Design of a Two-Wheeled Mars Rover with Sprawl Ability and Metal Brush Traction

Ivan Davydychev
Ronald S. Fearing, Ed.

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2019-145
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-145.html

December 1, 2019

Design of a Two-Wheeled Mars Rover with Sprawl Ability and Metal Brush Traction

Ivan Davydychev

January 29, 2017

I - Introduction

Purpose

This paper presents the design and development of a prototype of an expendable rover that is meant to explore the difficult-to-reach regions of Mars. The mission plan is to stack multiple such rovers in a container inside a parent rover such as *Curiosity*. When the parent rover encounters a feature on Mars that is worth exploring but is either too steep, too cramped, or too dangerous for it to approach, it can eject one of these expendable rovers to, as it were, "do the dirty work." If an expendable rover breaks or falls into some crevice and gets stuck, the greater mission is not jeopardized, because the parent rover (and its whole stack of other rovers) is still safe and functional. The locomotion strategy for the expendable rovers is to be minimalist to conserve space. Besides the locomotion electronics, each expendable rover is to have 1) a microscope and a camera for data gathering, 2) an antenna/processor for communication with the parent rover, and 3) a solar panel to recharge the rover's battery.

The ability to stack multiple such rovers in the parent rover without undermining their eventual mobility requires these rovers to be *collapsible* – to start out flat and to be able to unfold themselves. Hence, their given name: PUFFERs, pop-up flat-folding explorer robots. In particular, the PUFFERs were designed to have sprawlable wheels – wheels that would start prostrate but through some linkage action would unfold to turn upright. Two versions of PUFFER have so far been developed: a UC Berkeley prototype whose sole focus is mechanical performance, and a derivative JPL prototype that incorporates the advanced electronics listed above. This paper deals almost exclusively with the former; the latter is documented in Karras et al. (2017).

PUFFER was designed to optimize three general performance metrics: mobility, impact resistance, and collapsibility. Mobility was quantified by the maximum step height the rover could overcome, the minimum overhang height the rover could crawl under, and the maximum incline angle the rover could go up. For true mobility the rover also had to be able to flip itself over. Impact resistance was quantified by the maximum height from which the rover could consistently survive a fall (in its worst-case orientation). Collapsibility was quantified by the dimensions of the smallest box into which the rover could fit when collapsed (with an emphasis on the smallest dimension, *height*). Note that although the rover was optimized for applications on Mars, the above criteria make it equally useful for *terrestrial* applications such as disaster relief or terrain exploration. (It turns out that collapsibility, though not directly useful on Earth, greatly improves mobility.)

Overview

Figure 1 shows the final version of the UC Berkeley prototype of PUFFER – in both its sprawled (a) and unsprawled (b) state. When fully collapsed, the rover can fit in a 17cm×7cm×2cm box, whereas when fully expanded its bounding box is of size 14cm×14cm×7cm. These dimensions indicate a nearly six-fold change in volume, where (passive) tail expansion accounts for roughly a factor of two and (active) chassis expansion for roughly a factor of three. Both these components, as well as the rover's wheels, are discussed in detail in the next three sections. It is, however, instructive to give a short overview here.

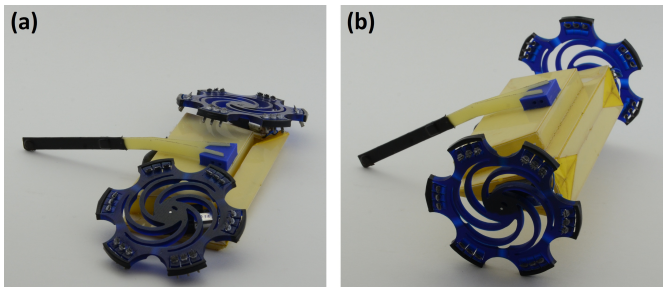

Figure 1: (a) sprawled and (b) unsprawled configurations of PUFFER. Note that although this is not shown here, the tail can be collapsed on top of the chassis for further reduction in dimensions. See Figure 6.

The chassis is a 3-D linkage that allows the rover to compress its shape for storage and also protects the rover's payload from dust and impacts. A novel aspect of the chassis is that it accomplishes the latter task by using both *rigid plates* and *tent-like structures*. From a mobility standpoint, the chassis allows the rover to change its sprawl angle based on what topography it has to get across. A high sprawl angle (i.e. wheels upright) makes it easier to go over rough, unsloped terrain, while a low sprawl angle (i.e. wheels prostrate) makes it easier to climb up inclines and crawl under overhangs. The chassis requires only one actuator.

Sprawl ability is not novel among mobile robots (Komsuoglu et al. 2008; Spenko et al. 2008; Zarrouk et al. 2013). Of those cited, STAR (Zarrouk et al. 2013) in particular has the same ability as PUFFER of being able to change the sprawl angle of every one of its wheels with just one actuator. What makes PUFFER so unique is that it uses SCM linkages (Wood et al. 2008; Hoover and Fearing 2008; Casarez and Fearing 2016) instead of gears to accomplish this task. Linkages (if designed correctly) are more robust than gears when it comes to impacts, especially when the links enclose most of the rover like the armor of a tank (see Figure 1). This way impact performance is partly evened out across all the possible impact angles, thereby increasing the maximum drop height.

The tail has three main functions: to keep the rover from flipping over during forward motion, to offset the rover's pitch-back moment during slope climbing, and to

minimize the rover's downward slides whenever it encounters a steep patch with bad traction. A novel aspect of the tail is that it is made of two halves that each have a high vertical stiffness but a low lateral stiffness. This design allows the tail to fold on top of the chassis for a reduced storage space yet to be able to bear all the required loads when expanded to its full length. The tail expands via stored elastic energy; no extra actuator is needed.

Most two-wheeled robots in the published literature likewise use a tail for similar functions (Stoeter et al. 2003; Murphy and Sitti 2007; Daltorio et al. 2009; Carpenter et al. 2015). The most analogous of these robots to PUFFER is Minnesota Scout (Stoeter er al. 2003), which likewise has a cantilever tail. This tail, however, has *low* vertical stiffness and *high* lateral stiffness, which is the opposite of what is found on PUFFER. The advantage of the PUFFER tail design is that folding stiffness is decoupled from load-bearing stiffness, meaning that with a given stress level the tail can fold into a smaller volume. Scout, however, has the advantage that with extra actuation its tail can be used to make the robot jump over obstacles.

The wheels are the traction and locomotion mechanism of the rover. They are designed to be able to go over the maximum variety of features possible, including steep slopes. The wheel rims are crenelated so that the rover can get over obstacles more than half the wheel diameter in height, while just inside the rims are clumps of metal wires that improve the rover's traction when the wheels are sprawled. These metal *brushes* were chosen as the traction system because they had the best performance of the many setups tested. Their main advantage is *load sharing* – having multiple points of contact with the ground.

Metal brushes as a traction method are akin to metal *spines*, which have been shown to perform well as climbing structures on multiple robotic platforms (Asbeck et al. 2006; Spenko et al. 2008; Daltorio et al. 2009; Carpenter et al. 2015). These robots all likewise take advantage of load sharing – by making the spines able to move relative to each other (either with compliant anchor mechanisms or with special linkages). Metal brushes are different in that they rely on the cantilever compliance of their own wires themselves to allow the relative motion required for load sharing. Because space-consuming anchor mechanisms are thus not needed, brushes can achieve load sharing over more concentrated areas compared to rigid spines. In this way the brushes are akin to gecko-inspired polymer adhesives (Aksak et al. 2007; Kim et al. 2008; Birkmeyer et al. 2012), though without the fouling problems. The main detriment of brushes in their current form, though, is that their wires must be oriented normal to the ground for optimal load sharing, which is not the best angle for asperity contact.

II - Chassis Design

Smart Composite Microstructures

The rover's chassis serves two main purposes: to give the rover the ability to sprawl and to protect the rover from dust and impacts. Sprawl ability is essential because it allows the rover to compress its shape for storage and because it improves the rover's climbing performance (see Section IV). To minimize actuation needs, the sprawl mechanism was chosen to be symmetric – one actuator to set the unloaded sprawl angles of both wheels to the same value. The sprawl transmission from the actuator to the wheels was chosen to be made with the smart composite microstructures (SCM) process, as it allowed the creation of complicated linkage structures without any incremental costs per linkage (Wood et al. 2008). Another benefit of the SCM process was that it made *plate-shaped* linkages, meaning that the links could also double up as the rover's armor.

The SCM process for the UC Berkeley prototype of PUFFER is relatively simple. It incorporates only the steps needed for mechanical performance, without any circuit traces (Hoover and Fearing 2008). Each PUFFER chassis is made of a stack-up with five layers: one layer of ripstop nylon sandwiched between two layers of thermal adhesive which are in turn sandwiched between two layers of PET. The nylon acts as the flexure layer; the PET as the rigid reinforcement. Each of these layers has patterns cut into it with a laser cutter to get the desired linkage. The layers are then laminated onto each other with the thermal adhesive. A final laser cutting step then frees the chassis linkage from the unused regions of the stack-up.

Chassis Linkage

The challenge with PUFFER chassis design was to make a 3-D linkage that could vary the wheels' sprawl angle ρ from 0° to 90° while also protecting the rover from impacts *at all of these sprawl angles*. The linkage also had to fit in the allotted space when collapsed (that is, when $\rho = 0°$). This meant in practice that the collapsed linkage could not have too many links stacked on top of each other, as each link had a thickness of 0.8mm[1] and most of the available height was reserved for actuators. The link thickness was chosen to be that value to give the chassis enough rigidity. One other requirement for the chassis was that it had to have a rigid cavity large enough to hold the payload/electronics.

These requirements were all satisfied with the following design: a planar four-bar linkage (the body) to which are coupled two spherical four-bar linkages (the sides). The combined linkage ends up having only eight links in total because the two largest links are shared by all three of the sub-linkages. Due to four-bar constraints, the combined linkage only has one degree of freedom in

[1] equal to the height of the five-layer stack-up.

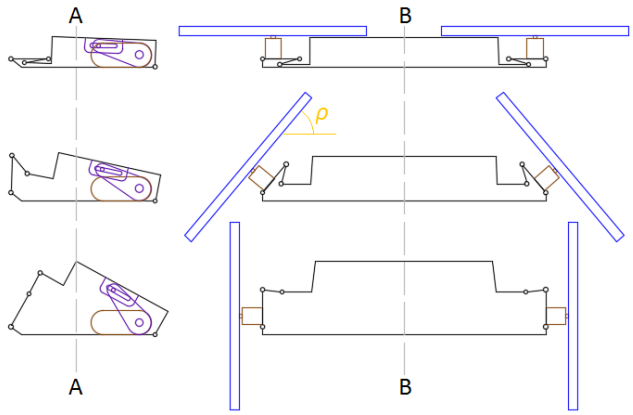

Figure 2: linkage diagrams of the chassis for various ρ. The cross sections on the left show only the planar four-bar linkage; the cross sections on the right show all three of the four-bar linkages. The circles indicate where the flexures (i.e. the connections between the links) are; all bends that don't have circles on them are rigid. Dashed lines indicate where the cross sections for the orthogonal views are taken.

expansion. Linkage motion is driven by a servo motor that controls the angle of the planar (i.e. the central) linkage. Figure 2 shows a set of linkage diagrams, taken at two orthogonal cross-sections through the rover, that illustrate how the sub-linkages are coupled. The wheels are shown in blue for reference; the tail is not shown for simplicity.

The servo mechanism that controls the chassis shape is seen in the left cross sections of Figure 2. It sets the angle between what are called the *bottom plate* and the *cavity link*, the two links that are shared by all three sub-linkages in the system. There is a slider connecting the servo horn to the cavity link because the servo axis is not aligned with the flexure connecting the two links. (Doing so would have been impossible without putting the servo partly outside the chassis.) The two short links on the left are called the *restrainer links*. They serve mainly to limit the expansion of the linkage, though they also protect the front side of the rover from impacts. Note that in the fully expanded state these links have a singularity. A small restoring spring (not shown) is needed to make sure that they fold back in the right direction.

The cross sections on the right show how the servo angle affects the sprawl angle ρ of the wheels. At the top, $\rho = 0°$; on the bottom, $\rho = 90°$. Note that because these cross sections now include *spherical* linkages (made with triangular links), incremental shifts in the plane A-A will cause the cross sections to change. In other words, the flexures are not all normal to the plane of the paper. Each set of three flexures converges at either the back right or the back left corner of the robot. (This convergence is what makes the corresponding four-bar linkages spherical in the first place.) In these cross sections the cavity also appears to bow out with increasing ρ. This deceptive geometry comes about as a result of a trapezoidal prism (the cavity) being cut by the cross-sectional plane A-A at an ever-increasing angle.

This design with three coupled four-bar linkages was chosen over other designs (such as one with two planar, orthogonal six-bar linkages) because it allowed the most space for the cavity. Here, only three of the rover's four bottom edges require stacks of four links in the collapsed state (see Figure 2), leaving the fourth edge free to be bordered by the cavity. A design with only planar linkages would have required such stacks on *all four* edges, leaving only a small space in the middle for the cavity. The design with spherical linkages is also easier to actuate, as rotary actuators tend to be easier to implement in a confined space than linear actuators (which is what the planar design would have needed). Finally, the design with spherical linkages only has two undesirable "gaps" in the chassis that are uncovered by rigid links (more on this later); a planar design would have had four such gaps.

The dimensions of the chassis (and of all the associated linkages) were chosen to maximize the volume of the cavity while meeting the rover size constraints set by JPL[2]. The optimization condition was chosen this way so that the rover could carry the maximum possible payload for its size. Note that the linkage design of Figure 2 has very favorable properties for maximizing cavity height, as it allows the rover's two thickest objects (the cavity and the wheel motors) to stack in parallel rather than in series. It also limits the height of the link stack at the cavity to just two links (i.e. the bottom plate and the cavity link), thereby leaving the maximum possible vertical space for the payload. In other words, had the optimization condition been to minimize the rover height given a certain payload height, no other design could have performed better.

The rigid cap that forms the cavity serves another purpose besides making space for the electronics. It strengthens the upper half of the rover by increasing that link's bending moments of inertia. The bottom plate, likewise, has short rigid extensions that increase its bending resistance. The rigid extensions on the bottom plate also help account for the nonzero link thickness as four links are stacked on top of each other near the rover's edges. A third purpose of these extensions is to reduce the bending angle of the crucial side flexures from 180° to 90°, thereby helping to increase these flexures' longevity.

Tent Flaps

As mentioned earlier, the chassis has two gaps that are not covered by rigid links. These gaps are at the front of the rover and are shaped like quadrangles (see Figure 1b). The gaps came about as a result of an inability to find an origami shape that could both cover the entire rover at all sprawl angles and also satisfy all the chassis objectives

[2]and other miscellaneous constraints such as leaving enough area on the spherical linkages for the wheel motors.

mentioned above. Fortunately a modification to the usual SCM process allowed the gaps to be filled with fabric – namely, the nylon layer of the SCM stack-up. The modification involved cutting away *plates* of PET (as opposed to just thin lines for flexures) – something that has not been done by previous SCM approaches. Although the resultant nylon flaps do not provide the best protection against impacts to these gaps, at least they protect the rover's cavity from dust and other contaminants. In other words, the nylon flaps act like tents, and hence are called *tent flaps*.

Impact Absorption

Another feature of the chassis is that even with a constant servo angle, the spherical linkages have some compliance, meaning that unequal loading on the wheels will cause each wheel's angle to be perturbed a bit from the intended sprawl angle. The compliance comes about mainly as a result of the non-infinitesimal width of the flexures (which causes the flexures to not behave as perfect pin joints and thereby introduces extra degrees of freedom into the system). The design rule here is that a larger ratio of flexure width to link length gives more compliance. A *built in* amount of compliance is not a bad thing, as it helps absorb rover impacts. PUFFER was designed to take advantage of this compliance. Other robots have taken advantage of SCM compliance for this purpose as well, and in many cases have achieved even better impact absorption than PUFFER because their more elastic linkage plates also contribute to the compliance (Birkmeyer et al. 2009).

For effective impact absorption at all sprawl angles, there must be a clear direction for all the linkages to collapse in when hit – in other words, there must be no linkage singularities. This is why in practice the servo is never actuated to the angle where it is constrained by the planar linkage (even with the restoring spring at the front of the rover). This is also why the spherical linkages are designed such that their shortest links are never parallel to the cavity link, even at $\rho = 90°$ (see bottom right of Figure 2). Otherwise a direct lateral hit to a wheel would not trigger an effective impact absorption motion.

2-D Layout

Figure 3 shows an AutoCAD layout of the chassis in its unassembled state. All along the chassis's border are the protrusions used to glue it together – either crenelations for rigid 90° bends or so-called *glue flaps* for flat connections. The glue flaps are noteworthy because they each have one of their two PET layers cut away so that they can mesh with their mate glue flaps without doubling the link thickness at the connection points. (To be clear, the gluing here is done manually during final assembly, not with a laminator as when processing the SCM layers.) Note that four of the glue flaps are attached to the tent flaps, showing that the tent flaps are fully integral objects of the SCM process; they can be enclosed by other SCM objects without any process issues. The tent flaps made this way, though, are fragile (because they only have one layer) and sticky (because they are covered in thermal adhesive), so in practice another step is added to the process wherein they are reinforced with two more sheets of nylon. This new step is introduced just before the final lamination step such that the reinforcements end up being made permanent by said final lamination step.

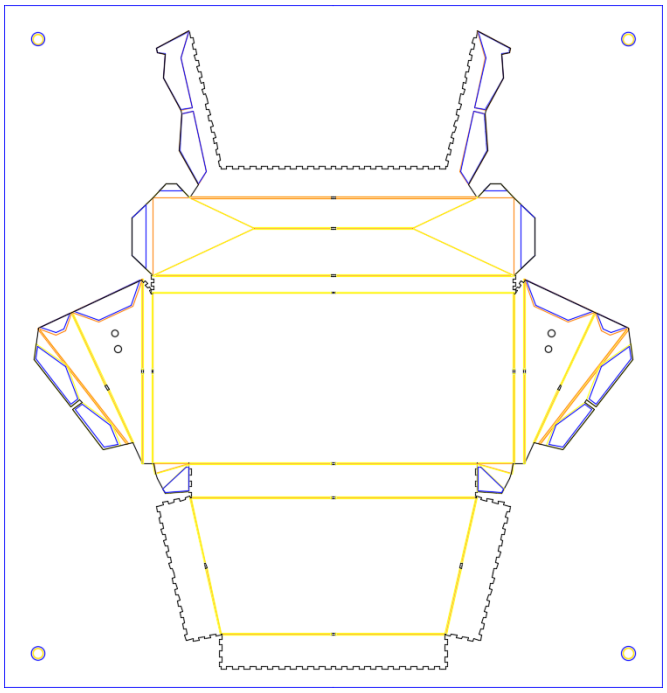

Figure 3: AutoCAD layout of the SCM pattern. Yellow lines indicate cuts through the bottom PET layer, orange lines indicate cuts through the top PET layer, blue lines indicate cuts through the nylon layer, and black lines indicate final cuts through all three layers. The four circles at each corner of the layout are for alignment purposes.

Most of the flexures in Figure 3 have different widths on the two different PET layers. This is done to give the flexures asymmetrical bending properties – to control their direction of bending. Those with closely spaced yellow-orange-yellow lines tend to bend inward; those with closely spaced orange-yellow-orange lines tend to bend outward. Another noteworthy fact about the flexures is that most of them have little breaks (represented by small black rectangles). These breaks were added to keep the linkages that were surrounded by flexures from falling out during processing. Note also that not all the flexures remain flexible after final assembly. Three of the flexures surrounding the bottom plate as well as the four flexures of the cavity cap end up being made rigid (manually) with hot melt adhesive.

III - Tail Design

Purpose of the Tail

For a two-wheeled rover whose wheels both move about a common axis, the main purpose of a tail is to prevent the rover's chassis from spinning about said axis when the rover tries to drive forward. The tail acts to counteract the reaction torques of the rover's motors; when the motors exert torques on their respective wheels to make the wheels spin in the forward direction, the reaction torques on the motors (and hence on the rover's chassis) tend to make the chassis spin in the opposite direction. This undesired backward spin ends up being neutralized if there is a tail to press into the ground, for the normal force on the tail cancels out the reaction torques on the chassis. The main conditions for the tail to be effective here are that it be long enough to extend past the wheel rims and that it be rigid enough to support the required force.

In other words, two-wheeled rovers almost by default satisfy the criterion of being able to flip themselves over. Only by adding a tail can one prevent the unwanted predicament where a rover does nothing *but* flip itself over. Fortunately having a tail does not make a rover lose its self-righting ability. To flip itself over, the rover just needs to be driven backward instead of forward. That way, all the torques change direction and the tail has to rotate all the way across the rover before it can counteract them. For such flip maneuvers to succeed, the combined moment due to tail and chassis weight must be small enough that the wheel motors can lift the tail off the ground – which in practice means that the tail can neither be too long nor too heavy.

Tailless Driving

A two-wheeled rover with *sprawl ability* (such as PUFFER) does not actually need a tail to move around on a flat surface. This is because when the wheels are no longer upright (i.e. $\rho < 90°$), they no longer rotate about the same axis, so there is no natural axis for the chassis to spin around. Depending on the compliance of the sprawl mechanism, the torque outputs of the wheel motors, the placement of the center of mass, and the wheel traction with respect to the ground (among other things), there is a critical sprawl angle ρ_c above which the rover continually flips itself over and below which it drives around on just two contact points. For the latest version of PUFFER, ρ_c was measured to be 65°.

This phenomenon of tailless driving, however, introduces too many problems to justify a tailless PUFFER design. First, it forces the rover to always drive around in a sprawled state – a constraint that clashes with the PUFFER design feature (see section IV) of using sprawl angle to switch between two traction mechanisms on the wheels. Second, it offers no protection against pitch-back moments whenever the rover goes up steep slopes. As measured on the latest version of PUFFER, for example, slope climbing with no tail is pitch-back limited at $\theta = 22°$, while slope climbing with the tail is traction limited at $\theta > 45°$.

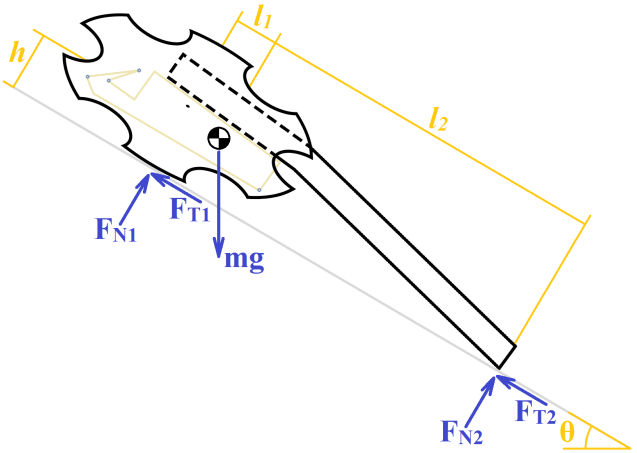

Figure 4: free body diagram of the rover on a slope. For simplicity, the chassis cross section is shown instead of a chassis side view. In this drawing the bottom plate has uniform clearance, but that is not true in general.

Optimal Sprawl Angle & Tail Mechanics

It being established that PUFFER requires a tail for adequate driving, it is appropriate to explain below how the tail parameters – placement, dimensions, and compliance – were chosen. First, the tail was placed on top of the chassis rather than on the bottom to leave the bottom plate smooth. This was done because of clearance issues; if the tail protruded from the bottom of the rover, it would tend to get the rover stuck on rocks. The disadvantage to having the tail placed this way is that the tail angle relative to the bottom plate is now a function of the sprawl angle rather than a constant. Because the tail tends to always touch the ground, this means in practice that the bottom plate is parallel to the ground at only one (or, with good design, two) sprawl angles. Hence, optimal ground clearance can be achieved at only one (or two) sprawl angles. Since slope climbing is the most challenging mode of operation for PUFFER, this optimal clearance was given to the sprawl angle best suited for climbing. Figure 4 shows the free body diagram used to derive this best sprawl angle.

For simplicity, the rover is oriented along the gradient of the angled surface. This way, symmetry reduces the number of unknowns from 9 to 4 and the number of equations from 6 to 3, making it possible to extract useful information out of the force balance without having to resort to modeling the internal stiffness of the rover. The two unknowns F_{N2} and F_{T2} are the normal and tangential forces on the tail, respectively, while the other two unknowns F_{N1} and F_{T1} are the combined normal and tangential forces on the wheels (lumped down from four to two unknowns due to symmetry). The distance variables l_1 and l_2 are the lateral lengths from the center of

mass to the wheel contact and to the tail contact, respectively. The distance variable h is the normal height of the center of mass off the ground. All three of these variables are complicated functions of the sprawl angle ρ, with h increasing strongly with ρ, l_2 decreasing weakly with ρ, and l_1 varying weakly in a way that depends on the geometry. The two variables independent of ρ are θ (the incline angle) and mg (the rover's weight).

Solving for the 4 unknowns as a function of the 5 variables gives the following result, where F_{T1} and F_{T2} cannot be decoupled because there are fewer equations than unknowns:

$$F_{N1} = mg\frac{l_2\cos(\theta) - h\sin(\theta)}{l_1 + l_2} \tag{1}$$

$$F_{N2} = mg\frac{l_1\cos(\theta) + h\sin(\theta)}{l_1 + l_2} \tag{2}$$

$$F_{T1} + F_{T2} = mg\sin(\theta) \tag{3}$$

In the worst-case scenario for traction, all the tangential load is borne by the wheels – i.e. $\mathrm{F}_{T2} = 0$ and $\mathrm{F}_{T1} = mg\sin(\theta)$. Assuming that the climbing angle θ is limited by traction (i.e. the tail is long enough to keep the rover from pitching back), the optimal sprawl angle is one that lets the wheels sustain the most tangential force without slipping. Because most traction materials have their tangential load-bearing ability increase with normal force, this criterion translates to maximizing F_{N1} as a function of rover geometry. To find out how F_{N1} varies with rover geometry, one must take the partial derivatives of equation (1) with respect to each of the three distance variables:

$$\frac{\partial F_{N1}}{\partial l_1} = mg\frac{h\sin(\theta) - l_2\cos(\theta)}{(l_1 + l_2)^2} \tag{4}$$

$$\frac{\partial F_{N1}}{\partial l_2} = mg\frac{h\sin(\theta) + l_1\cos(\theta)}{(l_1 + l_2)^2} \tag{5}$$

$$\frac{\partial F_{N1}}{\partial h} = -mg\frac{\sin(\theta)}{l_1 + l_2} \tag{6}$$

From equation (4) it is evident that decreasing l_1 (i.e. moving the center of mass toward the front of the rover) increases F_{N1} as long as $l_2\cos(\theta) > h\sin(\theta)$, which is true under the assumptions because this expression is the pitch-back inequality. From equation (5) it is evident that elongating l_2 (i.e. increasing the tail's reach) increases F_{N1} as long as $l_1 > 0$, which from a design perspective *must* be true to keep the rover from spontaneously tipping forward in its unsprawled state. From equation (6) it is evident that decreasing h (i.e. lowering the center of mass) increases F_{N1} as long as $l_1 + l_2 > 0$, which is true whenever the tail is oriented correctly. To summarize these results,

$$\frac{\partial F_{N1}}{\partial l_1} < 0 \tag{7}$$

$$\frac{\partial F_{N1}}{\partial l_2} > 0 \tag{8}$$

$$\frac{\partial F_{N1}}{\partial h} < 0 \tag{9}$$

under normal rover operating conditions. Although these results are useful for optimizing things like tail length during a rover's *design* phase, below they are used to see how changing the sprawl angle on a *given* rover perturbs l_1, l_2, and h (and hence varies F_{N1}).

As mentioned earlier, h increases strongly with sprawl angle ρ. This should be evident from Figure 2, as not only does the chassis itself expand with increasing ρ, but the unfolding wheels push the whole chassis off the ground. Hence,

$$\frac{\partial h}{\partial \rho} >> 0 \tag{10}$$

Similarly, one can use a side view of the rover (Figure 4) to deduce that l_2 decreases weakly with sprawl angle. As ρ increases, the tail rotates inward, but because its total length must remain constant, its lateral length must decrease. Hence,

$$\frac{\partial l_2}{\partial \rho} < 0 \tag{11}$$

The dependence of l_1 on the sprawl angle is the most complicated relationship of the three. Figure 4 is the best visual aid to refer to in the following discussion. As the tail rotates inward with increasing ρ, its contribution tends to move the center of mass toward the front of the rover. However, the tail represents very little of the rover's mass; most of the mass is attached to the lower half of the chassis. The kinematics is such that the lower half of the chassis rotates (i.e. pitches either forward or backward) in the opposite direction as the tail with increasing ρ – *as long as the tail is longer than a critical length*, which in practice it is (more on this later). Whether or not this rotation moves the center of mass toward the front or the back of the rover depends on the angle of the chassis's bottom plate relative to the ground. If the angle is positive (i.e. the ground clearance decreases from the front to the back of the chassis), the center of mass moves toward the back of the rover. If the angle is negative, the center of mass moves toward the front. The latter is the worst-case scenario for this analysis, so for now let

$$\frac{\partial l_1}{\partial \rho} < 0 \quad \text{(worst case)} \tag{12}$$

As explained earlier, the purpose of this analysis is to find the sprawl angle ρ that maximizes F_{N1} for a given rover. From the rules of continuous optimization, the ρ that maximizes F_{N1} is either at the upper end of its range ($\rho = 90°$), at the minimum ρ such that the conditions of the problem are satisfied (the lowest ρ at which only the wheels and the tail touch the ground), or is such that $\partial F_{N1}/\partial\rho = 0$. The full expression for $\partial F_{N1}/\partial\rho$ is written below:

$$\frac{\partial F_{N1}}{\partial \rho} = \frac{\partial F_{N1}}{\partial l_1}\frac{\partial l_1}{\partial \rho} + \frac{\partial F_{N1}}{\partial l_2}\frac{\partial l_2}{\partial \rho} + \frac{\partial F_{N1}}{\partial h}\frac{\partial h}{\partial \rho} \tag{13}$$

All six partials on the right side of this expression have been discussed in the previous paragraphs, with three of them known outright [equations (4) - (6)] and the other three not explicitly calculated but with information about their signs. Substituting in the signs (with the worst case for $\partial l_1/\partial\rho$) yields the following ambiguous result:

$$\frac{\partial F_{N1}}{\partial \rho} = (-)(-) \quad + \quad (+)(-) \quad + \quad (-)(+) \tag{14}$$

The term representing lateral center of mass movement tends to make the expression go positive, while the two terms representing tail movement and normal center of mass movement tend to make the expression go negative. Because of the conflicting contributions, no conclusions can be drawn from this result.

Fortunately one can get a better limit on $\partial l_1/\partial\rho$ in place of inequality (12). Due to the reduced changes in chassis pitch compared to the tail angle and due to the lower lengths involved, one can say for reasonable designs[3] that if $\partial l_1/\partial\rho < 0$, then $|\partial l_1/\partial\rho| < (l_1/(l_1+l_2))|\partial l_2/\partial\rho|$. In the other case where $\partial l_1/\partial\rho > 0$, one can of course say that $\partial l_1/\partial\rho > \partial l_2/\partial\rho$ because $\partial l_2/\partial\rho < 0$. It is easy to collapse these two cases into the following expression:

$$\frac{\partial l_1}{\partial \rho} > \frac{l_1}{l_1+l_2}\frac{\partial l_2}{\partial \rho} \tag{15}$$

Now if one substitutes equations (4) - (6) and inequality (15) into equation (13), one gets the following form for the optimization gradient $\partial F_{N1}/\partial\rho$:

$$\frac{\partial F_{N1}}{\partial \rho} < mg\left(\frac{2h\sin(\theta)}{(l_1+l_2)^2} + \frac{l_1^2\cos(\theta)}{(l_1+l_2)^3}\right)\frac{\partial l_2}{\partial \rho} + mg\left(\frac{-\sin(\theta)}{l_1+l_2}\right)\frac{\partial h}{\partial \rho} \tag{16}$$

The coefficient of $\partial l_2/\partial\rho$ is positive for all θ, while the coefficient of $\partial h/\partial\rho$ is negative for all θ. The sign form of the full expression is now:

$$\frac{\partial F_{N1}}{\partial \rho} \quad < \quad (+)(-) \quad + \quad (-)(+) \tag{17}$$

Now the right side of the expression is negative for all ρ and θ in their respective allowed ranges for this problem. The substitution of inequality (15) into the $\partial F_{N1}/\partial\rho$ expression shows that the effect of lateral center of mass movement is overshadowed by the two other terms. Because $\partial F_{N1}/\partial\rho < 0$ for every ρ in its permissible range, the function $\mathrm{F}_{N1}(\rho)$ has no local maximum and gets larger as ρ decreases. Hence, even with the complications due to l_1 and l_2 perturbations, equation (16) shows that the best sprawl angle for traction is *the minimum ρ at which only the wheels and the tail touch the ground.*

Depending on the wheel design, this optimal angle is somewhere between 20° and 25°. (Setting ρ below this angle lifts the wheels off the ground, causing F_{N1} to drop rapidly to 0.) Unfortunately this optimal ρ is not suitable for being the default sprawl angle used for climbing, as the ground clearance of the bottom plate with this ρ is by definition zero. In practice the rover would get stuck on even the smallest rocks if it had to climb in a way that maximized F_{N1}. From the point of view of tail design, then, the default sprawl angle used for climbing was set to 30° – to give the bottom plate a few millimeters of clearance:

Tail Dimensions

With this default sprawl angle set, tail design was no longer indeterminate. First, the tail was shaped in such a way that the bottom plate was parallel to the ground at $\rho = 30°$ – to maximize ground clearance at the default sprawl angle. This is why the tail is bent slightly downward, as shown in Figure 1. Second, the tail was designed to have almost no compliance in its up-down direction. This was done to preserve perfect clearance at all incline angles, as the normal force on the tail varies somewhat with θ.

According to equation (5), the tail should be made extremely long to maximize F_{N1}. In practice this is impractical, though, because with a long enough tail the rover would not be able to flip itself over. Fortunately lengthening the tail to optimize F_{N1} is subject to diminishing returns, meaning that near-maximal F_{N1} is possible even with a relatively short tail. With $l_1 = 5$mm and $h = 10$mm, for example, and $\theta = 45°$, using $l_2 = 19l_1$ gives 85% as high an F_{N1} as using $l_2 = \infty$. Hence, a secondary criterion is needed to optimize tail length. (The values listed in the example, incidentally, are the real values used for PUFFER.)

It was mentioned in the discussion of chassis rotation (as a function of ρ) that there is a critical length l_2 above which the lower half of the chassis pitches *forward* with increasing ρ. A tail with a shorter l_2 would cause the chassis to pitch *back* with increasing ρ – i.e. such that the clearance at the front increased more compared to that at the back. Because clearance is made to start uniform at the lowest functional ρ, such a design would be undesirable from a clearance point of view. When driving up a rough incline, the rover would allow rocks under its chassis that would tend to get stuck at the back of the chassis. In such a case the third contact point would be on the bottom plate instead of the tail tip. Hence, the effective l_2 of the rover would be reduced, limiting the traction output of the wheels. If the clearance at the back were greater than at the front, there would no longer be a problem because any rocks that would get stuck at the back would not make it past the front of the rover

[3]Designs where the bottom plate is parallel to the ground at some angle ρ where only the wheels and the tail touch the ground. Otherwise a design is not optimized for clearance at some best sprawl angle ρ, which is the whole point of this analysis. Also, the tail length must be of the same order of magnitude as the chassis dimensions for the inequality to be true.

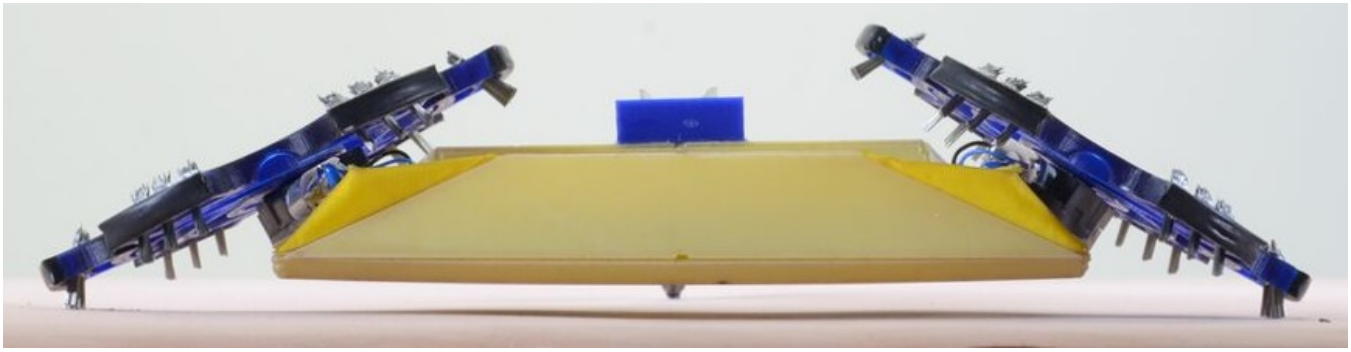

Figure 5: rover clearance at the default ρ for climbing.

anyway.

One can get a good estimate of the critical l_2 by taking a rover with a bent tail (designed to make the rover have uniform clearance at its default $\rho = 30°$), expanding the chassis such that $\rho = 90°$, and then lengthening the anchored part of the tail such that uniform clearance is again true. (Because the tail tip's position relative to the chassis has two degrees of freedom in the relevant plane, it is possible to satisfy the latter objective without sacrificing the former.) The l_2 found this way is special because with it the height change of the rotating tail is well matched to the height change of the expanding wheels/chassis. A complete match for all ρ is impossible because of nonlinearietes in the chassis linkage mechanism, so the l_2 found using this method is essentially the *optimal l_2 for maintaining close-to-uniform clearance at all ρ.*

The actual l_2 used for PUFFER is 95mm, which is 10% longer than the critical l_2 for the same chassis and wheels. The chosen l_2 was a compromise between maintaining near-uniform clearance at non-default ρ and maximizing F_{N1}. As explained earlier, having the actual l_2 be longer than the critical l_2 does not cause problems, as the chassis ends up rotating in the correct direction with increasing ρ. With $l_2 = 95$mm, the pitch-back limit for the entire rover (based on setting equation (1) equal to zero) is very high at $\theta = 84°$, so the assumption of traction-limited climbing does indeed hold true.

Sideways Tail Compliance

One design variable that has not been mentioned yet is the sideways compliance of the tail: compliance in and out of the page of Figure 4. The main benefit to having sideways compliance is that it lets the tail fold in such a way that it conforms to the back of the chassis, thereby reducing the rover dimensions during storage. Adding sideways compliance does not clash with any of the tail analysis done previously, as the free body diagram of Figure 4 assumes zero sideways force on the tail. In practice sideways tail compliance is actually good for climbing rough slopes, as the tail can better conform to surface protrusions. It is easy to independently vary the normal and sideways tail compliance by changing the height and width of the tail cross section. For PUFFER the cross-sectional tail height was set to 7mm to give the tail negligible normal compliance, while the tail width was set to consist of two 0.4mm thick (shaped) fiberglass sheets. The tail material was chosen to be fiberglass due to its good elastic range at the given stress levels.

Figure 6 (a) and (b) shows the tail in its collapsed and its expanded state, respectively. The noteworthy design feature here is that the tail is split into two halves to take advantage of symmetry. If the tail were made of only one piece and the robot were stored in a compact state for many months, creep deformation would cause the tail to be off-center when it popped back into its relaxed state. With two equal and opposite pieces set at an angle to the rover's bisection line, though, any creep deformation is canceled out during deployment as long as the deformation angle is lower than the placement angle (see tail holder).

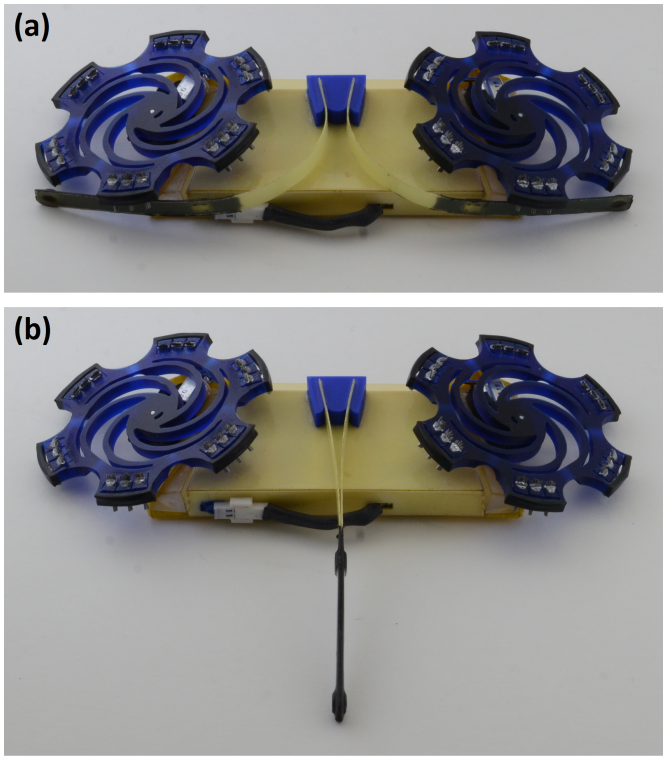

Figure 6: rover with its tail (a) collapsed and (b) expanded.

Another advantage of the two-piece design is that it allows the tail to have more than double the torsional stiffness. With only one narrow beam serving as the tail, the tail tends to undergo torsional buckling when its tip is in contact with the ground. With two narrow beams that are separated at their base by a gap (see tail holder), however, buckling is no longer an issue because the torsional stiffness is augmented by the torsional version of the I-beam effect. The beams must not be allowed to slide relative to each other after deployment for this effect to occur. This criterion was fulfilled by gluing magnets on the end of each tail piece. As soon as the two pieces spring out during deployment, they stick together and stay that way.

When the tail is in its *collapsed* state, each half of the tail is held in place by the wire clumps on the wheels. It is worth noting that the rover can deploy the tail from this state just by actuating its servo and motors – without the need for an extra actuator or any outside help.

Tail Traction

The current tail is not fully optimized for surface traction. Its friction coating consists of two layers of tool-dip rubber (which are cut along the line of symmetry to enable the tail to collapse). One hypothesis is that putting spines or brushes on the tail will increase the climbing performance of the rover. The next section discusses such traction mechanisms in more depth, though in relation to the wheels and not the tail.

IV - Wheel Design

Impact Absorption

The main purpose of the wheels is to let the rover get over obstacles and go up steep slopes. Hence, the main design requirement for the wheels is that they have good traction with the ground. A less important purpose of the wheels is to absorb impacts (to some degree) when the rover falls from large heights. It was this second purpose that dictated the bulk material of the wheels, as traction is a surface property and can be optimized with surface modifications at the points of ground contact.

The material chosen for the wheels was high impact strength polyurethane[4] because of its high strain limit and low density. Each wheel was made with a three-step casting process. First, a negative of a mold was 3-D printed out of PLA. Second, silicone rubber was poured into this negative mold and left to cure to get the actual mold. Third, polyurethane was poured into the silicone mold and left to cure *under pressure* (to minimize the size of the bubbles). Figure 7 illustrates the three steps of the casting process.

Many wheel designs were tested on PUFFER, most differing only in their manner of traction. One trait that all designs have in common is that they have a set of curved spokes connecting the hub of the wheel to the rim. These spokes were designed for impact absorption; they have enough stiffness to not droop much during normal rover motion, yet enough compliance to absorb a fraction of the rover's kinetic energy during impacts. Admittedly this impact absorption strategy is not ideal, since the impact forces (though mitigated) still pass through the fragile gears of the motor assembly. The motor assemblies end up being the limiting factor when it comes to the maximum survivable drop height (= 0.4m for the whole rover).

A much better design is to have each wheel revolve about a large cylindrical hub in a way that removes the gears and motors from the path of impact force transmission. The premier version of PUFFER, built by JPL, is set to use such a design (Karras et al. 2017). Field tests have shown that JPL's PUFFER (which uses a similar chassis as ours) can survive falls from more than 2m on Earth, thereby affirming the chassis's effectiveness at impact absorption. Cylindrical hubs were not used in the UC Berkeley prototype because they would have required custom motor assemblies. Nonetheless, the impact damping portion of the wheels is decoupled from their traction system, so all further discussion of the wheels could apply equally well to either hub design.

Wheel Diameter

Another trait that all the wheel designs have in common is that their outside diameter is very close to 7cm. This value came about as the lower limit to a crucial design objective: the wheels had to be big enough to enclose the whole chassis between them (with no protrusions as seen on a side view), even at low sprawl angles like $\rho = 30°$. Otherwise the front (or back, if driving backward) edge of the chassis would tend to get stuck on step-shaped obstacles. The wheel diameter was not made *larger* than this value because of storage size constraints and motor torque limits.

Testing Surfaces

Each of the many wheel designs was tested for performance on two specially-crafted gypsum surfaces: a rough surface with macroscopic features on the order of 10mm, and a smooth surface with no macroscopic features. Both surfaces had asperities on the order of 100μm in size. The

Figure 7: from left to right: 3-D printed negative of the mold; silicone mold itself; polyurethane part of the wheel.

[4] TP-4014 from Innovative Polymers, Inc.

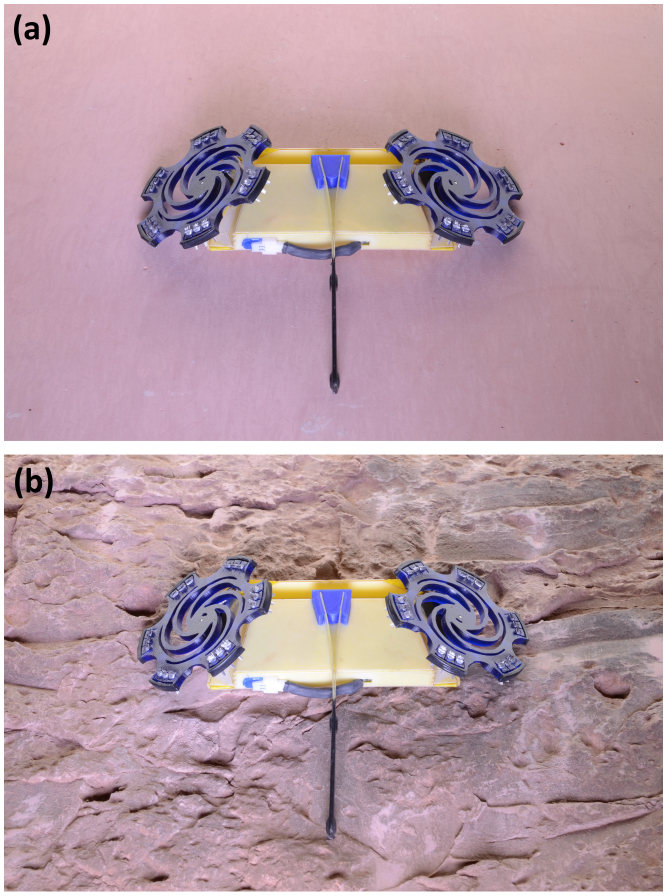

Figure 8: PUFFER on (a) the smooth surface and (b) the rough surface.

rough surface was made with rock molds taken from natural rock faces, while the smooth surface was made by sanding down a solidified puddle of gypsum. Figure 8 shows PUFFER perching on both surfaces (at an incline angle of 45°).

Testing was done for each set of wheels on each surface to determine the *cutoff angle* – the maximum incline angle up which PUFFER could reliably climb with said wheels. Cutoff angles are defined as θ_r for the rough surface and θ_s for the smooth surface. For a certain climbing trial to be judged as a success, the rover had to make it up 50cm along the length of the surface without once tumbling to the bottom. In all cases the rover was controlled manually by the same pilot. A climbing trial at the cutoff angle would tend to take several minutes because of multiple short-distance slips down the incline. Hence, the cutoff angle should be thought of as the incline angle at which the average upward rover velocity tends to zero.

Cutoff angles (as opposed to some other angle metric) were used to compare the different wheels because it was possible to measure the cutoff angles with great precision. It was the several-minutes-per-trial aspect that made each measurement so precise: if it took a rover several minutes to get up a short distance, one could be sure that the rover's average upward velocity was very close to zero. In other words, precision was proportional to trial time. Testing thus boiled down to finding the incline angle at which a successful trial would take the longest time possible. Adding even a fraction of a degree to that angle would flip the sign of the rover's average upward velocity, meaning that elevation gains would no longer counterbalance slipping losses. There would then be a clear transition in the average direction of rover motion.

Another advantage of the long trial times was that they tended to average out any irregularities due to manual control. [Automatic control would have made testing more repeatable, true, but programming the controller in a way that would give as good a performance as manual control would have been extremely difficult. In practice one must use a set of complicated maneuvers such as jigging to get the rover up the rough surface at high incline angles. Manual control was thus the only viable option to measure the true climbing ability of PUFFER.]

Exploratory Designs

Of the many wheel designs tested on PUFFER, ten are worth discussing in this paper: five exploratory designs to determine a good general configuration for the wheels, and five advanced designs to empirically optimize that configuration. Figure 9 shows the inner side (i.e. the side touching the ground) of each of the exploratory designs, while Table 1 compares their θ_r and θ_s performance. All wheels are made of the same material; they are just dyed differently.

Table 1: exploratory designs.

design	up rough incline	up smooth incline
#1	41°	32°
#2	38°	34°
#3	30°	27°
#4	40°	41°
#5	43°	40°

Design #1. The first design tested has six curved protrusions whose extremities are coated in tool-dip rubber – in other words, an impact absorption design described earlier but without the rim. Although the design has decent performance up the rough incline ($\theta_r = 41°$), it has the problem that its high aspect ratio protrusions tend to get stuck in nooks and crannies on the rough surface. Each time the rover gets stuck, it has to be driven backward to get it unstuck, leading to valuable losses in elevation. The logical conclusion here is that any wheel protrusions should have either low aspect ratios or low stiffness.

Design #2. The wheels here have a circular rim with six circular cutouts, yielding essentially six low aspect ratio protrusions. The purpose of the protrusions is to let the rover get over step-shaped obstacles whose height is more than the wheel radius (the rover's sprawl angle being 90° in this case). Each cutout is shaped like a halfcircle because that is the deepest shape that is geometrically

Figure 9: from left to right, designs #1 to #5.

unable to get stuck as long as only one cutout touches the ground. As with design #1, each protrusion is coated in two layers of tool-dip rubber. Performance on the rough surface is worse than with design #1, but with the qualitative advantage that the wheels no longer get stuck. To get good performance, one must evidently combine this design with a better traction mechanism. Note: design #2 was partly inspired by the crenelations of Stoeter et al. (2003).

Design #3. Identical in all relevant regards to design #2, but without the rubber coating on the six protrusions. This is basically the "control" design that shows how bad performance gets without any traction optimization. The traction surface here is made of polyurethane, the same as the bulk of the wheels.

Design #4. Identical to design #2, but with a pin attached at an angle to each protrusion. The pins are oriented relative to each wheel such that they dig into the ground as the wheels spin forward. Performance is quantitatively better than with design #2, but again there is the qualitative problem that the wheels tend to get stuck because of high aspect ratio protrusions (this time the pins). To take full advantage of such pin designs, one must find a way to keep their traction benefits while discarding their propensity to get stuck.

Design #5. Identical to design #2, but with a metal brush attached to each protrusion. The brushes are oriented such that their planes of contact are parallel to the ground when the rover is at its default sprawl angle of $\rho = 30°$. In other words, each brush is angled 30° away from the axis of the wheel. Each brush has hundreds of steel wires with a diameter of 80μm that are arranged in a circular cluster with nearly 100% (that is, hexagonal) packing. Designed for use with power tools, the brushes themselves are durable; but individual wires tend to plastically deform if loaded disproportionately.

As evident from Table 1, the brushed design has the best quantitative performance of the five exploratory designs. It also does not suffer from the qualitative problem of getting stuck in macroscopic concavities. The reason brushes tend not to get stuck in concavities is that the high aspect ratio structures (the wires) are clumped; individual wires are kept from getting lodged in concavities by the wires around them.

The brushes used in design #5 have the benefit that their wires are relatively compliant, but that does not need to be true for brushes in general to avoid getting stuck. Compliance is crucial, though, to promote *load sharing*, which is the main reason brushes are good at traction. Load sharing allows multiple wires on the same brush to engage with surface asperities, thereby increasing the force needed to cause slip.

Due to these benefits of brushes, it was decided to ignore all the other traction paradigms and to focus all design efforts on making and optimizing custom brushes. This decision, though, should not be taken as a repudiation of all these other traction paradigms. Individual spines, for example, have been shown to work well with anchor mechanisms that introduce compliance (Daltorio et al. 2009; Asbeck and Cutkosky 2012; Carpenter et al. 2015). Rather, this was an attempt to investigate only minimalist designs and to reduce the design space to manageable levels.

Sprawl-Switchable Traction

Design #5 illustrates one of the novel paradigms of PUFFER: sprawl-switchable traction. At low sprawl angles (such as $\rho = 30°$), only the metal brushes touch the ground. At high sprawl angles ($\rho \Rightarrow 90°$), only the rubber-dipped crenelations touch the ground. By changing its sprawl angle, the rover can switch between two different modes of traction without having to change its wheels. These two modes of traction correspond well to their respective sprawl angles: metal brushes are good for incline climbing, which is done best at low sprawl angles; crenelations are good for getting over obstacles on otherwise flat terrain, which is done best with upright wheels. (The maximum surmountable obstacle height for typical wheels is $r{\cdot}\sin(\rho)$, where r is the radius of the wheels. Non-circular wheels can exceed this limit, however; see section V.)

Note that a previous wheeled robot (STAR) has demonstrated the paradigm of changing its sprawl angle to optimize its shape for a given terrain (Zarrouk et al. 2013). This robot, however, does not have a variable traction surface as a function of ρ. Hence, the robot does not take full advantage of changes in its sprawl angle. Other robots have demonstrated the ability to change their own modes of traction, but either passively (Daltorio et al. 2009; Lee et al. 2016) or at the cost of one or more actuators per leg (Spenko et al. 2008; Herbert et al. 2015). The former is not ideal in general because the traction surfaces tend to interfere with each other, while the latter is un-

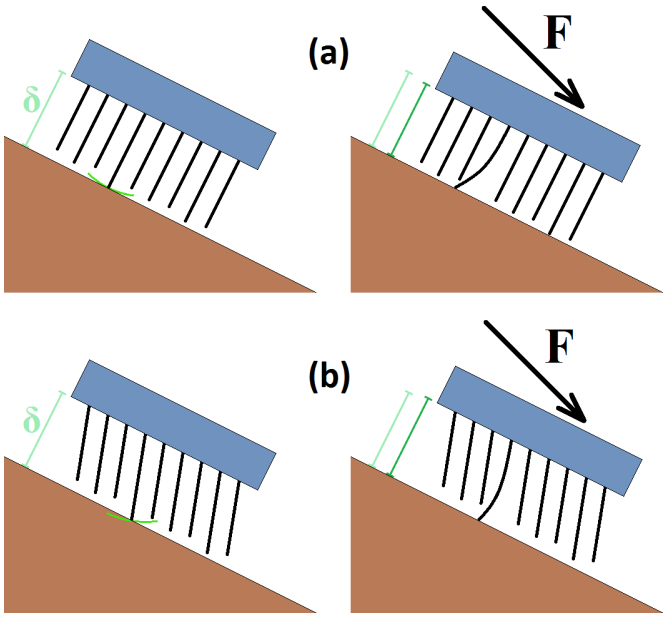

Figure 10: brush mechanics on inclines.

suitable for small robots because of size and complexity constraints. To the author's knowledge, PUFFER is the only robot that can actively change its mode of traction without having to use extra actuators.

Besides improving the rover's versatility, sprawl-switchable traction also helps with rover durability. In practice the metal brushes degrade with use (especially if the individual wires start out sharpened). By not having to use the brushes on even terrain, the rover is able to conserve them until they are needed. This is similar to how cats and other felines are able to retract their claws to limit unnecessary wear.

Brush Angles

The brushes of design #5 are oriented such that each wire that contacts the ground is perpendicular to the ground when the rover is at its default sprawl angle of $\rho = 30°$ (and when the ground is smooth). This geometry is illustrated back in Figure 5, though with a more advanced brush design. One might think that adding a tangential slant (as with design #4) to the brushes would help with traction, as each wire would tend to dig into the ground more. A tangential slant would actually hurt performance (even if the brush surface were planarized relative to the ground), as it would discourage load sharing. Figure 10 illustrates the problem with adding a tangential slant.

In practice the wires on each brush always have an imperfect length distribution; inevitably some wires are longer than others. As a given brush gets closer to the ground, the wire that contacts the ground first tends to be the longest one. In Figure 10, the longest wire in each case is the fourth one from the left. In case (a) the wires are all normal to the surface, while in case (b) the wires all have a slight slant. The gap between the brush base and the ground is marked as δ. Also marked (with a green arc) is the deformation range of the tip of the longest wire relative to its own base. Let the brush on each left be unloaded and the brush on each right be loaded with a force such as those seen by the wheels. Let each brush be constrained such that its base remains parallel to the ground, as would be approximately the case if it were attached to a wheel on PUFFER. Let the longest wire in each case have its tip fixed to an asperity in the ground.

In case (a), the only way the longest wire can bend is in a way that decreases δ. Hence, the shorter wires are able to get a foothold on the ground and provide load sharing. (The same bending effect happens with the second longest wire, and so forth.) If on the other hand there is a nonzero slant to the wires, as in case (b), the longest wire now tends to bend in a way that *increases* δ. The shorter wires now lift *away* from the ground, meaning that all the load ends up being carried by just one wire/asperity pair. In practice the brush ends up slipping because the wire bends too far back and causes the asperity contact to fail.

This phenomenon was verified experimentally with a variant of design #5 that had tangentially angled brushes. (The brushes were planarized such that the contact surface of each brush touching the ground was parallel to the ground, as in Figure 10b.) These brushes had poor performance ($\theta_s < 30°$) compared to their tangentially unangled counterparts ($\theta_s = 40°$).

Advanced Designs

With the angles of the brushes decided upon (oriented 30° away from the wheel's axis, with no tangential slant), there were still many design parameters that had to be optimized to make ideal brushes. The most important parameters were wire material, wire length, wire radius, wire spacing, and the number of wires per brush. Another relevant parameter was the number of brushes per wheel. It was decided to maximize this last parameter over all the available wheel space, as there was no good reason not to do so.

The wire material was chosen to be superelastic nitinol due to its good resistance to plastic deformation and its high surface hardness (in the league of tool steel). The former trait would protect the wires from splaying; the latter from abrasion. Aluminum ferrules were used to clamp low diameter wires into clumps in cases where close packing of the wires was desired. Here, aluminum was chosen due to its low weight.

Wire length was chosen to be maximized with the constraint that the wires could not scrape the chassis except at very low sprawl angles. This optimization gave a wire length of 5mm (excluding the root of each wire embedded in the base of its brush). Long wires were seen as a benefit because each brush would need more wires of the same radius to achieve the same bulk stiffness. Each

Table 2: advanced (nitinol) designs.

design	wire diameter	wires per clump	packing	up rough incline	up smooth incline
#6	16mil (0.41mm)	1	N/A	47°	42°
#7	9mil (0.23mm)	1	N/A	43°	45°
#8	9mil (0.23mm)	14	50%	46°	45°
#9	9mil (0.23mm)	14	50%	48°	47°
#10	6mil (0.15mm)	50	70%	47°	42°

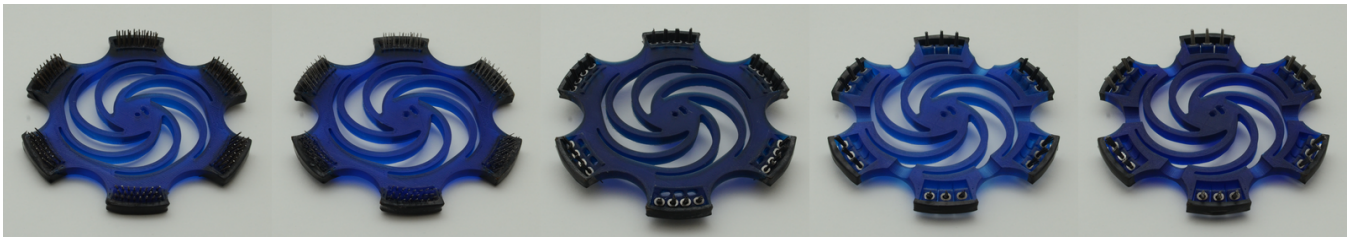

Figure 11: From left to right, designs #6 to #10

brush would thus have more surface contact points (and hence better engagement with asperities) without any undesirable trade-offs.

The remaining three parameters (wire radius, wire packing, and wire number per clump) were optimized empirically. Simulations were not used because there were too many unknowns with regard to wire contact with asperities to adequately guide an effective optimization. Figure 11 shows each of the advanced brush designs, while Table 2 gives their specifications and compares their performance on the two surfaces.

Both designs #6 and #7 use arrays of individual wires instead of having the wires clumped in ferrules; the only difference between designs #6 and #7 is the wire diameter. The lower diameter wires (design #7) perform worse on the rough surface (deemed the more important surface) because they tend to bend back too far to hold adequate asperity contacts. This is less of a problem on the smooth surface because load sharing is more effective with no surface roughness. The higher diameter wires (design #6) have the reverse properties; they perform *worse* on the smooth surface because their wires are too stiff for load sharing.

Design #8 uses wires of the same diameter as those of design #7, but the wires are clumped in ferrules such that each clump has the same bending stiffness as the individual wires from design #6. (The bending stiffness of wires is proportional to the fourth power of the diameter, so one would expect to need $(\frac{16}{9})^4 = 10$ wires per clump to balance the stiffness. It was discovered that not filling the ferrules to their full capacity [28 wires in this case] increases the effective length of the wires, however, so in practice 14 wires are needed in each ferrule to match the bending stiffness.) Based on Table 2, design #8 evidently combines the smooth surface performance of design #7 with the rough surface performance of design #6, thereby taking the best of both worlds. This shows the main advantage of having clumped wires: one can get the benefit of multiple asperity contacts without the disadvantage of low effective stiffness. In other words, the "packing" parameter adds another design degree of freedom to let one optimize the traction properties of the wheels.

Design #9 uses the same brushes as design #8, but this time the polyurethane wheel itself was made more compliant. It was hypothesized that if one puts each brush on its own cantilever beam and also adds compliance between the six protrusions holding the brushes, one would get more traction because more of the brushes would be able to contact the ground simultaneously. The hypothesis was inspired by the multi-level compliance setup of gecko toes (Autumn et al. 2000). Unfortunately not much compliance could be added without degrading the structural integrity of the wheel; the compliance at the wheel rim (in the axial direction) ended up being only 1mm under half the weight of the rover. Nonetheless, the results of Table 2 show a clear improvement over the performance of design #8. It is likely that adding even more compliance (by some more effective means like having a metal ring embedded in the wheel) could improve performance further.

Design #10 uses the same compliant polyurethane structure as design #9, but with brushes whose wires have even lower diameter. Again the wire number was tuned such that the bulk bending compliance of each brush was the same as before. The individual wires were also not specially cut; whereas the wires of designs #6 to #9 were cut to make their tips sharp, the wires of design #10 were cut to have flat ends. (Figure 12 compares the brush ends of designs #9 and #10 on both an angled and an unangled surface.) According to the results, design #10 performs worse than design #9 at climbing. However, design #10 was optimized for durability in addition to climbing. Durability is discussed in the next section.

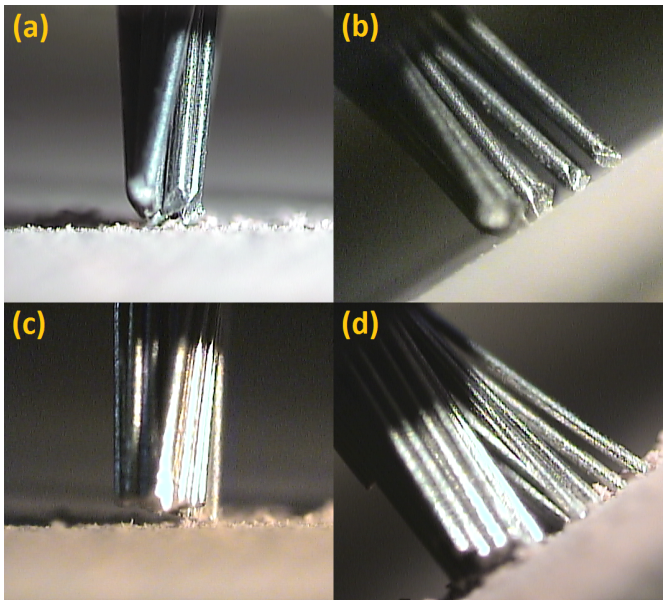

Figure 12: Microscope pictures of the brushes on an unsloped and on a sloped surface. (a) and (b) are from design #9, (c) and (d) are from design #10.

V - Performance

Climbing Characteristics

One surprising fact about Tables 1 and 2 is that almost all the wheel designs perform better on the rough surface than on the smooth. This can be explained by tail action. If the rover momentarily has bad traction on the *smooth* surface, there is nothing macroscopic that prevents it from sliding downward. On the rough surface, though, there is always some protrusion that the tail can end up catching on. In such an event the indeterminacy of equation (3) can resolve such that the tail carries *all* the tangential load. It thus works out that the surface protrusions limit elevation losses and give the rover second chances to get up inclines.

Notwithstanding, at most angles below the cutoff the rover is able to go up smooth inclines *faster* than up rough inclines. This is illustrated in Figure 13, which shows the velocity vs. incline angle plot for design #10 on both inclines. (The shapes of these plots are very similar for the various wheel designs; the plots are just stretched relative to each other based on the cutoff angles.) For each data point, the rover was tested five times to see how fast it would get up 50cm along the length of an incline. The error bars for each data point represent the sample standard deviation of these five trials. (The variability comes about in large part due to manual control.) Note that variability is low near the cutoff angle – consistent with the discussion in section IV. At low incline angles the velocity on the smooth surface tends to be higher than that on the rough surface because the rover does not need to relegate its kinetic energy toward overcoming obstacles.

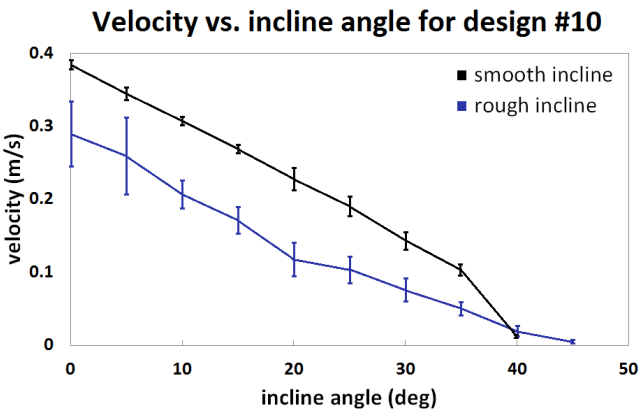

Figure 13: climbing performance for design #10. In all trials, the rover's sprawl angle was close to its default value of 30°. The pilot allowed it to vary slightly on rough terrain, though, whenever clearance was an issue.

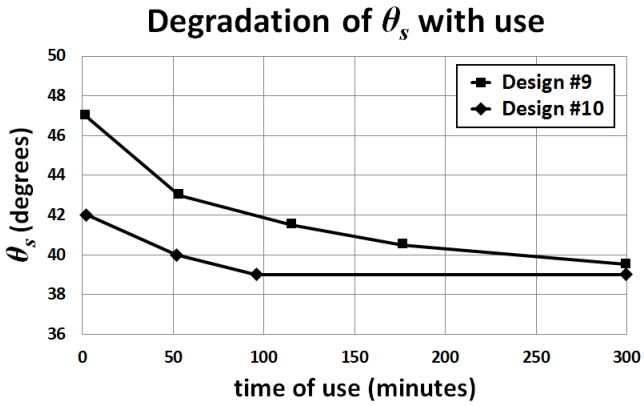

Figure 14: Degradation of designs #9 and #10 on the smooth gypsum surface. Uncertainty for each data point is ±0.5°.

Brush Durability

The performance results listed in Table 2 are only valid when the brushes are brand new. Unfortunately brush performance degrades with use due to abrasion. Degradation is most pronounced on rock surfaces, as the rocks act as grinding blocks to dull the individual wires. Figure 14 shows a plot of degradation over time of designs #9 and #10 on the smooth gypsum surface. In both cases the rover was run near the cutoff angle at full throttle, which is the worst-case scenario for degradation. (That way the brushes tend to scratch the surface instead of gripping it, thereby hastening abrasive action). The experiments were done on the smooth gypsum surface instead of the rough because it was easy to periodically sand it down to avoid conflating brush abrasion with surface abrasion. The degradation in cutoff angle on the rough surface turned out to be similar, though, over the same 5 hour time span: 48° ⇒ 44° for design #9 and 47° ⇒ 44° for design #10.

Figure 14 suggests that degradation in cutoff angle *plateaus* after a few hours of use. This observation

is consistent with spine theory; spines with lower tip radii have better surface adherence (because there are more available asperities (Asbeck et al. 2006)), and abrasion cannot degrade a spine beyond a maximum tip radius (which is equal to the radius of the spine itself). Design #9 approaches the plateau slower than design #10 because its wires a) started out sharp and b) have a larger stalk radius – both factors that increase the abrasive action needed to get the wires to their final tip shape. These factors are also the reasons why design #9 has more net degradation in cutoff angle: 8° compared to 3°. It is difficult to separate the effects of these two factors using current data; nonetheless, it seems evident that thinner wires (when clumped) have more resistance to degradation than thicker ones. Extremely thin wires, for example, are expected to have zero degradation in cutoff angle with use – since any abrasion will leave the wires with the same extreme sharpness.

Terrain Versatility

So far the discussion of performance has been limited to gypsum inclines (which represent natural rock slopes). Table 3 gives climbing performance data for a more varied set of surfaces. Only wheel design #10 was tested on these surfaces, as it was considered the "final" design because of its superior durability. The surfaces are arranged in order of increasing brush performance (i.e. increasing cutoff angle at $\rho = 30°$). It is evident from the table that brush performance correlates well with some combination of surface roughness and asperity size. On polished glass (the only surface with neither surface roughness nor asperities), brush performance is abysmal – a cutoff angle of just 11°. On tufted carpet, on the other hand, the cutoff angle is a relatively impressive 64° – helped in part by the carpet loops which act as asperities. Note that this trend is only expected to hold while the surface disturbances are smaller than the dimensions of the rover. With larger disturbances the surface can be thought of as having its θ vary with location.

When the wheels are upright ($\rho = 90°$), performance varies little across the eight different surfaces – from a low of 26° on polystyrene to a high of 33° on carpet. This lack of variability comes about because unlike brushes, rubber does not rely on asperities to maintain contact. The cutoff angle with upright wheels has a lower limit across the eight different surfaces than with sprawled wheels (33° vs. 64°) – mostly because in the latter case the center of mass is closer to the ground. Nonetheless, on the three smoothest surfaces the upright wheels do perform better. These results are a good example of the benefits of sprawl-switchable traction: on inclines where the primary mode of traction does not work well, the secondary can be deployed (and vice versa). In other words, because the rover can change its ρ on the fly, the effective cutoff angle for each surface is the *maximum* of the two values listed in Table 3.

Most climbing robots in the published literature are designed to have superb performance on a specific surface – without much regard for generalized terrains. There are at least three robots, for example, that can climb polished glass at $\geq 90°$ inclines whose traction systems get contaminated after just a few seconds of use on a dusty terrain (Daltorio et al. 2005; Kim et al. 2008; Murphy et al. 2011). PUFFER takes the alternate approach of being able to climb a wide variety of surfaces, though with only moderate performance on each given surface. This is why so much emphasis is placed on sprawl-switchable traction and on making the traction systems both rugged and durable.

Step & Overhang Height

When its wheels are upright, PUFFER can surmount steps of height up to 41mm, which is 117% of its wheel radius. Rovers with simple (non-crenelated) wheels, in comparison, have an upper step height limit of only 100% of their wheel radius (assuming there is no adhesion). There are rovers, however, that have reached step height values of 125% (Zarrouk et al. 2013) and even 150% (Morrey et al. 2003) of their respective wheel radii. These rovers use what are called "wheel-legs" – high aspect ratio protrusions similar to those of design #1 – to surmount the edges of steps. As discussed in section IV, such high aspect ratio protrusions have a tendency to get stuck in concavities on irregular terrain. Hence, there is an indirect trade-off between maximum climbable step height and maximum climbable incline angle on irregular terrain. The rovers mentioned in this paragraph represent different points on a hypothetical trade-off curve between these two performance metrics.

When right-side up and sprawled, PUFFER can crawl under overhangs as low as 28mm, which is 80% of its wheel radius. When *upside down* and sprawled, it can crawl under overhangs as low as 25mm – that is, 71% of its wheel radius. No other rover in the literature has better normalized overhang performance; the closest one is STAR (Zarrouk et al. 2013), with a corresponding value of 89%. In comparison, rovers with no sprawl ability can only crawl under overhangs that are higher than 200% of their wheel radius. There is a discrepancy between right-side up and upside down performance for PUFFER because when the wheels are right-side up, they need a larger ρ to touch the ground; larger ρ means that the wheels take up more vertical space.

VI - Conclusion

Summary

This paper presented the design of a simple, compact two-wheeled rover that can go up steep and rugged inclines while using a near-minimalist actuation strategy. The rover has three actuators: two for its wheels and one

Table 3: performance on different surfaces for design #10.

surface:	cutoff angle at $\rho = 30°$	cutoff angle at $\rho = 90°$	notes:
polished glass	11°	28°	no asperities above micron-scale
smooth plywood	20°	29°	much of surface polished to near-shine
smooth cardboard	26°	27°	surface features on the order of 50μm
smooth gypsum	42°	27°	asperities on the order of 0.1mm
rough gypsum	47°	28°	surface features on the order of 10mm
constrained pebbles	47°	30°	pebble diameter on the order of 10mm
floral polystyrene	49°	26°	surface roughness on the order of 1mm
tufted carpet	64°	33°	polypropylene loops of height 3mm

to change its shape. Shape-changeability is useful in large part because it allows the rover to fit in a constrained space for storage – required for the rover's applications on Mars. (The rover's collapsible tail design also helps with this criterion.) Changes in the rover's shape are actuated with a servo-driven chassis linkage that is made with SCM technology. Other functions of the linkage are that it protects the rover from dust and impacts and that it allows the rover's wheels to sprawl. The latter ability is crucial for climbing steep slopes, as the rover benefits from having a reduced center of mass.

The wheels' sprawl ability also allows the use of two different traction surfaces as a function of sprawl angle – a useful design feature known as sprawl-switchable traction. When the wheels are upright, traction is accomplished with rubber-coated[5] crenelations (which are good for overcoming obstacles and driving over smooth terrains). When the wheels are sprawled, traction is accomplished with nitinol brushes (which are good for climbing rocky, steep terrains). Brushes were chosen over other traction mechanisms like spines (Spenko et al. 2008) or gecko adhesives (Kim et al. 2008) because they have a good combination of ruggedness and simplicity.

Future Work

It must be noted that the wheel design was not fully optimized, as the optimization was done empirically and with only a limited set of wheels. The conclusions found in section IV – that wires should be clumped into brushes to allow stiffness regulation, that brushes should have wires normal to the ground to allow load sharing, that wheels should have some compliance to allow multiple brushes to touch the ground – are good rules of thumb for future designs, but are not enough to constrain all the design variables. For an optimal wheel design one would need to build a simulation that accurately models brush tip interactions with the ground. One would then need to iterate over all the design variables (on multiple surfaces) to find out which configuration gives the best performance.

[5]Note that rubber does not work well on Mars due to low temperatures, so a different material would have to be used for that application.

So far all brush designs have used *straight* nitinol wires, which is what brings about the "wires normal to the ground" criterion. One of the benefits of superelastic nitinol, though, is that it can be annealed to fit any shape necessary. If the wires were designed to have a *bend*, they would be able to have a more favorable contact angle relative to surface asperities – without having to sacrifice load sharing. In this case care would have to be taken to ensure that the wires have enough torsional stiffness not to twist out of the plane along which they are supposed to act. Another way to bypass the "normal to the ground" criterion would be to add axial compliance to the wires, as in (Wang et al. 2016).

A very likely (but not entirely proven) fact about the brushes is that sharper wires tend to increase the climbing cutoff angle. The caveat with sharpened wires, though, is that they end up getting dull after a few hours of use. One idea for future research is to give the rover the ability to sharpen these wires. Two small whetstones could be attached to the top of the chassis such that they would contact the brushes only when the sprawl angle was close to 0° (i.e. with the brushes not touching the ground). To sharpen the wires, the rover would then simply need to spin its wheels such that the brushes scraped the whetstones in the opposite direction as they scraped the ground. If this idea were a success, the rover would thus have the ability to self-service its brushes (like a cat with its claws) without the need for extra actuators.

Another idea for future research is to make the rover amphibious – for applications on Earth, of course, not on Mars. If the chassis were covered on the inside in a waterproof layer of thin plastic, the rover would be able to float. (The rover's sprawl angle would have to be high enough to give the rover a lower effective density than that of water; higher sprawl angle means more air in the chassis.) Assuming that all the other problems (hydrodynamic stability, nonwaterproof motors, etc.) are not too difficult to solve, the rover would then be able to navigate water – again without the need for extra actuators. (With proper modification, the wheels could serve as paddles.) The versatility of such a rover could far exceed that of all published mobile platforms (Nie et al. 2013).

References

Aksak B, Murphy MP, and Sitti M (2007) Adhesion of biologically inspired vertical and angled polymer microfiber arrays. *Langmuir* 23(6): 3322-3332.

Asbeck AT and Cutkosky MR (2012) Designing compliant spine mechanisms for climbing. *Journal of Mechanisms and Robotics* 4(3): 031007 1-8.

Asbeck AT, Kim S, Cutkosky MR, et al. (2006) Scaling hard vertical surfaces with compliant microspine arrays. *The International Journal of Robotics Research* 25(12): 1165-1179.

Autumn K, Liang YA, Hsieh ST, et al. (2000) Adhesive force of a single gecko foot-hair. *Nature* 405: 681-685.

Birkmeyer P, Gillies AG, and Fearing RS (2012) Dynamic climbing of near-vertical smooth surfaces. In: *IEEE international conference on intelligent robots and systems*, 7-12 October 2012, Vilamoura, Portugal: IEEE, pp. 286-292.

Birkmeyer P, Peterson K, and Fearing RS (2009) DASH: a dynamic 16g hexapedal robot. In: *IEEE international conference on intelligent robots and systems*, 11-15 October 2009, St. Louis, MO: IEEE, pp. 2683-2689.

Carpenter K, Wiltsie N, and Parness A (2015) Rotary microspine rough surface mobility. *IEEE Transactions on Mechatronics* PP-99: 1-16.

Casarez CS and Fearing RS (2016) Step climbing cooperation primitives for legged robots with a reversible connection. In: *IEEE international conference on robotics and automation*, 16-21 May 2016, Stockholm, Sweden: IEEE, pp. 3791-3798.

Daltorio KA, Horchler AD, Gorb S, et al. (2005) A small wall-waling robot with compliant, adhesive feet. In: *IEEE international conference on intelligent robots and systems*, 2-6 August 2005, Edmonton, Canada: IEEE, pp. 3648-3653.

Daltorio KA, Wei TE, Horchler AD, et al. (2009) Mini-whegs$^{\text{tm}}$ climbs steep surfaces using insect-inspired attachment mechanisms. *The International Journal of Robotics Research* 28(2): 285-302.

Herbert P, Bajracharya M, Ma J, et al. (2015) Mobile manipulation and mobility as manipulation – design and algorithms of RoboSimian. *Journal of Field Robotics* 32(2): 255-274.

Hoover AM and Fearing RS (2008) Fast scale prototyping for folded millirobots. In: *IEEE international conference on robotics and automation*, 19-23 May 2007, Pasadena, CA: IEEE, pp. 886-892.

Karras JT, Fuller CL, Carpenter KC, et al. (2017) Pop-up Mars rover with textile-enhanced rigid-flex PCB body. In: *IEEE international conference on robotics and automation*, 29 May - 3 June 2017, Singapore: IEEE.

Kim S, Spenko M, Trujillo S, et al. (2008) Smooth vertical surface climbing with directional adhesion. *IEEE Transactions on Robotics* 24(1): 65-74.

Komsuoglu H, Sohn K, Full RJ, et al. (2008) A physical model for dynamical arthropod running on level ground. In: *IFRR international symposium on experimental robotics*, 13-16 July 2008, Athens, Greece: IFRR, pp. 303-317.

Lee JS, Fearing RS, and Cho KJ (2016) Compound foot for increased millirobot jumping ability. In: *International conference on climbing and walking robots*, 12-14 September 2016, London, UK: World Scientific, pp. 71-78.

Morrey JM, Lambrecht B, Horchler AD, et al. (2003) Highly mobile and robust small quadruped robots. In: *IEEE international conference on intelligent robots and systems*, 27-31 October 2003, Las Vegas, NV: IEEE, pp. 82-87.

Murphy MP, Kute C, Mengüç Y, et al. (2011) Waalbot II: adhesion recovery and improved performance of a climbing robot using fibrillar adhesives. *The International Journal of Robotics Research* 30(1): 118-133.

Murphy MP and Sitti M (2007) Waalbot: an agile small-scale wall-climbing robot utilizing dry elastomer adhesives. *IEEE Transactions on Mechatronics* 12(3): 330-228.

Nie C, Corcho XP, and Spenko M (2013) Robots on the move: versatility and complexity in mobile robot locomotion. *IEEE Robotics and Automation Magazine* 20(4): 72-82.

Spenko MJ, Haynes GC, Sanders JA, et al. (2008) Biologically inspired climbing with a hexapedal robot. *Journal of Field Robotics* 25(4): 223-242.

Stoeter SA, Burt IT, and Papanikolopoulos N (2003) Scout robot motion model. In: *IEEE international conference on robotics and automation*, 14-19 September 2003, Taipei, Taiwan: IEEE, pp. 90-95.

Wang S, Jiang H, and Cutkosky MR (2016) A palm for a rock climbing robot based on dense arrays of micro-spines. In: *IEEE international conference on intelligent robots and systems*, 9-14 October 2016, Daejeon, Korea: IEEE, pp. 52-59.

Wood RJ, Avadhanula S, Sahai R, et al. (2008) Microbot design using fiber reinforced composites. *Journal of Mechanical Design* 130(5): 052304 1-11.

Zarrouk D, Pullin A, Kohut N, et al. (2013) STAR, a sprawl tuned autonomous robot. In: *IEEE international conference on robotics and automation*, 6-10 May 2013, Karlsruhe, Germany: IEEE, pp. 20-25.

Appendix: PUFFER Circuit Diagram