Battle of the Rockets

Clayton Penney and James C. Squire

The team of nine students from the Virginia Military Institute (VMI) collectively held their breath during the countdown. The result of their yearlong project—an 11-ft-tall, 60-lb rocket—stood on the launchpad, holding an autonomous soil-collecting robot they had designed. There was little doubt the engines would ignite, developing 450 lbf of thrust within 200 ms; their worry was whether the robot would eject at apogee and survive the resulting 1,000-ft fall.

These were the opening minutes of the Battle of the Rockets Competition, hosted by the Federation of Galaxy Explorers, in Culpepper, Virginia, in April 2019. The competition allows high school and college students to enter different contests of varying degrees of difficulty that test their skills in rocket and robot design. There are three different contests; in order of complexity, they are the Target Altitude, Sounding Rocket, and Mars Rover events. This article describes the VMI's entry in the Mars Rover event, which involves designing an autonomous robot launched from a rocket. No teams have fully completed all facets of the Mars Rover challenge since the rules were made more stringent in 2015.

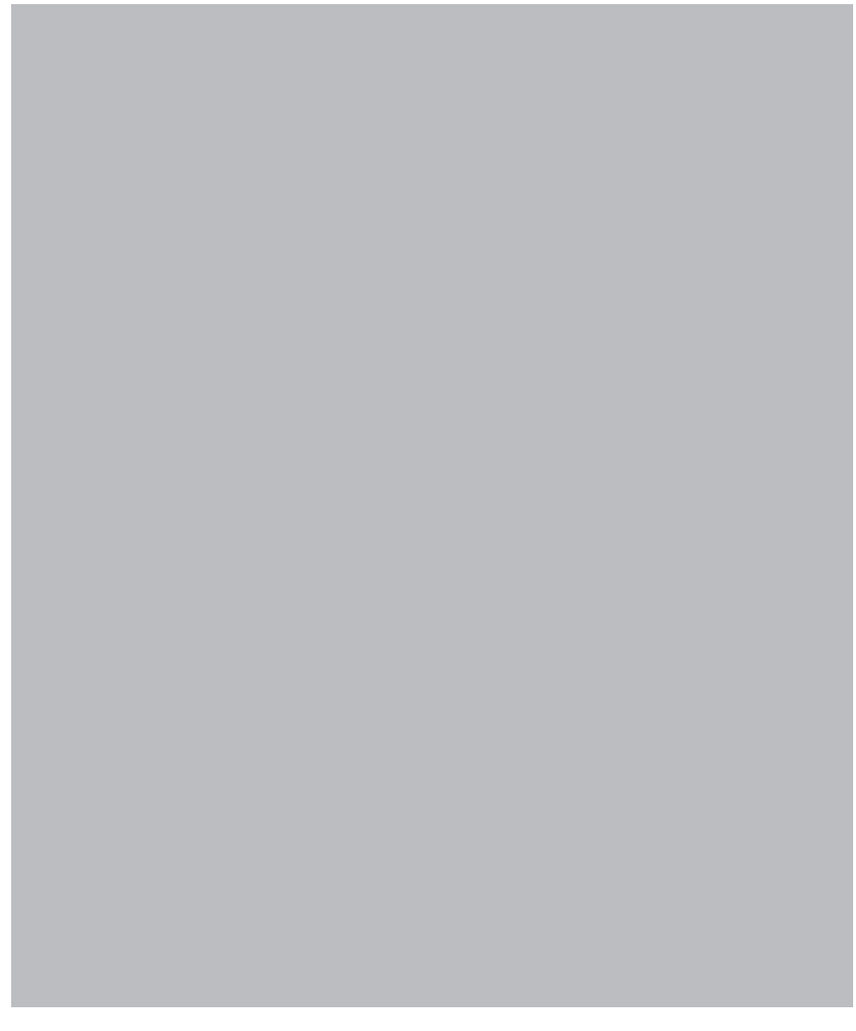

xxxxxx

The rules governing the Mars Rover competition are simple to state but difficult to achieve: to design a robot and rocket system that launches in the air, deploys, lands on the ground, collects soil, and performs

Digital Object Identifier 10.1109/MPOT.2019.2934015
Date of current version: 6 July 2020

 0278-6648/20©2020IEEE

telemetry (transmits data). These parallel the tasks several American, Chinese, and European agencies are planning in Mars exploration missions.

Specifically, teams must design a Mars rover robot and a rocket to launch it to at least 1,000 ft. The rover, weighing no more than 2 kg, must be enclosed within the rocket before launch and safely return to the ground by controlled descent, presumably by parachute. After the rover lands, it is required to disconnect from its recovery device, travel at least 3 ft in any direction, and await a command to continue its mission. A handheld wireless device must be designed to issue this command with a single button press. Once given the command, the rover has 5 min to collect at least 5 g of soil and place it in a detachable container. After the soil is collected, the team issues a second command from the wireless controller, directing it to take a picture of the collection site and send it by wireless telemetry to be displayed on the handheld device.

The rocket

The rocket consists of four sections (Fig. 1). From aft to fore, these are the booster, avionics (AV) bay, payload bay, and nose cone. The booster section houses the motor in the motor tube and the small drogue parachute that deploys at apogee (its point of highest altitude), and it provides an attachment point for the fins to be mounted to the rocket. The AV bay joins the booster segment to the payload section. It houses redundant flight computers that fire ejection charges to separate the booster when the rocket reaches apogee and fire more ejection charg-es after the rocket falls to within sev-eral hundred feet of the ground.

The payload section is located directly below the nose cone. It has two purposes: it contains the rocket's main parachute that deploys once the rocket falls from apogee to a point closer to ground level, and it houses the rover during the flight. The nose cone is the foremost section; in addition to providing a smooth aerodynamic shape, its forward weight helps stabilize the rocket's flight during ascent.

The motor

The VMI team wanted plenty of space for the rover, which implied a large rocket diameter and, therefore, a heavy rocket. The contest rules place no limit on the maximum rocket size or weight, but they do limit the total impulse (the integral of thrust over time) of the rocket motor to a maximum of 2,560 N·s. Accordingly, the design of the rocket was driven by the contest's constraint on motor size. To ensure the rocket leaves the launch rail with sufficient speed to allow the fins to keep the rocket stable, general guidelines suggest that the engines must develop at least five times more thrust than the mass of the rocket (i.e., it must have at least 5 g of acceleration). This implies that the rocket motor should develop a high thrust (greater than five times the rocket mass) for a short time to keep the total impulse under 2,560 N·s, rather than develop a lower thrust for a longer time.

A K1999N motor was selected because its 2,540-N·s total impulse is the closest available to the maximum allowed and it burns very quickly (in only 1.25 s), developing an average thrust of 1,999 N, or about 450 lbf. This sets an upper limit on the weight of the rocket at 90 lb. A survey of available high-power rocket kits reveals that this translates into an 8-in-wide fiberglass rocket. Specifically, simulations of the Madcow DX-3 kit show it will reach 1,500 ft when carrying a 2-kg payload with a K1999N motor. This kit comprises only a nose cone, fiberglass tubing, cut fins, and has no electronics, suggested motors, or even instructions. All design choices are left to the builder.

High-power rocket motors, such as the K1999N, are composed of ammonium perchlorate propellant in a rubber binder, such as was used in the space shuttle boosters. This formulation has roughly eightfold higher energy density than the black powder composites used in smaller model rocket motors, although they are considerably more difficult to ignite.

Stability

Wind and aerodynamic asymmetries push the rocket from a straight path upward and create wobble during flight; in extreme conditions,

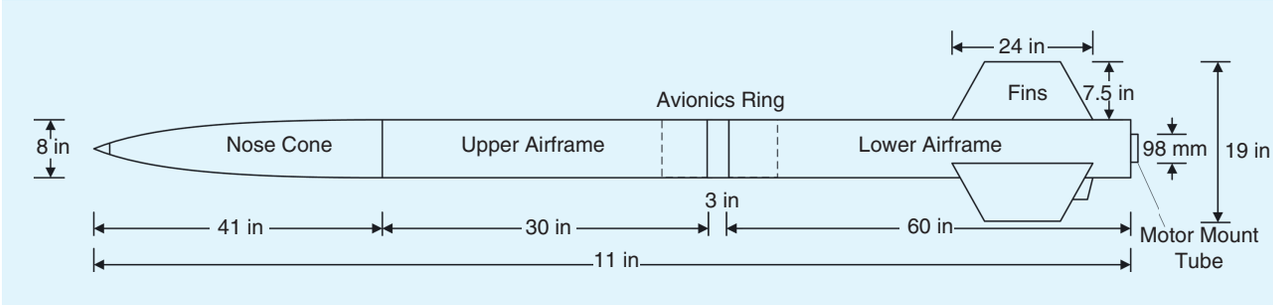

FIG1 The rocket's major sections: a nose cone that separates from the payload bay by ejection charges triggered by flight computers in the avionics bay. A separate set of charges releases the booster, which falls under its own parachute.

an unstable rocket will follow a highly erratic trajectory that rocketeers call *skywriting*. The stability of the rocket is determined by the relative location of the rocket's center of gravity (CG) and center of pressure (CP). When in flight, the rocket rotates around its CG; however, the lift and drag forces act around the CP. Although the drag force remains mostly vertical during small angle changes during flight, the lift force is parallel to the ground.

If the CP is located below the CG, the lift and drag forces act in a restoring manner and cause the rocket to angle itself vertically again. If the CP is above the CG, the lift and drag forces will act in a destabilizing manner and cause the rocket to further deviate from the correct path of travel. This deviation quickly compounds, and then the rocket begins to skywrite. The rocket described in this article has a CG that is 8.9 in higher than its CP, as calculated by the RockSim simulation package (Apogee Components), ensuring flight stability.

Recovery

We used a dual-deployment recovery system, meaning that the rocket deploys different parachutes at two different times. Although this system is complex, it increases the chance that the rocket will return safely to the ground without significant lateral drift. This system begins when the flight computer in the AV bay detects that the rocket has finished its ascent and arced over at apogee. At this moment, the flight computer sends an electrical current through an igniter buried in 4.25 g of black powder in the aft side of the AV bay, causing a sudden increase in pressure in the booster that separates it from the rest of the rocket.

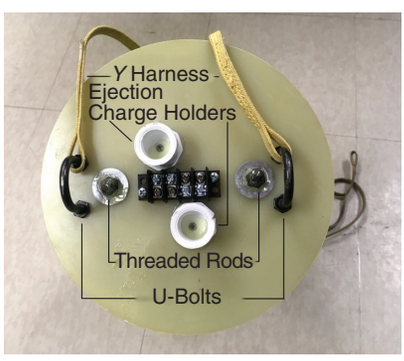

FIG2 The bottom of the AV bay, showing the white wells that hold black powder ejection charges, electrical bus that will be connected to e-matches to ignite the ejection charges, and epoxied ends of the all-thread rods that join the caps of the AV bay together.

The booster remains connected by a 40-ft shock cord, which has a small drogue parachute tied along its length. The drogue deploys, and it, coupled with the nonsmooth shape of the separated booster and separated AV bay/payload/nose cone section, slows the rocket descent speed to approximately 72 ft/s. This speed is fast enough that the rocket body will not drift much laterally but slow enough that the main parachute will not rip apart when it deploys.

We programmed the flight computer in the AV bay to fire a second set of ejection charges, this time in the fore direction through the payload bay, when the rocket is 400 ft from the ground. This pushes three things out of the payload section: the nose cone (which, unconnected from the rest of the rocket, flies away freely under its own small parachute), the rover (also unconnected, which descends under its own small parachute), and the rocket's main parachute (which is tethered to the rest of the rocket by another 40-ft shock cord).

These three parachutes allow the two rocket parts and rover to each descend gently to Earth at approximately 20 ft/s. This process is entirely dependent on the flight computer: if it fails, the rocket and rover will have a catastrophic l anding. To re-duce this possibility, two electrically independent flight computers are used, each of which is connected to redundant ejection charges.

Simulations

RockSim simulations, including a mass representing the rover, show that the rocket leaves a 12-ft launch rail at 57 ft/s, which is sufficient for stable flight. It maintains stability through the powered phase of flight, developing up to 6.5 g of acceleration for 1.24 s, at which time it is traveling at 326 ft/s, or approximately 225 mi/h. It then commences its unpowered, coasting phase of 0-g flight for 11.4 s before arcing over at an apogee of 1,495 ft, well in excess of the required 1,000-ft floor. The simulated rocket deploys its drogue parachute at that time, and, 26.4 s later, it reaches 400 ft from the ground, at which time it deploys its main parachute. This slows it down to a 24-ft/s landing, 72 s after launch.

AV bay

The AV bay is a 2-ft-long tube mounted between the payload bay and the booster (Fig. 2). It is capped by bulkheads on both sides connected by threaded rods, so the inside is pneumatically isolated from ejection charge pressures in the adjoining lower booster and upper payload sections. The outward-pointing faces of the upper and lower bulkheads each have two black powder charge holders and two U-bolts. A fire-resistant Kevlar shock cord connects the fore-side U-bolts to the main parachute. A second Kevlar cord connects the AV bay's aft-side U-bolts to the rocket's booster and the drogue parachute.

A hole called a *static port* is drilled through the exterior of the AV bay to ensure that the outside air pressure equalizes with the interior, which enables barometric pressure sensors on the flight controllers to calculate the altitude of the rocket. The placement of the static port was given careful consideration since small projections above it could influence the surrounding air pressure.

The primary and redundant flight controllers are housed inside the AV bay. These are connected by wires, that run through the fore and aft AV bay bulkheads to the black powder charge holders. The aft ejection charge fires when the rocket

reaches apogee, separating the booster from the rest of the rocket and releasing the drogue parachute. The fore ejection charge fires when the rocket falls to approximately 400 ft above ground level, separating the nose cone (which falls under its own parachute), ejecting the robot (which falls under its own parachute), and also ejecting the main parachute attached to the rest of the rocket.

As the rocket ascends, decreasing external air pressure often causes air trapped in the lower rocket to separate at its junction with the AV bay and, similarly, tends to cause separation at the payload bay/nose cone junction. These events should be initiated only when the flight controller triggers an ejection charge to fire. To prevent this, we connected these junctions with small plastic screws called *shear pins*. These pins are designed to shear under the pressure difference induced by the ejection charge but not under the smaller pressure differences caused by rocket ascent.

To further discourage premature separation, small vent holes are placed in the payload and booster sections, large enough to reduce altitude-induced pressure differentials but small enough to allow the ejection charge to create a pressure differential large enough to shear the pins. Much larger-diameter screws are used at the junction between the AV bay and payload to prevent separation under any flight conditions.

Two additional holes placed through the AV bay wall allow external access to two safety switches that ensure the ejection charges cannot be fired before the rocket is safely on the launch rails. These disconnect safety switches are a contest requirement.

Two separate flight controllers redundantly monitor the attitude (angle) of the rocket and its altitude to determine the optimal time to deploy the fore and aft ejection charges. The primary flight controller is the Eggtimer TRS (Eggtimer Rocketry), a kit built from surface-mount components that requires soldering. The backup flight controller is the Telemetrum Easymini (Altus Metrum). Both controllers are powered by batteries that are secured by aluminum brackets and redundantly secured with zip ties to reduce chances that the 6.5 g of flight acceleration will dislodge them.

The amount of black powder used for each ejection charge is critical. If there isn't enough, the rocket does not separate, and the corresponding parachute fails to deploy. If there is too much, the booster or payload ruptures. We experimentally determined that 4.25 g for the fore side and 2.75 g for the aft side, proportional to the volume of free space in each section, worked well.

Typically, flight computers are secured onto a plywood sled and are connected to the batteries, mechanical safety switches, and ejection charges through wire harnesses. A common failure mechanism for such designs is for wires to disconnect under the heavy rocket acceleration, resulting in a lawn dart: a rocket whose ejection charges do not deploy and that lands at high velocity, nose downward.

To mitigate this cause of failure, we designed a printed circuit board (PCB) on relatively thick 1/8-in fiberglass to mount the primary and backup flight computers, their batteries, and the disconnect switches, thus eliminating both the wiring harness and the sled as a separate structural component (Fig. 3). Like the more typical plywood sleds, the PCB is attached to two threaded rods that are epoxied into the AV bay's endcaps. The placement of the antenna for the Eggtimer TRS that provides GPS telemetry was given special attention to avoid attenuation by being electromagnetically shadowed by the threaded rods.

The rocket's flight is filmed using two video cameras. These are mounted just aft of the AV bay on the outside of the booster section on opposing sides of the rocket to prevent asymmetry. They are attached to the external rocket body using a 3D-printed mount that permits access to the video-recording switch.

Mars rover

The Mars rover competition has a number of strict design constraints. One is physical size: although the rocket has an 8-in diameter, the rover must fit into the narrower 7.75-in coupler tube joining the nose cone and pay-load bay. Second, contest rules require the rover to weigh fewer than 2 kg. These quantitative con-straints are joined by the derived requirement of durability, since the rover is exposed to 6.5 g of launch acceleration, approximately 15-g acceleration from the ejection charge at apogee, and significant landing forces.

Once the rover lands safely, it must detach itself from the recovery system and travel 3 ft autonomously. Contest rules discourage the use of wheels or tracks, so the rover uses a combination of wheels and legs called *whegs*. These spindle-like projections from the hub, like a wagon wheel without the rim, traverse rough terrain better than standard wheels or tracks. They are capable

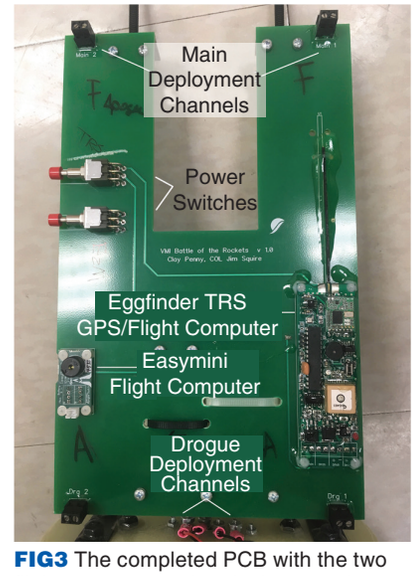

FIG3 The completed PCB with the two flight computers, electrical connections for the e-match ejection charge ignitors, and separate disconnect switches labeled.

of traversing obstacles almost twice their height, which is helpful when navigating a harvested cornfield, and are used in the competition (Fig. 4). The whegs were designed to rotate in the same direction to move the required 3 ft. Then they counterrotate, causing bottle-top nails to scoop dirt, which, as the wheg rotates, is then deposited into a channel in the wheg that directs it to a collection tray on top of the rover.

FIG4 The Mars rover staged in a cornfield for a full systems test. The whegs rotate together for horizontal movement and then counterrotate, causing the bottle tops to scoop dirt into angled channels in the whegs that ultimately deposit it into the red top-mounted trays. The system is vertically symmetric; should the rover land upside down, it will work in the same manner. The camera is visible at the front and the parachute-detachment mechanism in the rear.

Competition day

The morning was cool and damp on the day of the Battle of the Rockets competition. Weather forecasts predicted 2 h of calm weather before the wind would pick up, but we were not concerned; countless system tests in our lab had taught us how to prep the subsystems in minutes. Four people retested the Mars rover, while the rest of the team began assembling the rocket. We secured the motor tube retention strap to the AV bay while the ejection charges were packed. Once the flight computer verified charge ignitor continuity, we loaded the parachutes and confirmed that all of the harnesses were attached correctly.

Unfortunately, the rover testing team reported back with problems. The Mars rover, which was thoroughly tested the night before, was now refusing to turn its rear whegs, although its front pair worked normally. After 45 min of troubleshooting, we located an intermittent open circuit in the motor wires. We resoldered all of the motor wires onto a duplicate PCB to avoid the chance of a second wire failing. Using a cooler as a flat surface, we kneeled in the dirt to get close enough to a car's cigarette adapter, which was our power source for the soldering iron. Imagine trying to quickly resolder small, critical connections in an increasingly windy and dusty environment while staying aware of the flight paths of the other rockets being flown. It was not easy.

Once we replaced the PCB, we tested the rover again. It still did not work. Another 20 min of troubleshooting revealed that the voltage regulators in the spare PCB were –5 V instead of 5 V. Among all of the subsystem testing we had done, we forgot to test the spare PCB. The only correct regulator we had was in the original PCB. We carefully desoldered it and then resoldered it into the new PCB without breaking its delicate legs or overheating it. Thankfully, we were successful, but it was one of the most nerve-wracking soldering experiences of our lives. Once that was completed, all of the rover systems tested perfectly; we quickly loaded it into the rocket. It was finally time to fly.

We carried the rocket to the remote launchpad, mounted it to the launching rail (Fig. 5), and turned on the flight computers. Everything came up smoothly until an error message about charge continuity appeared on our LCD receiver, although we double-checked continuity during assembly. Power-cycling yielded the same error message. We needed to take the rocket back down and open up the AV bay to investigate, which was complicated by the fact that we were now out at the remote launchpad, away from our rocket stand and tools, and with the wind picking up.

Three people balanced the rocket on their knees, while two others pulled the pieces apart very carefully so we didn't disrupt the payload bay

FIG5 The rocket mounts to the launch rail by two buttons that slide into a C-shaped gap machined along the length of the rail. This requires the launch rail to lie horizontally during the mount procedure. Here, the rail and mounted rocket are being raised into launch position.

or ejection charges. We were able to get everything opened up with just enough room to slide one hand inside the AV bay. Using our fingertips, we moved two sets of ejection charge wires to the redundant flight computer, while smaller rockets continued to launch behind us and the wind blew dust everywhere. We reassembled the rocket and had it back on the launch rail in about 10 min.

We were finally ready for launch. The countdown began, and the launch button was pressed. Silence; nothing happened. The process was repeated: again, nothing. Because the flight com-puters were armed, only the range safety officer and student team leader were permitted to approach the rocket. Was it just a bad ignitor or a more fundamental problem? Although the original igniter looked fine externally, a quick continuity test revealed it had internally opened when electric current was applied during the launch sequence. We replaced it quickly and returned to the firing line: the countdown began again. By now, the wind had picked up considerably and was gusting at 15 mi/h.

Launch! The K-class motor took 1 s to build up its full 450-lb thrust, a time that seemed like forever; then, there was a thunderous roar, and the rocket spat off the launch rail in a smoky rush. It looked so beautiful going up. The rocket reached apogee, and everything ejected correctly. The nose cone's parachute deployed fully, and it landed safely in a nearby field.

However, the high winds whipped the 10-ft diameter main parachute around, and, like a thrown bola, it swept into the rover's parachute. The two entangled, and continued to wrap around each other while descending, with shroud lines immobilizing two of the rover's whegs. The entwined parts landed in a cow field, where ground wind blew the main parachute and rocket/rover into an electric fence. Wind gusts caused the main parachute to rotate, repeatedly pummeling the attached rover against the ground.

FIG6 The team with the VMI Mark 1-9 rocket on competition day. The author is in the front row, second from left. ***<AU: Please provide the name of the author being referred to in the photo.>***

Once we determined that the rover could not free itself, we carefully crawled through the electric fence and untangled everything. We decided to run the rover just to see what would happen. It worked flawlessly. However, it was disappointing to see it work so well since the failure to detach from the recovery system disqualified the entire attempt.

The tangling was unfortunate and not the ending we envisioned. However, we were proud of the way we performed under pressure and that every subsystem we designed worked independently. The team accomplished a successful launch, nothing broke, and the rover worked (Fig. 6). The experience taught us more about both project management and engineering than any class-based course we have taken, and we are grateful for the support of our advisor, department, and the Battle of the Rockets organizers.

About the authors

Clayton Penney (penneycd19@mail.vmi.edu) earned his B.S. degree in electrical and computer engineering from the Virginia Military Institute in 2019. He currently works in defense contracting and plans to attend graduate school in fall 2020.

James C. Squire (squirejc@vmi.edu) earned his B.S.E.E. degree from the U.S. Military Academy and his Ph.D. degree from the Massachusetts Institute of Technology. He is the Jamison-Payne Professor of Engineering at the Virginia Military Institute. He was awarded a Bronze Star in the U.S. Army during Operation Desert Storm and was selected as Virginia's Rising Star professor. He is a licensed professional engineer in Massachusetts and Virginia and maintains an active consulting practice.