater or lesser slope than it really is. This problem affects the accuracy of all three components of the rover's position estimate, but it stands out mostly in the elevation (Z-component).

The Rover Planners also generate sequences of commands to control the Instrument Deployment Device (IDD), which is a robotic arm with various science instruments at its end effector[11]. The IDD FSW also has safety checks which run continuously during motion of the arm, and one of these checks monitors the rover tilt. This check, however, monitors the tilt derived from the FSW's attitude estimate rather than accelerometer readings. This difference in which tilt is monitored, between the IDD and driving limits, is one more thing for the Rover Planners to keep in mind when sequencing commands for driving or operating the IDD.

Furthermore, when commanded to move the IDD, the IDD FSW performs droop compensation when the end effector is commanded to a location in Cartesian space. The FSW adjusts the final joint positions of the arm from those which were initially computed using the inverse kinematics, to account for the pull of gravity on the arm. The IDD FSW uses the attitude estimate, which has the tilt error due to clock drift, to determine the direction of gravity, so there is a slight error in the droop compensation due to the clock drift.

D. Imaging

The camera pointing for images taken with the rover's cameras can be specified in several different frames, some fixed to the rover and some fixed to the surface of Mars, and in either three-dimensional Cartesian coordinates or azimuth/elevation angles[12]. The clock drift does not affect the pointing for images specified in rover-fixed or celestial coordinate frames. However, the FSW's attitude estimate is used to determine the transformation between the camera-fixed frame and those frames fixed to the surface of Mars. Therefore, the attitude error that results from clock drift affects the pointing of images in Mars-fixed frames.

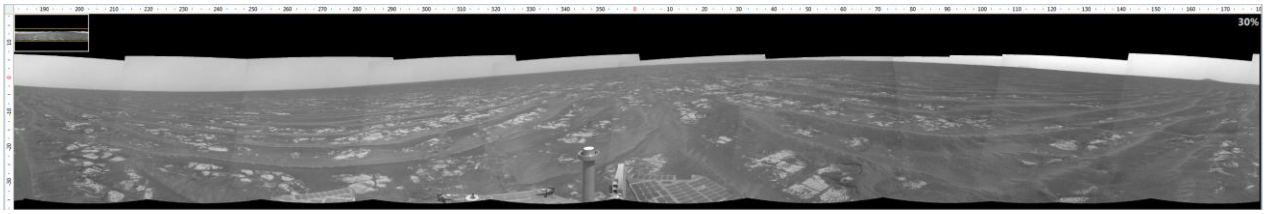

Figure 4. A 360 degree panoramic image mosaic taken by Opportunity on sol 2646. The images were pointed in a Mars-fixed frame. The actual terrain is almost entirely flat, with the exception of the small ripples in the foreground and a hill in the distance, which appears as a small bump on the horizon on the rightmost side of the image. The tilt error causes the cameras to point higher than intended in one direction and lower in the other, giving the horizon a sinusoidal appearance.

Fortunately, because of the way the ground operation team targets most of their images, this pointing error is only a problem some of the time. Normally, after driving to a new location, the operations team will take a panorama of the terrain with the Navcams (see Fig. 1), which have a wider field-of-view than the Pancams, to use for further pointing of the cameras. These panoramas, which are pointed in a Mars-fixed frame, do suffer from pointing error due to the tilt error. This can be seen most notably in full 360 degree panoramas, such as the one in Figure 4. The tilt error causes the cameras to point higher than intended on one side of the rover and lower on the other, causing the horizon to appear sinusoidal. After the images in the panorama are downlinked to the ground, the operators will target further Pancam images on particular locations of interest seen in the panorama. The ground tools used to view the images use the attitude estimate embedded in the image metadata, which has the tilt error included, to display the Mars-frame coordinated of each pixel in the image. The ground operators choose their targets, determining the Cartesian coordinates in the Mars-fixed frame using the erroneous transformation that the ground tools display. However, since the FSW that is used for pointing the cameras uses this same erroneous transformation, since it has the same tilt error, so the cameras end up properly pointed at the target.

E. Power Analysis

Engineers on our Power Subsystem team produce predictions of the rover's available energy each sol. These predictions are based on a model that takes several pieces of information as input, one of which is the rover attitude. The attitude is important, since the rover is solar powered, and the angle of the sun on the solar panels affects the energy produced. The Power Engineers use the telemetered attitude estimate as the input for their model. Since this attitude estimate has tilt error due to the clock drift, the output of the model will be incorrect.

The Power Engineers accommodate the error in the tilt estimate by adjusting another input to the model, the rover's longitude on Mars. As mentioned in Section III (A), the error in the attitude due to clock drift appears the same as error in its knowledge of its longitude. The Power Engineers regularly compare the performance of their model with the actual power telemetry, and they update the rover longitude used in their model to accommodate the attitude estimate error.

IV. Relationship Between Clock Drift and Temperature

The Mission Clock in the MER rovers uses a custom hybrid crystal oscillator from Q-Tech Corporation. The clock drift rate is related to the temperature of the clock[3]. The general trend of the drift rate expected over the temperature range that the rovers experience shows that the clock should run faster at the lower end of the temperature range and slower at the higher end.

Figures 5 and 6 show plots of clock drift versus temperature, based on telemetry from the rovers. The drift is calculated for each time period, of at least one sol, between receipt of time packets. Then, the drift was normalized to a 10 sol period, to facilitate comparison, since we do not receive time packets on a regularly occurring schedule. The normalized drift is plotted against the mean temperature, recorded by the closest temperature sensor to the mission clock, over the time period. The FSW can only record the temperature from this sensor while the rover is awake (while the FSW is running), so at night, when the rover is asleep, no data is collected. Since the temperatures are coldest during the Martian nighttime, the average of the measured temperatures will be higher than the true average temperature for the sol. Further error is introduced because the rover does not always wake up and sleep at

---

[3] Q-Tech provided us with a polynomial relating the error to temperature, as well as error bounds, but we cannot share that information.

the same times each sol. Over the most recent Martian winter, when the temperatures were coldest, Opportunity has spent much of each sol asleep, even during much of the daylight hours. This contributes to extra noise and a skew in the plot for Opportunity at lower temperatures, where high drift rates appear at temperatures higher than they would if Opportunity had been awake most of the daylight hours.

Although the plots of the clock drift versus the mean temperature are very noisy, for the typical operating temperatures, the drift exhibits a definite trend for both rovers. The drift has been more positive, indicating the clock is running faster, at lower temperatures, and more negative at higher temperatures. The plots also seems to indicate that Spirit's clock required a slightly higher temperature to run at the desired rate, with no drift. Whereas this may be the case, it may also be due to the fact that the temperature readings are only collected while the rover is awake, and the two rovers had different sleep/wake periods.

Figures 7 and 8 show Spirit and Opportunity's cumulative clock drift over the mission and every measured temperature from the temperature sensor nearest the Mission Clock since the rovers landed on Mars. Note that the temperatures are only recorded while the rover is awake. These graphs demonstrate that the clocks drift in the positive direction the most during the coldest periods and the least during the warmest, with the direction of the drift reversing during the warmest periods. The sudden dip in temperatures around the second summer solstice after landing was due to a planet-wide Martian dust storm.

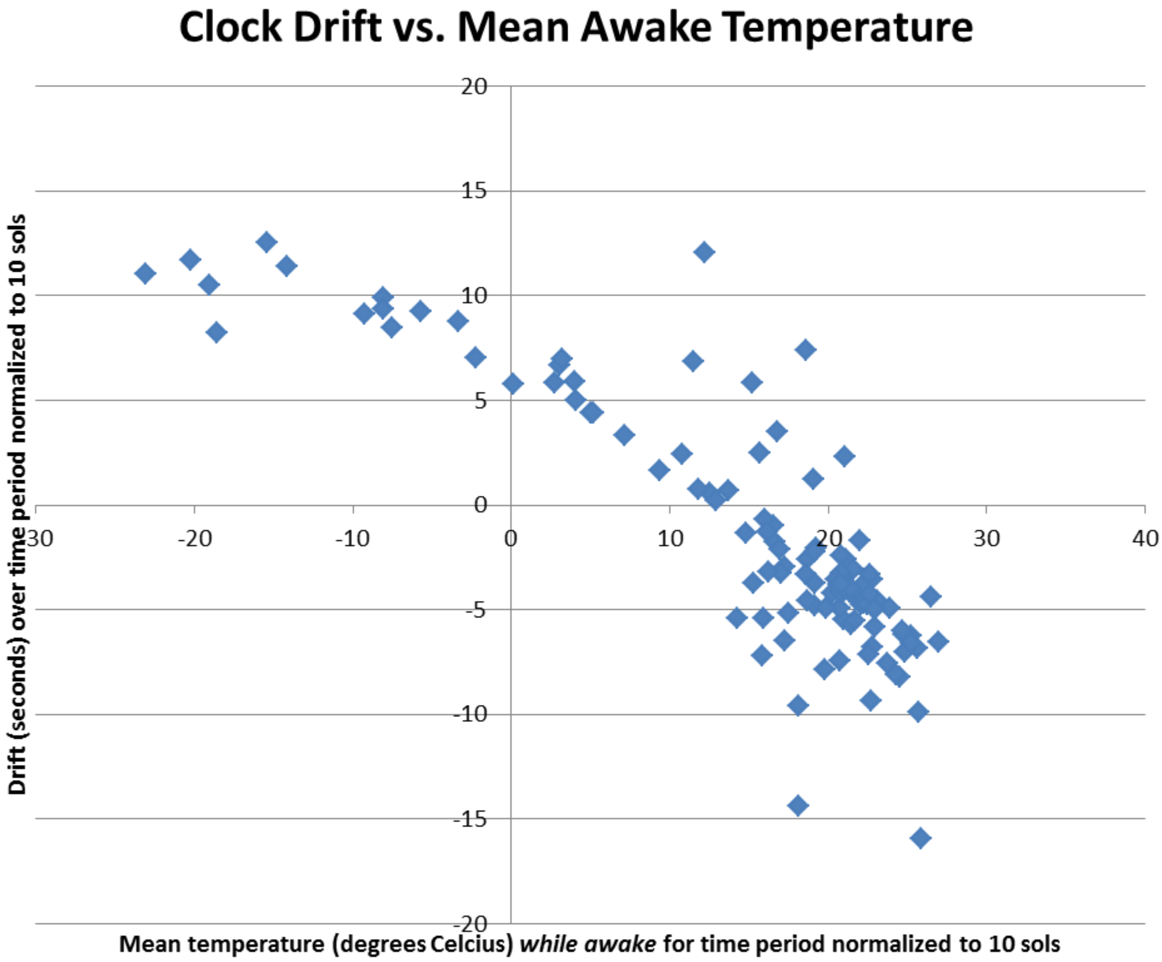

Figure 5. Spirit's Clock Drift vs. Mean Awake Temperature. The clock drift is normalized for a 10 sol time period. The temperatures are from readings only when the rover was awake.

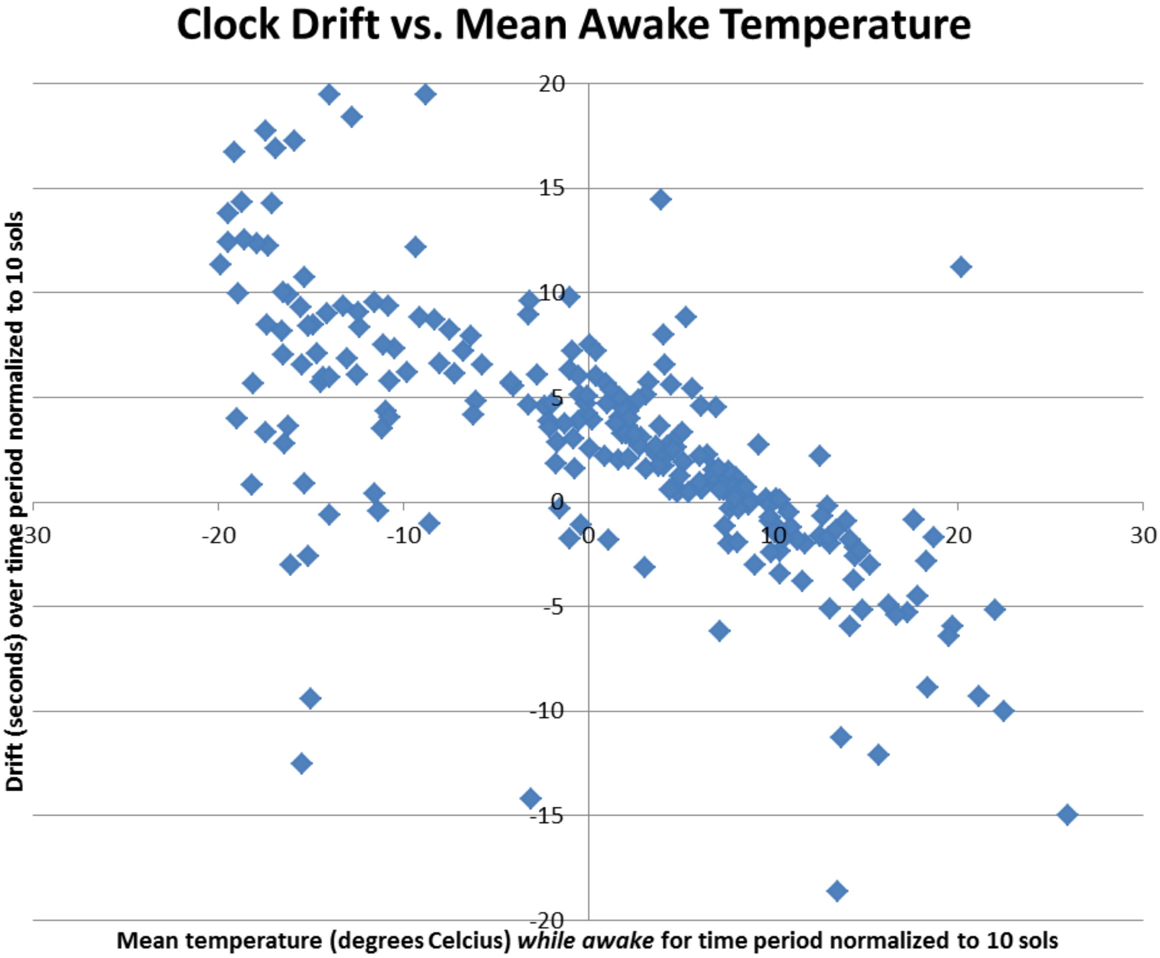

Figure 6. Opportunity's Clock Drift vs. Mean Awake Temperature. The clock drift is normalized for a 10 sol time period. The temperatures are from readings only when the rover was awake.

V. Techniques for Reducing the Clock Error

We have considered three methods to reduce the clock error and drift. Two techniques involve setting the clock to the correct time, either all at once or in small increments. The third technique involves controlling the amount of clock drift by controlling the temperature of the clock.

The FSW has five commands that the operations team can use to set the clocks on the rover. Three of those commands set the Spacecraft Clock and the Mission Clock simultaneously. These three commands differ in terms of how they are applied. One command sets the time to a value relative to the current Spacecraft Clock value. This time adjustment takes place at the next one second boundary after the command is processed. The other two commands that set both the Spacecraft Clock and Mission Clock set the new time value to an absolute time. One of them takes effect at the next eigth of a second boundary, and the other takes effect at a Spacecraft Clock time specified in the command arguments.

A fourth command synchronizes the Mission Clock and Spacecraft Clock, setting the Mission Clock time value to be the same as the current Spacecraft Time value. This takes place at the next one second boundary after the command is processed.

The fifth command sets only the Mission Clock, without changing the Spacecraft Clock. The Spacecraft Clock is seeded from the Mission Clock each time the FSW is initialized, and after that, only the Spacecraft Clock is used by the FSW to determine time. So, the effects of this command will only be seen by the rest of the FSW after the next cold boot.

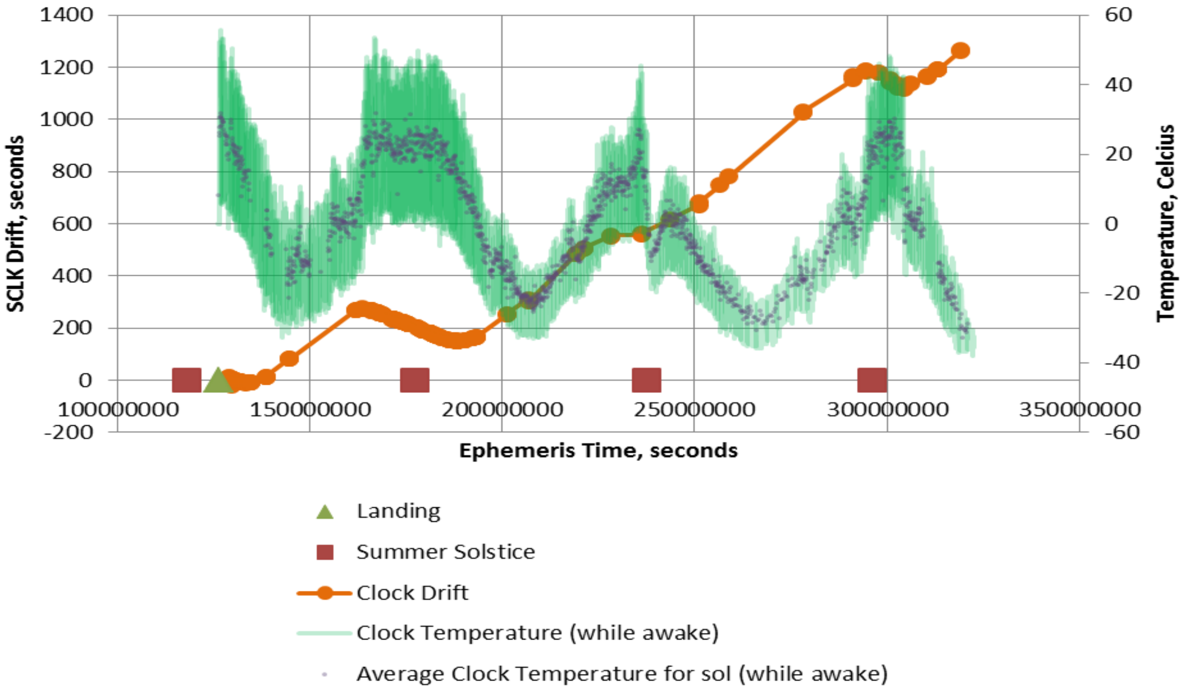

Figure 7. Spirit's cumulative clock drift and the temperature of the mission clock over the mission. The temperatures are from the sensor closest to the Mission Clock, and are recorded only while the rover is awake.

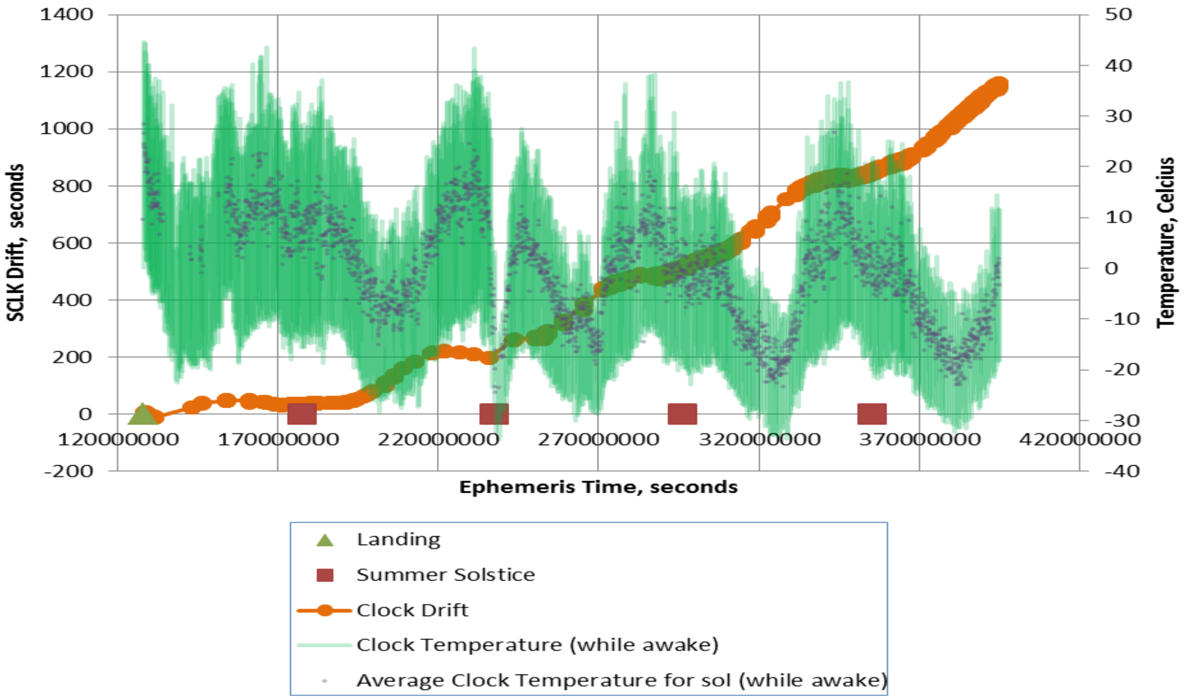

Figure 8. Opportunity's cumulative clock drift and the temperature of the mission clock over the mission. The temperatures are from the sensor closest to the Mission Clock, and are recorded only while the rover is awake.

F. Resetting the Clock

The commands to set the clocks could be used to set the the desired time, eliminating the clock drift all at once. However, since the rover's time estimate is ahead of our time estimate on the ground, resetting the rover's clocks would involve jumping the clock to a time that the FSW believes to be in the past. Setting the clock backwards in one large jump carries risk and has some undesireable features.

If the FSW sees a backwards jump in time, it could cause an unintended reset of the FSW. Portions of the FSW assume that time only moves in one direction. Whereas jumping the clock backwards does not actually change the direction of time flow, to the FSW it would appear to do so. Depending on the actual timing of events around the clock reset, the FSW might detect the change as an anomaly and reset itself. We would like to avoid unintended and nondeterministic FSW resets.

Another problem is related to the timestamps on telemetry data. The telemetry from the rover is timestamped with the Spacecraft Clock time. The ground operators use this timestamp to sort the telemetry chronologically and determine the order of events on the rover. If the rover's clock were set backwards, there would be a confusing period when the telemetry from different real times had overlapping timestamps. This period of time would have a duration equal to the amount the clock was set backwards.

There is a mitigation strategy for both of the problems caused by the backwards direction of the necessary clock adjustment. The command which only sets the Mission Clock could be used, setting that clock to the correct time, but leaving the Spacecraft Clock untouched. This would have no effect until the rover wakes up from its next sleep. So long as the duration of the sleep is longer than the magnitude of the backwards jump of the Mission Clock, then there would be no chance of an unintended FSW reset and no telemetry mess.

There are some further problems related to changing the clock time by a large amount, which are present regardless of the direction of the adjustment. First, the rover's attitude estimate would need to be updated after the time correction, but before the next HGA communication pass. The current correct pointing of the HGA depends on the fact that the determination of the rover's attitude and the prediction of the Earth's position both have the same time error. After setting the clock, the attitude estimate will have been determined with one time error amount, and the Earth's position at the time of an HGA pass will be determined with a new, presumably smaller, time error. The larger the *difference* in the time error, the larger the error in pointing the HGA will be. If the HGA pointing error is large enough, which it would be if the clock were changed by the current time error amount, then the ground operations team would not be able to command the rover through the HGA. Although there are other communication paths for commanding the rover, they are much less desireable. HGA commanding is the nominal and preferred method.

Due to the large change in time, the attitude estimate update would need to be done using the GetFineAttitude technique, rather than the QuickFineAttitude technique we normally use. This is because it is unlikely the sun would be close enough to the center of the initial image of the QuickFineAttitude such that it would remain in the image throughout the ten minute Sungaze[6]. Since the prior attitude estimate was determined using the time before the update, and the expected Sun position is determined using the time after the update, the image pointing would be off and the sun would not be in the center of the image. With the current clock error, it is possible that the sun could move off the side of the image while the rover was watching it during the QuickFineAttitude. If the attitude estimate update after the clock update failed, it might be several sols before another one could be re-commanded, due to the difficulty in commanding through the HGA with the old attitude estimate after the time correction.

After a large update to the rover's clock, the ground operations team would need to update the SCLKvSCET file. Without this update the ground tools could not translate the SCLK times for events like communication passes to the correct SCET time. The proper way to do this update involves sending Timing Packets from the rover, which require a successful HGA communication pass. So, this would need to be done after successfully updating the rover's attitude.

There are a few other "bookkeeping" issues that would need to be dealt with. All of the communication pass windows that had previously been loaded onto the vehicle would need to have their start times updated. The communication window times would need to change to reflect the new SCLK time representing the time the Deep Space Network (DSN) would be listening for the communication. This could be done only after the SCLKvSCET file update. Additionally, the longitude input to the power analysis ground software would need to be changed to reflect reality. Finally, the rover has a notion of the local time during the day, which is computed from the SCLK time and a parameterized reference time, which represents midnight on the first sol. The reference time has been updated several times during the mission to deal with the clock drift. It would need to be updated again after the clock error was reduced.

G. Incrementally Jumping the Clock

Several small incremental adjustments to the rover's clock can be used to mitigate the problems associated with a single large adjustment. To illustrate why this is true, an extreme case will be considered. If the clock is currently drifting forward a certain amount each sol, then it is possible to adjust the clock backward each sol by an amount smaller than the forward drift. Looking at the drift for a single sol, it would appear to drift forward a certain amount and then jump backwards part of that amount. Considered from a per-sol perspective, however, this would have the appearance of reducing the forward drift, but not eliminating it. If the adjustments were made only to the Mission Clock, as described in the previous section, then there is no risk of unintentional FSW reset. Furthermore, the ground operations team already has the experience and tools in place to deal with the slowly fluctuating rate of the clock drift.

By increasing the magnitude of the clock adjustment slightly, the drift can be effectively cancelled or reversed, when considered from a per-sol perspective. So long as the effective rate of the backwards drift is no greater than the forwards drift the clock has experienced, then the operations team can deal with that drift with no change to its existing procedures. Attitude updates, updates to the files relating SCLK and SCET times, updates to communications windows, and updates to the rover's notion of local time could continue to happen at their current frequency, without risk of losing contact with the rover.

The incremental clock adjustments could happen each sol, or slightly larger adjustments could happen every several sols. Adjustments of one minute would only affect the attitude estimate by approximately 0.24 degrees, an amount that would not cause serious HGA pointing problems, assuming the starting attitude estimate was reasonably accurate and the attitude estimate was updated before making the next adjustment. Furthermore, the effects of one minute adjustments could be tolerated from the perspective of communication window times. From a spacecraft operations perspective, however, infrequent updates would be favorable, due to workload and complexity considerations. We are currently in the process of determining the best magnitude and frequency for the adjustments.

H. Reducing the Clock Drift by Controlling the Clock Temperature

Since the rate of the clock drift is dependent on the temperature of the clock, the drift rate can be controlled if the clock temperature can be controlled. As Figures 7 and 8 show, when the average Mission Clock temperature is warm, the drift rate levels off. For particularly warm temperatures, which are still within the safe operational range for the hardware, the clock drifts in the other direction.

The Mission Clock is located in the Rover Electronics Module (REM). There are no heaters near the Mission Clock nor in the REM which can be turned on or off by ground commands. The temperature in the REM near the Mission Clock tends to be warmest when it is summertime at the rover's location. This is due to a combination of higher external temperatures and increased rover activity. Since the rover is solar powered, it can be more active when there is more sunlight during the summer, and the increased activity level keeps the REM warmer.

Although ground operators cannot turn on heaters to warm the Mission Clock, they do have some control over the activities the rover engages in. Certain activities are more effective at heating up the REM. The two most effective means of heating the REM are keeping the Inertial Measurement Unit (IMU) powered on, such as when the rover is driving, and conducting Ultra-high Frequency (UHF) communication passes. Due to power considerations, the ground operators power on the IMU only when it is needed, and it cannot be left powered on overnight.

Ideally, the rover would drive during the day, and then have both an afternoon and early-morning UHF communication pass. This is occasionally done during the Martian summertime. Unfortunately, during the Martian winter, when the Mission Clock needs heating the most, there typically is not enough power for either driving or a morning UHF pass. During the Martian summer, the rover often drives most of the day, and thus has the IMU powered on already, although there are some sols when the rover is not driving and the IMU could be powered on for a portion of the sol. Furhtermore, morning UHF communication passes are already used at almost every occasion when the power levels permit. So, there would be only a small amount of extra benefit that could be gained by attempting to control the clock temperature by controlling the activity level.

There is some risk associated with trying to increase the clock temperature. The hardware in the REM can be damaged if the temperature is too high. Although the FSW's fault protection behavior will shut the rover down if it gets too hot, this will also deactivate any sequences of commands the rover is running. We prefer not to surrender control of the rover to the autonomous fault protection behavior. So, it would require careful planning to keep the REM temperature warm, but not too warm. Due to the limited benefit, we have not currently attempted to do this.

VI. Current State and Plans

To date, no attempts at reducing the clock drift have been conducted. A few experiments were conducted to determine how effectively the ground operations team can raise the clock temperature by powering on the IMU, but there has been no actual effort to reduce the clock drift.

The current plan is to correct the drift of the Mission Clock on the Opportunity rover by conducting a series of small incremental clock adjustments. This effort is still in the planning stages. The size and frequency of the updates has not been determined. If the plan is implemented, it should be effective at eliminating the clock error, but the clock would continue to drift, especially during the colder times of the Martian year. Therefore, occasional additional incremental clock adjustments would still be needed.

Acknowledgments

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. It was funded by the Mars Exploration Rover (MER) Project.

The authors would like to thank Ryan Mukai and Alfonso Herrerra for obtaining and providing information about the Mission Clock and IMU on the MER rovers. We would also like to thank Q-Tech Corporation for providing details about the hybrid crystal oscillator used in the Mission Clock on the MER rovers.

References

[1] Arvidson, R. E., Bell III, J. F., Bellutta, P., Cabrol, N.A., Catalano, J. G., Crumpler, L., Des Marais, D. J., Estlin, T., Farrand, W., Gellert, R., Grant, J. A., Greenberger, R., Guinness, E. A., Herkenhoff, K. E., Herman, J. A., Iagnemma, K., Johnson, J. R., Klingelhöfer, G., Li, R., Lichtenberg, K. A., Maxwell, S., Ming, D. W., Morris, R.V., Rice, M., Ruff, S., Shaw, A., Siebach, K., de Souza, P., Stroupe, A., Squyres, S. W., Sullivan, R. J., Talley, K., Townsend, J., Wang, A., Wright, J., Yen, A., "Spirit Mars Rover Mission: Overview and Selected Results from the Northern Home Plate Winter Haven to the Side of Scamander Crater", *Journal of Geophysical Research*, Vol. 115, 2010, E00F03,doi:10.1029/2010JE003633.

[2] Arvidson, R. E., Ashley, J. W., Bell III, J. F., Chojnacki, M., Cohen, J., Economou, T. E., Farrand, W. H., Fergason, R., Fleischer, I., Geisler, P., Gellert, R., Golombek, M. P., Grotzinger, J. P., Guinness, E. A., Haberle, R. M., Herkenhoff, K.E., Hermann, J. A., Iagnemma, K. D., Jolliff, B. L., Johnson, J. R., Klingelhöfer, G., Knoll, A.H., Knudson, A. T., Li, R., McLennan, S. M., Mittlefehldt, D. W., Morris, R.V., Parker, T. J., Rice, M. S., Schröder, C., Soderblom, L. A., Squyres, S.W., Sullivan, R. J., Wolff, M. J., "Opportunity Mars Rover Mission: Overview and selected result from Purgatory Ripple to Traverses to Endeavour Crater", *Journal of Geophysical Research*,Vol. 116, 2011, E00F15, doi: 2010JE003746.

[3] Standish, E. M., "Time scales in the JPL and CfA ephemerides", *Astronomy and Astrophysics*, Vol. 336, 1998, pp. 381-384.

[4] Hughes, D. W., Yallop, B. D., and Hohenkerk, C. Y., "The Equation of Time", *Monthly Notices of the Royal Astronomical Society*, Vol. 238, 1989, pp. 1529-1535.

[5] Reeves, G. E., Snyder, J. F., "An Overview of the Mars Exploration Rovers' Flight Software", *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics*, 2005.

[6] Ali, K. S., Vanelli, C. A., Biesiadecki, J. J., Maimone, M. W., Cheng, Y., San Martin, A. M., Alexander, J. W., "Attitude and Position Estimation on the Mars Exploration Rovers", *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics*, 2005.

[7] Vanelli, C. A., Ali, K. S., "High Gain Antenna Pointing on the Mars Exploration Rovers", *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics*, 2005.

[8] Biesiadecki, J. J., Maimone, M. W., "The Mars Exploration Rover Surface Mobility Flight Software: Driving Ambition", *Proceedings of the IEEE Aerospace Conference*, 2006.

[9] Leger, C., Trebi-Ollennu, A., Wright, J., Maxwell, S., Bonitz, R., Biesiadecki, J., Hartman, F., Cooper, B., Baumgartner, E., Maimone, M., "Mars Exploration Rover Surface Operations: Driving Spirit at Gusev Crater", *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics*, 2005.

[10] Biesiadecki, J., Baumgartner, E., Bonitz, R., Cooper, B., Hartman, F., Leger, C., Maimone, M., Maxwell, S., Trebi-Ollennu, A., Tunstel, E., Wright, J., "Mars Exploration Rover Surface Operations: Driving Opportunity at Meridiani Planum", *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics*, 2005.

[11] Baumgartner, E., Bonitz, R., Melko, J., Shiraishi, L., Leger, P., "The Mars Exploration Rover Instrument Positioning System", *Proceedings of the IEEE Aerospace Conference*, 2005.

[12] Maki, J. N., Bell III, J. F., Herkenhoff, K. E., Squyres, S. W., Kiely, A., Klimesh, M., Schwochert, M., Litwin, T., Wilson, R., Johnson, A., Maimone, M., Baumgartner, E., Collins, A., Wadsworth, M., Elliot, S. T., Dingizian, A., Brown, D., Hagerott, E. C., Scherr, L., Deen, R., Alexander, D., Lorre, J., "The Mars Exploration Rover Engineering Cameras", *Journal of Geophysical Research*, Vol. 108, No. E12, 2003.