

MARS PATHFINDER MICROROVER

A SMALL, LOW-COST, LOW-POWER SPACECRAFT

Henry W. Stone

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109

ABSTRACT

On December 5, 1996, NASA will launch its first Discovery Class Mission, the Mars Pathfinder. On July 4, 1997 the 450 kg spacecraft will enter into the Martian atmosphere and descend directly to the surface using a viking style aeroshell, parachutes, RAD rockets, and airbags to slow its decent. Once on the surface the tetrahedrally-shaped lander will open like a flower and release a 10.5 kg rover previously to the inside surface of one of the lander's three deployable petals. The rover will then drive off the lander and begin to perform a wide range of scientific and technological experiments.

The significance of the rover is that although it is formally part of the mission's instrument payload, it is in reality a small spacecraft. It performs all the functions that a typical spacecraft performs including: navigation; command and data handling (command execution, data acquisition, telemetry packetization); power generation, distribution, and control; thermal control; telecommunications; and instrument, sensor, and actuator control. The rover is also quite unique in that its total cost of development as well as support for mission operations is \$25M [4]. This is more than an order of magnitude less than the cost of any previously flown interplanetary spacecraft. The rover is also being developed and built in only 3 years, which includes conceptual design on up to delivery and launch of the flight article.

The design of the rover has been influenced by a wide range of mission, environmental, and programmatic constraints. These have included severe limitations upon mass and volume; the need to survive the launch and landing loads; the desire to be as independent from the lander as possible; and the need to operate on the surface of Mars. The mass and volume constraints were particularly difficult to deal with and had an immediate impact upon how much power could be generated onboard. As a consequence power, or the lack thereof,

quickly became a major design constraint impacting nearly every aspect of the rover's design in particular the control system.

This paper provides an overview of the design of the Mars Pathfinder Microrover with an emphasis on how its integrated design enables it to perform all of the functions normally associated with entire spacecraft. The impact that mass, volume, and power constraints have had upon the design will also be discussed.

INTRODUCTION

The Mars Pathfinder Microrover is a 10.5 kg, 6 wheeled robotic vehicle which will be delivered to Mars aboard the Mars Pathfinder Spacecraft on July 4, 1997. The rover, which is formally part of the spacecraft's instrument payload, will perform scientific and technological experiments, and take images of the lander. Although the rover is formally an instrument, it in fact performs all of the functions that a typical spacecraft performs. It is self-contained, self-powered, and operates autonomously in a hazardous and unknown environment. Most importantly, the rover, using its onboard computer and navigation sensors, is capable of executing high level commands (e.g., Go-to-Waypoint X,Y) and navigating around terrain hazards without human intervention.

The flight rover depicted in Figure 1 is 68 cm long by 48 cm wide by 28 cm high when fully deployed. While encapsulated in the lander during launch, cruise, entry and decent, the rover is stowed in its pre-deployed configuration with a height of only 18 cm. The size of the rover, which was dictated by the size and configuration of the lander, posed a major challenge to its designers. In addition to the effort needed to design, test, and qualify light weight mechanisms, sensors, and structural components, the vehicle's limited size essentially dictated how much power it could generate. The resulting limitations on the amount of power which

could be generated had a fundamental impact upon the design of the vehicle's overall control system.

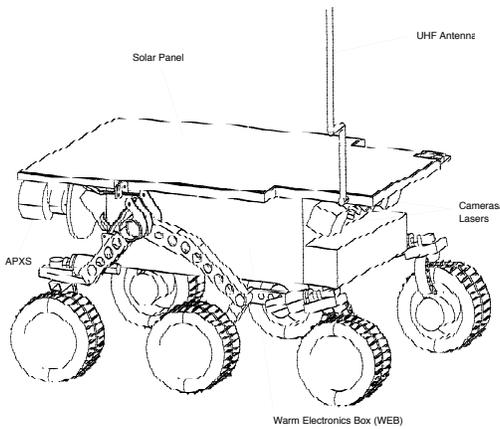


Figure 1: Mars Pathfinder Microver (MFEX)

MOBILITY

The basic mechanical design of the rover is based upon a 6 wheeled rocker-bogie suspension system [1]. This system enables the rover to traverse across the many different types of terrain conditions which are expected to exist at the landing site (e.g., rocks, sand, fine dust, and combinations thereof).

The rocker-bogie suspension system consists of a pair of rigid linkages on each side of the vehicle which are attached to one another by a passive rotary joint. The front and middle wheels are rigidly attached to each end of the forward linkage, the bogie. The rear wheel is attached to the rear end of the rear linkage, the rocker. The forward end of the rocker is attached to the middle of the bogie through a rotational joint. The rocker, and hence the three wheels, is attached to the vehicle's chassis by a second rotary joint on the rocker located forward of the rear wheel. The rockers on each side of the vehicle are connected to one another and the chassis by a member to create a differential between the two sides of the vehicle.

The rocker-bogie suspension system, in combination with the design of the wheels and wheel treads, provides the rover with an extremely high degree of mobility. As the vehicle drives, the wheels are free to move up and down independently of one another and to follow the contour of the terrain. The kinematics of the design are such that the weight of the vehicle remains nearly equal across all six wheels. Testing has verified that the vehicle can safely

climb obstacles 1 1/2 wheel diameters in height. It can also climb slopes to within 3 degrees of the angle of repose of the soil.

The four corner wheels are independently steerable and enable the vehicle to drive along an arc or to turn in place (i.e., about its center). The latter is particularly useful for navigating in highly constrained environments.

POWER GENERATION

The rover is powered by solar energy and/or energy stored within three LiSOCl₂ batteries [6]. The batteries are housed within the rover's Warm Electronics Box (WEB) and are capable of providing at least 150 Whr. The 0.22m² GaAs/Ge solar array [6], which covers the entire top of the rover, is capable of generating a maximum of 16 watts on Mars at the Ares Vallis landing site (15°N lat, 165°W long) assuming nominal atmospheric conditions. The current .vs. voltage (IV) curves of the panel for a Tau of 0.2 and a tilt of 0° is plotted in Figure 2 and the typical maximum daily power output of the panel is shown in Figure 3.

Solar energy is assumed to be the primary means for powering the vehicle during the mission. The batteries are intended to be a backup to the solar array should the array suffer severe damage during landing and/or failure for any reason. The energy storage capability of the rover's batteries is such that the rover can perform its entire 7 day primary mission without any solar power. The use of rechargeable batteries was precluded by the fact that rechargeable batteries have a relatively high mass to energy storage ratio.

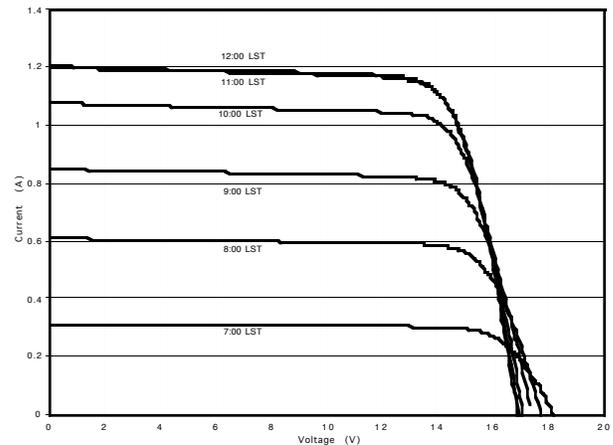


Figure 2: Solar Panel IV Curves (Tau = 0.2, Tilt = 0°)

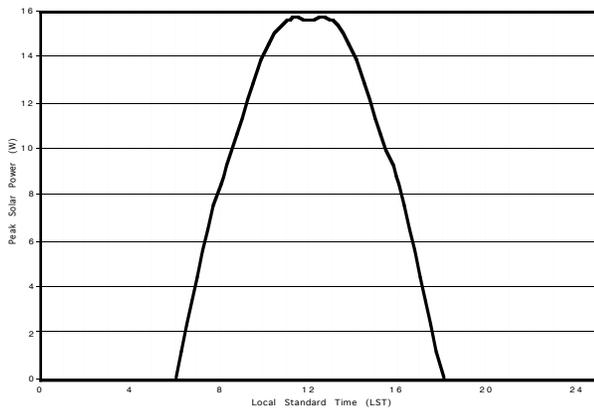


Figure 3: Maximum Solar Panel Power Output .vs. Time (LST)

Under nominal conditions, the batteries will be used to perform certain experiments at night, such as the gathering of long integration-time spectra with the Alpha Proton X-Ray Spectrometer (APXS) and routine health checks. They will also be used, if needed, to provide short bursts of energy if extreme terrain conditions are encountered. If the batteries were to fail or be completely discharged either prior to deployment or during surface operations, the rover will still be able to perform an entire mission, including a 30 day extended mission, using only solar power.

POWER MANAGEMENT

The rover contains a relatively complex power distribution and power management system as illustrated in the functional block diagram in Figure 4. This is required due to a variety of factors including: the need to provide numerous controlled supplies for the many onboard sensors (e.g., cameras, gyro, accelerometers), the need to maintain a high degree of reliability in the event of a device failure, the need to drive in unknown terrain for which it is impossible to accurately predict how much power will be needed, and most importantly, the need to operate using only solar power.

While a complete discussion of the details of the power distribution and control system is beyond the scope of this paper, several features are worth mentioning.

- The current limiter is a hardware circuit which is designed to insure that the Core Bus will not dip below 13.5 volts regardless of how loaded the motors become. Basically, it protects the CPU from

a brownout condition caused by excessive loading by the actuators.

- The battery bypass switch enables the vehicle to power the motors using battery power in the event of a solar panel failure.
- The “3 Battery String Switch” (3BSS) is a latching relay which is closed only after the vehicle has landed on Mars. This switch ensures that battery strings B and C remain offline during the 7 month cruise. This preserves their passivation layer, providing greater energy retention, and protects the batteries from being accidentally depleted during cruise.
- The latching relay on the input side of the 9.0 V converter enables the vehicle to power the APXS directly off the batteries and hence acquire spectra at night even though the vehicle’s CPU has been shutdown.
- The power monitor is a hardware circuit, which in combination with S/W allows the vehicle to gracefully shutdown at end of each day as the sun goes down.
- The load shedder is a hardware circuit which monitors the voltage of the Core Bus and immediately turns off all devices and supplied if the voltage drops below a preset value. It protects the CPU from unexpected power surges and informs the CPU of the event via an interrupt.
- The “Lander Controlled Power Switch” (LCPS) is a magnetically coupled to the lander so that the lander can periodically wakeup the rover for brief periods during cruise to perform simple health checks using battery power.
- The “Alarm-Clock Controlled Power Switch” (ACPS) is controlled by an independently powered alarm clock circuit. When the alarm clock goes off the rover wakes up. This switch is used to wakeup the vehicle if it didn’t wakeup when it was nominally expected to. For example, if the solar panel failed during the night this switch would cause the vehicle to wake up on battery power.

Although not explicitly shown, the rover’s power distribution and control system contains 58 independently controlled power switches. Managing the configuration and activation of these switches based upon the amount of measured power available from the solar panel and amount

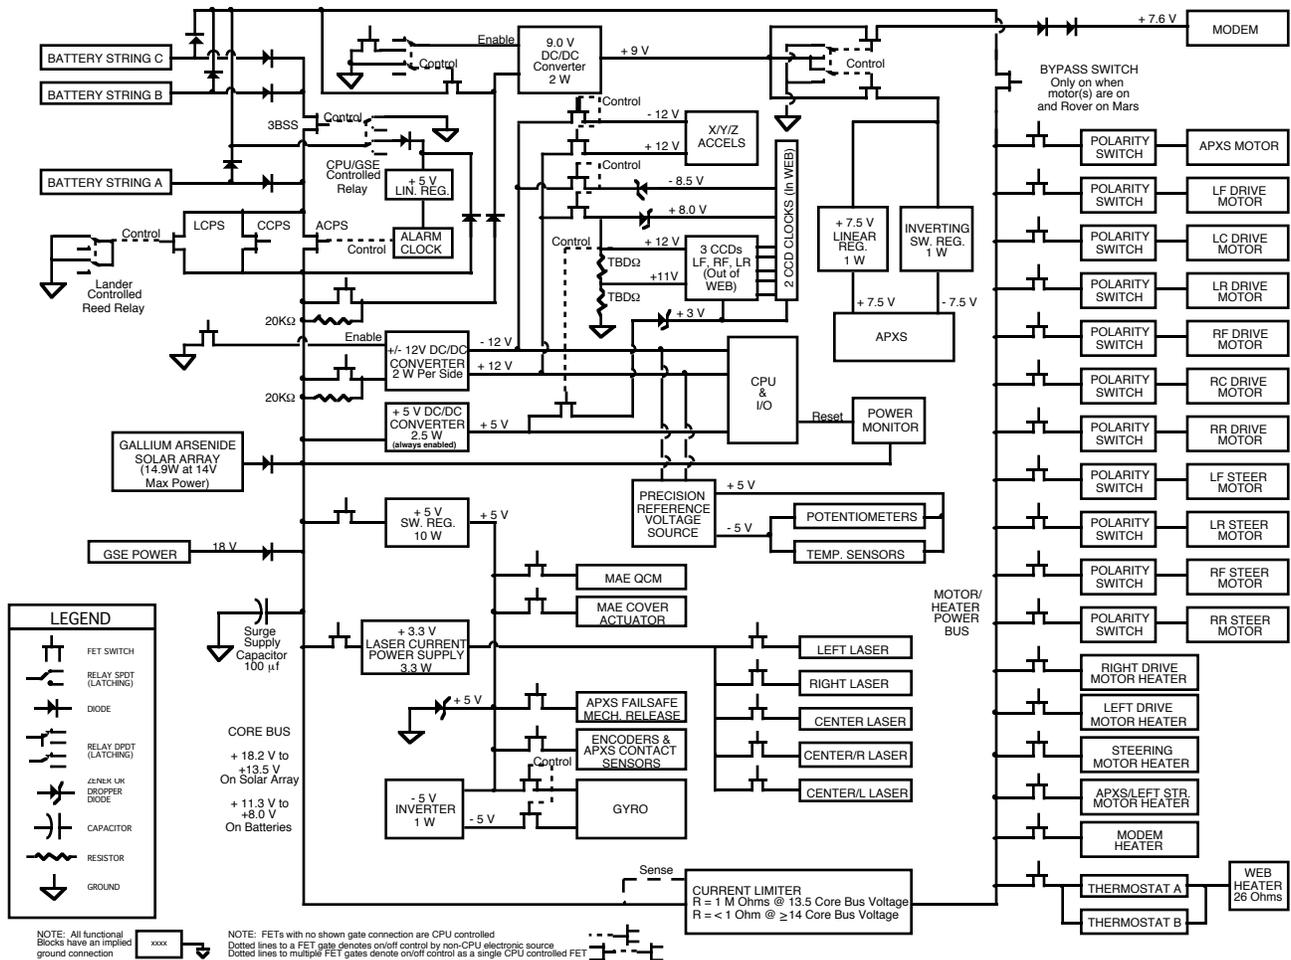


Figure 4: Rover Power System Functional Block Diagram

of energy left in the batteries is a significant task for a vehicle as computationally limited at this one.

The amount of solar power produced by the rover's solar array will vary as a function of sun angle (i.e., time-of-day), vehicle tilt, panel temperature, and atmospheric conditions. Under nominal conditions, however, the amount of solar power available will vary primarily as a function of the time-of-day as illustrated in Figure 3. Experiments and/or activities such as driving which require significant amounts of power must be performed as close to noon as possible. To make as much use of the day as possible, the rover has been designed to wake up and initiate operations with as little as 4.0 W of power.

THERMAL DESIGN AND CONTROL

The Martian thermal environment is extremely harsh. Night time temperatures are expected to go below -100°C

and day time temperatures are likely to reach only -20°C. In order for the rover's electronics and critical sensors to survive such an environment, they must be housed within a protective thermal enclosure. For the rover, this enclosure called the Warm Electronics Box (WEB) is also the chassis of the vehicle.

The WEB is a box like structure whose walls are formed from thin fiberglass face sheets bonded to "Z" shaped aluminum spars. The 1.0" gap between the face sheets is filled with blocks of an ultra low-mass version of solid silica aerogel which provide the primary means of insulation. The aerogel has a density of only 20 mg/cc and exceptional thermal properties.

The WEB, in conjunction with the vehicle's active and passive thermal control system, is designed to keep the rover's electronics, modem, batteries, APXS electronics, and other critical sensors between -40°C and 40°C

throughout the Martian day and night. This is accomplished by conserving the thermal energy which is pumped into the WEB throughout the day by an internally mounted heater. By the end of each day, the inside of the WEB is expected to be around 40°C. During the night, the internal temperature will drop. Based upon the design and thermal performance of the WEB, however, the temperature inside will not fall below -40°C before the sun has risen, the vehicle has woken up, and the control system is again actively heating the inside of the WEB.

While the WEB is key to the thermal design of the rover, the current limiter circuit described earlier also plays a important role. This circuit allows the WEB heater to be turned on and automatically draw just enough current to keep the overall vehicle operating at or near the solar panels peak power point. In other words, it permits all unused solar energy to be routed to the WEB heater without causing the Core Bus voltage to drop and thereby brown out the electronics.

Electronic components such as the rover's cameras, lasers and wheel encoders, and electromechanical components such as the wheel drive motors obviously cannot be housed within the WEB and hence must be capable of withstanding the ambient thermal environment of Mars. In the case of the cameras, lasers, and encoders, a rigorous environmental test plan was developed in order to qualify commercial components which, based upon their design, showed promise of being able to survive the environment. A similar approach was taken to qualify the motors and gear boxes used in the drive mechanism. In the latter case, however, the commercial components were modified in order to achieve the necessary performance and reliability.

TELECOMMUNICATIONS AND TELEMETRY

The rover receives commands from its ground-based operators via the lander. Similarly, it transmits telemetry data collected during the execution of each and every command to the lander which later forwards it to Earth¹. The lander and rover communicate with one another through a pair of UHF modems. The modems operate at 459 Mhz and have range of approximately 500 meters. They communicate at 9600 bps and achieve an effective data transmission rate of approximately 2000 bps. Both

¹ This mission does not include an orbiter. The lander can only communicate with Earth when the Earth and lander are in line-of-sight of one another (i.e., after Earth-rise and before Earth-set).

modems are thermally controlled, albeit not very tightly, to reduce drift within their transmit and receive oscillators and to minimize the bit error rate. The lander and rover are both equipped with omni-directional whip antennas. The rover's antennae is depicted in Figure 1 in its deployed configuration.

Although the rover transmits telemetry data to the lander immediately after it is collected, it is not immediately forwarded to Earth. There are two reasons for this. First, due to power limitations and the need to recharge its secondary batteries using solar energy, the lander only communicates with Earth twice per Sol, once in the early morning and once in the late afternoon. Second, rover telemetry and lander telemetry packets are prioritized according to the urgency with which the data is needed on the ground. Data critical to the planning of the next Sol's command sequence is downlinked first. In the case of the rover, such data would include an "end-of-day" image of the rover (taken by the lander) and rover telemetry which indicates where the rover thinks it is and how successful it was at executing its daily command sequence.

The communication protocol between the rover and lander, while conceptually straightforward, is quite complex and includes several layers of error detection and correction in order to achieve a high degree of reliability and data integrity. At the highest level, the rover collects telemetry data in the form of messages. Each message contains the data obtained during the execution of a single rover command. Messages are transmitted to the lander in the form of CCSDS formatted packets. The data field within a packet contains all or a portion of the message, depending upon the message's length, and the packet's headers contain packet field position information and error detection codes. At the lowest level, the data link layer, each packet is transmitted to the lander by first dividing into frames which have a maximum length of 2kbits. Frames are transmitted using handshaking-based protocol with its own level of error detection.. Thus, error detection is performed at both the frame and packet level. Errors are corrected in this system by retransmitting corrupted data blocks.

It is worth noting that the lander does not process the data contained within the messages sent by the rover. It simply forwards them on to Earth. Similarly, the lander does not process the rover command sequences which it receives from Earth and passes on to the rover. The rover thus processes all of its own telemetry data and command sequences in the same fashion as a full up spacecraft.

ONBOARD SENSORS

The rover contains an extensive suite of sensors which are used for navigation, hazard and obstacle detection, power management, fault protection, and collection of measurements for particular technology experiments. The types of sensors and their primary function(s) are listed in Table 1. Many of these sensors are used for more than one purpose such as the cameras which serve as both a principle component of the vehicle's proximity-based hazard detection system and a general terrain imaging system. The accelerometers, which measure the vehicle's attitude (i.e., pitch and roll), are used to determine if the vehicle is tilting excessively, to compute heading (via compensation of the rate gyro measurements), and to identify changes in mission state (e.g., pre-launch, cruise, landed-ops).

Qty	Sensor	Primary Function(s)
3	Accelerometers	Hazard Detection
1	Rate Gyro	Dead Reckoning
2	Bumper (chassis)	Collision Detection
4	Bumper (solar)	Collision Detection
2	CCDs (front)	Imaging; Proximity
1	CCDs (rear)	Imaging; Target Validation
3	Temp (CCDs)	CCD Calib.; Scientific Imaging
2	Temp (Motor)	Motor Performance Evaluation
7	Temp (WEB, Modem)	Thermal Control & Characterization
1	Temp (Solar)	Array Performance Evaluation
1	QCM	Dust Adherence; Particle Mass
2	Solar Cells	Power Management
11	Current (Motor)	Torque Monitoring; Fault Protection
3	Current (Batt)	Power management
5	Current (Reg)	Power Management
2	Bogie Position	Hazard Detection; Mobility
1	Differential Pos.	Hazard Detection; Mobility
6	Wheel Position	Dead Reckoning
4	Steering Angle	Direction Control
10	Voltage (Bus, Reg)	Power management/Fault Protection
1	Solar Cell	Wheel Abrasion
1	ADM Position	Mechanism Deployment Control
3	Bumper (ADM)	APXS Emplacement; Release Indicator
2	Real-Time Clock	Time/Alarm Clock

Table 1. Onboard Sensors

COMMAND AND CONTROL STRATEGY

The strategy for commanding and controlling the rover is based upon a combination of operator-based waypoint designation [7] and onboard behavior control [2]. The waypoint designation component deals with the ground-based planning of the activity sequence and the interactive selection of the locations through which the vehicle should travel in order to reach the desired activity sites. During this step in the process of controlling the rover, it

is the responsibility of the human operator to designate paths (i.e., sequences of waypoints) which are free of obstacles and/or hazards which could threaten the safety of the vehicle and jeopardize the mission. The behavioral component of the control strategy corresponds to the onboard algorithms which autonomously and safely navigate the vehicle from one waypoint to the next. These algorithms enable the rover to respond in real-time to the uncertainties inherent in navigating through an unknown terrain and in the presence of obstacles and/or hazards which were undetected by the operator.

The task of designating a waypoint is conceptually quite simple. Using the display capabilities built into the Rover's Control Station, the operator looks at the local Martian terrain in 3-D and chooses how best to get from one location to another. The 3-D images are obtained from the lander's stereo imaging system which is located approximately 1.5 meters above the ground. (The lander's cameras have an imaging resolution of 1 mrad/pixel) A joystick is then used to position a 3-D graphical model of the rover at locations (i.e., the waypoints) which, when connected by straight lines, define the nominal path through which the vehicle should travel. If the terrain contains numerous obstacles and/or hazards, the operator can space the waypoints relatively close to one another (e.g., 0.25 meters), whereas if the terrain is relatively benign only a few waypoints may be needed. The choice of how many waypoints to designate is up to the operator and the experiment team.

The primary advantage of the aforementioned control strategy is the inherent separation between the planning and control functions which require significant processing capabilities and can be performed on the ground, and those which require relatively little computational capability and can be implemented aboard the rover.

The locations through which the rover will travel and the activities that will be performed at specific sites are chosen by the experiment team based upon the results of previous activities, the most recent images of the rover and surrounding terrain, and the relative importance of the experiments yet to be completed. Once selected, the operator uses the interactive capabilities of the rover control station to construct the actual command sequence to be uplinked to the rover via the lander. The resulting executable command file is called a Rover Activity Sequence File, or RASF.

SOFTWARE DESIGN

The vehicle's control software is organized as a single control loop, with interrupt handling for a few

asynchronous “reflex” events (such as bumper contact) to which the rover must react quickly [5]. This loop dispatches periodic functions (e.g., thermal management, automatic vehicle health checks, and command upload requests) as indicated by software timers, and invokes command handlers as directed by the uploaded command sequence.

The command handlers all follow a common format:

- extract and validate any command parameters
- verify that the command is allowed based on the current rover state (e.g., mission phase and power availability)
- set a timeout limit for completion of the command
- perform the command
- format and send telemetry containing command execution results
- return a completion status

The architecture of the vehicle’s control system was largely dictated by the fact that the vehicle does not have enough power to “walk and chew gum at the same time.” For instance, there simply isn’t enough power to run the RF transmitter and wheel drive motors; or to run the wheel drive motors and the cameras and lasers; or to run the APXS instrument and the modem transmitter. In fact, the vehicle is so power starved, it must stagger start the wheel drive motors, especially when they are cold, to insure that each has access to enough power the instant they turn on and go to stall.

As a consequence, there was no real advantage to using a more conventional real-time, multitasking system architecture. In fact, the small benefits which could be realized would be overwhelmed by the number of issues which would arise concerning the reliability of multitasking kernel which would have to be developed from scratch for this vehicle’s computationally challenged 80C85.

NAVIGATION

The GoTo-XY command is the principle means for instructing the rover to move from one waypoint to the next while simultaneously avoiding obstacles and hazards. As described, the execution of the GoTo-XY command, along with the rest of the commands in the RASF, is performed completely autonomously. In addition to the GoTo-XY command, there are several other commands which can be used to reposition and/or reorient the rover. These “discrete-motion” commands, however, are intended for diagnostic purposes and do not

invoke feedback beyond that of servoing the drive and steering motors to specified positions.

The implementation of the GoTo-XY command consists of three primary components. The first component is the dead-reckoning algorithm which is used to compute an estimate for the location (i.e., the X-Y position and heading) of the rover relative to the lander. The heading estimate and current heading setpoint are used by the low level servo system to control the vehicle’s steering angle. The second component is the underlying navigation algorithm which simply drives the vehicle directly towards the specified waypoint. Ideally, if no obstacles and/or hazards are present and the vehicle has perfect traction, the vehicle will travel along a straight line to the waypoint. The third component is the behavior set which utilizes information from the onboard sensors to detect the presence of obstacles and/or hazards and generate steering/drive setpoints which override those generated by the underlying navigation algorithm. Once the vehicle no longer senses the presence of an obstacle/hazard and it has completed executing its avoidance maneuver, the behaviors return to generating a “null” output and the underlying navigation algorithm takes control of the vehicle.

OBSTACLE/HAZARD DETECTION

The Martian environment poses many potential threats to the safety of the rover including large rocks, complex boulder fields, cliffs, ravines, escarpments, steep slopes, and dust pits, just to name a few. In response, the rover is equipped with a variety of sensors for detecting the presence of critical physical hazards. These sensors range from simple potentiometers which measure the kinematic configuration of the mobility subsystem to the more sophisticated proximity sensing system which is comprised of 5 laser stripe projectors and 2 CCD cameras. While it is beyond the scope of this paper to describe all of these sensors, a brief description of the proximity sensing system is in order

The stripe projectors and cameras are mounted to the front of the vehicle just below the solar panel as indicated in Figure 1. The cameras have a focal length of 4.0 mm and an extremely wide field of view, 1.7 radians in the horizontal direction and 1.4 radians in the vertical direction. Each CCD contains 767 x 484 with a horizontal and vertical resolution of 2.9 mrad and 3.4 mrad, respectively. The CCDs, including the rear mounted color CCD, are controlled directly by the rover’s CPU and associated digital I/O ports. While images are scanned out quite slowly due to the computational limitations of the rover’s 80C85 processor, the effects of

dark noise are minimal due to the naturally cold temperatures at which the CCDs will normally be operated (e.g., -60°C to -20°C).

The stripe projectors generate vertical planes of light which create stripes on the surface of obstacles and the terrain in front of the rover. Each projector consists of a 860 nm, 200mW laser diode and a single cylindrical plano-convex lens which collimates the elliptical output pattern of the diode along the pattern's minor axis. The major axis of the diode's beam pattern is left uncollimated to form the vertical extent of the projector's output beam with a Full Wave Half Maximum (FWHM) of 30 degrees.

The stripe projectors are mounted so that they project vertical planes of light in a criss-cross pattern in front of the vehicle as illustrated in Figure 5. When viewed by the cameras, each stripe appears as a ragged line due to the irregularities in the terrain. Detection of the stripes is enhanced by the fact that the cameras are equipped with very narrow band filters centered at the wavelength of the laser diodes.

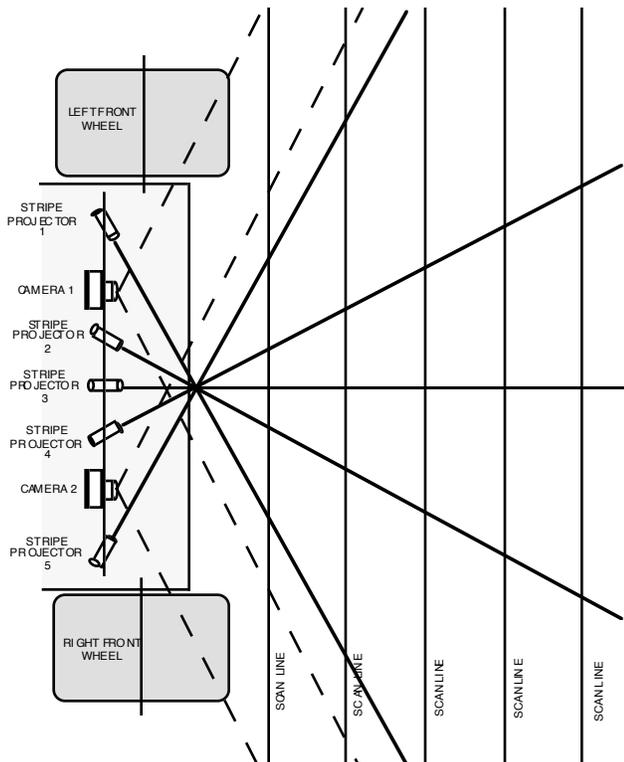


Figure 5: Plan view of stripe projector beam pattern.

The detection of the presence of hazards using this arrangement of cameras and stripe projectors is conceptually based upon triangulation and thresholding much like the focus on a point-and-shoot camera. In this approach triangulation is used to measure the range to an obstacle and/or a surface patch upon which a light stripe has landed. Thresholding is then used to determine whether or not the range measurement signifies the presence of a hazard.

When the vehicle is navigating, both the right and left cameras are used to sense for hazards and validate whether or not the vehicle can safely turn in place. As the vehicle drives forward it periodically stops, approximately every wheel radius, turns on the first stripe projector, acquires five scan lines worth of image data, turns the next stripe projector, acquires five scan lines of image data, and repeats this process until all five stripes have been imaged. The scan lines are then processed using a simple peak detector to determine where in each scan line the various stripes occur. The result is a 5 by 5 array of range measurements. This array is then analyzed in several ways to determine if an obstacle or hazard exists in front of the rover. In principle, analysis involves a simple comparison of the 25 pixel locations (i.e., range measurements) to the ranges at which these stripes would fall if the vehicle were on flat terrain. Range values significantly above the flat terrain model indicate the presence of an obstacle. Values significantly below that of the model indicate the presence of a drop off (e.g., trench or cliff). The thresholds used to signify the presence of a hazard are determined empirically during ground-based testing and contain a range of values each of which corresponds to different level of conservatism with regard to how safe it is to drive over or across the hazard.

OBSTACLE/HAZARD AVOIDANCE

The rover avoid obstacles and other navigation related hazards by invoking preprogrammed behaviors which override the default straight-line navigation algorithm. The particular behavior invoked depends upon which hazards have been encountered. For example, if proximity sensing system detects an obstacle to the right of the vehicle, the vehicle will stop, turn to the left and drive forward a short distance. If no other hazards are then detected, the vehicle re-invokes the straight-line navigation algorithm and again starts heading directly towards the next waypoint.

If the rover is in extremely rough terrain, it may not be able to reach the next waypoint before timing out. If this occurs, the rover will generate an internal error message (i.e., flag) which will then be used to determine which