Early Stage Innovations (ESI) FINAL REPRT **NASA Grant Number NNX15AD74G-Agogino** 1/30/2019 Precision Hopping/Rolling Robotic Surface Probe Based on Tensegrity Structures University of California at Berkeley	
Principal Investigator	Alice M. Agogino
Co-PI	Adrian Agogino until Jan. 6, 2017
NASA Research Collaborator	Terry Fong NASA Ames Research Center

1. RESEARCH OBJECTIVE

Show that a 10 kg tensegrity ball probe can quickly and precisely deliver a 1 kg payload over 1 km distance on the Moon using a simple gas thruster. The probe is expected to be robust to many terrain conditions. The research is conducted primarily in simulation, but hardware is used to test and validate structural integrity concepts.

2. MOST SIGNIFICANT TECHNICAL ACHIEVEMENT(S)

We completed all proposed tasks and more. We evaluated a range of thruster concepts, evaluated the design trade-off space and expanded simulations to include thruster-based control and mobility. We simulated hopping profiles on both smooth and hilly terrain, and have expanded the visualization capabilities of the NTRT (NASA Tensegrity Robotics Toolkit). We completed the development and testing of five prototypes. A small-scale version called TT-4_{mini} was developed to increase ease of testing and became the first untethered spherical tensegrity robot to be able to successful "walk" on an uphill slope.

A medium-scaled versions of the robot, TT-$5_{meso\text{-}impact}$ for impact testing was developed (0.5 m rods) and TT-5_{meso} with 6-12 actuators (0.72 m rods), and designed to complement the full scale TT-5 (1 m rods) and the TT-5_{mini} (0.305 m rods) for drop tests, which were used to evaluate robustness with respect to both electrical and mechanical hardware. The TT5-meso prototype features a lighter weight structure that can be carried by a drone, in addition to TPU 3D-printed end caps and strain-stiffening latex tubing for better impact deformation characteristics. Custom-designed modular motor gearbox assemblies also allow for varying the number of actuators (e.g., adaptable from 6- to 12-motor actuation) to easily test different control policies rapidly-developed in software. A central payload containing the custom modular electronics designed in the lab is attached using 12 passive elastic elements connected to each of the rod ends and is used to control all actuation on the robot. To further evaluate the payload protection capabilities of the tensegrity structures, we completed a number of rigorous drop test studies on both the TT-5 Meso as well as the TT5-mini and its variety of different elastic lattice designs to see how different design choices affected the impact-resilience of the robots.

We developed and analyzed different control strategies to allow the robot to achieve dynamic rolling, resulting in more robust and faster locomotion. In terms of hardware testing, we completed two testbeds which enabled us to investigate the impact, payload interaction, and payload protection characteristics of the robot: a horizontal launcher, a gimbaled thruster testbed, and the TT-4_{impact} for vertical drop testing. Finally, 12-bar forms of spherical tensegrities were investigated in parallel for their locomotive capabilities and impact characteristics.

A method for incorporating the restitution behavior of the tensegrity robot into a mission level path-planning scheme was developed. This involved the formulation of a simplified restitution model for the robot's bouncing behavior and the use of this model in the estimation of landing locations and landing zones during path-planning using thruster based hopping to achieve risk-aware and conservative paths. As a supplement or potential alternative to deep reinforcement learning, another approach to control multi-cable actuation using model predictive control was demonstrated in simulation.

Finally, in our no-cost extension year, we worked with Squishy Robotics, Inc. to use their hardware for further refinement and testing. Squishy Robotics, Inc., a spin-off of our ESI research to commercial tensegrity robots for disaster response on Earth. Due to our previous difficulty in reconciling the stiffness of the tensegrity structure with the power needed to actuate the cables, the new robot was divided into an "active" system focused on locomotion and control policy testing, and a "passive" system for drop test experiments, with the goal of better understanding their independent behaviors and ultimately merging the two prototypes in the future. A new control board was developed for the active system, and features a wide array of scientific sensors, long-range radio communication, and versatile motor driver capabilities. While the full active system prototype is currently being manufactured, the passive system went through multiple iterations of its structural design after repeated drop tests, and can currently withstand falls of up to 400 ft (122 m), the highest drop legally allowable from a drone.

3. ACTIVITIES AND ACCOMPLISHMENTS

We have proposed and completed the following tasks for the previous award year (Year 1, Q1-Q4).

- Task 1: Evaluate appropriate thrusters (Q1-Q2).
- Task 2: Expand simulations to include thruster-based mobility (Q2-Q4).
- Task 3: Evaluate hopping profiles for smooth terrain (Q1-Q2).
- Task 4: Evaluate hopping profiles for hilly terrain (Q2-Q3).
- Task 5: Develop control algorithm for gimbaled thruster (Q2-Q4).

We have proposed and completed the following tasks for the previous award year (Year 2, Q5-Q8).

- Task 6: Develop control algorithm for tensegrity cable-based thruster orientation (Q3-Q5).
- Task 7: Develop control policy for navigation (Q4-Q5).
- Task 8: Simulate complete mission profile on smooth terrain (Q4-Q5).
- Task 9: Simulate mission profile on hills and craters (Q5-Q6).
- Task 10: Manufacture and assemble robot ball hardware (Q6-Q8).

We have proposed and completed the following tasks for the final award year and no-cost extension (Year 3, Q9-Q12 & Year 4, Q13-Q16).

- Task 10 (continued): Manufacture and assemble robot ball hardware (Q10-Q15).
- Task 11: Test control algorithms on tensegrity ball hardware (Q9-Q15).
- Task 12: Test payload protection under impact profiles consistent with mission (Q9-Q15).

We first summarize the technical work performed in the previous award years 1 and 2 and discuss the final year's work afterwards. We made significant progress in achieving TRL 3 to show proof-of-concept with mission profiles tested in simulation with hardware validation.

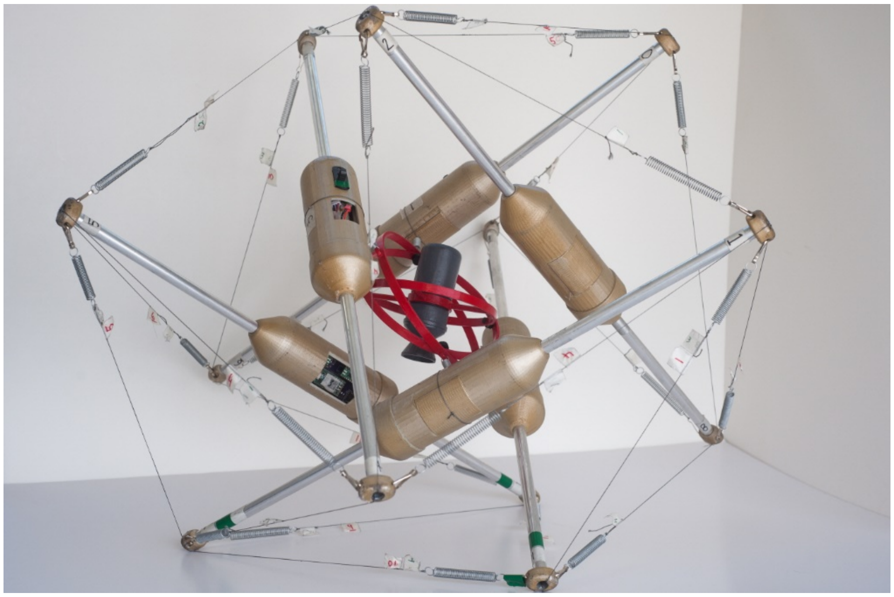

Figure 1. The third version of a rapidly prototyped six-bar tensegrity robot (TT-3) at UC Berkeley. The six golden capsules, located at the centers of each rod, contain distributed controllers as well as electronic components. At the center of the robot is a mock-up of a gimbal-enclosed thruster connected to the outer structure by additional cables.

3.1. Tasks for the Previous Award Year 1

3.1.1. Task 1: Evaluate appropriate thrusters (Q1-Q2)

Several options are present for propulsion systems, namely, solid rockets, monopropellant and bipropellant propulsions, cold gas thrusters, etc. In this work, we have chosen cold gas thrusters for the following reasons: a) they are safe to operate in university research settings, b) they are of low system complexity, c) they are inexpensive and readily available, d) they provide low thrust propulsion, yet their thrust levels are sufficient for hopping of lightweight tensegrity robots.

Several propellants are available for cold gas thrusters; among them, nitrogen, helium and carbon dioxide are the most popular. While helium has high specific impulse, it also has a very low density, which necessitates a large volume for storage, which is not favorable for the robot having a limited space at its center. Carbon dioxide can also be problematic because it is stored in mixed gas and liquid phases and liquid carbon dioxide needs extra care in handling. Overall, in our initial analysis, nitrogen is deemed to be a promising choice for the robot's thruster. Our preliminary analysis has shown that, although nitrogen has relatively low specific impulse, it can still provide sufficient thrust for hopping of the robot.

A preliminary analysis on the feasibility of a nitrogen thruster was performed in Q1, using design and derived parameter values summarized in Table I. These parameters include: nozzle efficiency λ, propellant mass flow rate dm/dt, specific impulse of nitrogen I_{sp}, nozzle inlet pressure P_0, nozzle's throat area A_t, characteristic exhaust velocity C^*, speed of sound in nitrogen a_0, nozzle's expansion ratio ε, specific gas constant of nitrogen R, specific heat ratio of nitrogen γ, temperature T and estimated thrust of F_t.

Table I. Parameters and Thermodynamic Quantities

λ	0.95	a_0	204 m/s
P_0	400 psi	ε	5
A_t	1.20e-05 m2	C^*	252 m/s
γ	1.4	$\dot{m}$	0.132 kg
R	297 J/(kg-K)	I_{sp}	46.5 s
T	100 K	$\boldsymbol{F_t}$	50 N

The trajectories of the robot at a 45-degree angle, initial mass of 10 kg, thrust of 50 N and a burn time of 9 seconds are shown in Figure 2. The robot traveled more than a kilometer with this single hop. The total mass of propellant required for this hop is 1.19 kg. Including propellant, tank and robot weight we would expect our total weight to be within our 10 kg limit. Therefore, it is expected that the robot will be able to travel a farther distance than what is shown in Figure 2. However, the payload will receive significant impact forces if only a single hop is used. Analysis in the next section using the NASA Tensegrity Robotics Toolkit (NTRT) shows performance under a variety of multi-hop scenarios.

3.1.2.Task 2: Expand simulation to include thruster-based mobility (Q2-Q4)

Simulations in this work were primarily done with the NASA Tensegrity Robotics Toolkit (NTRT). NTRT is an open-source simulator developed by the NASA Ames Intelligence Robotics Group to foster researches related to tensegrity robotics. NTRT provides all the core methods to model, simulate and control broad types of tensegrity robots.

As a first step of simulation, a six-strut tensegrity structure was modeled in NTRT by adopting physical parameters from UC Berkeley's rapidly prototyped robot (Figures 1 & 3). To simulate thruster-based mobility, a single vector of thrust force was applied to the center of the payload which is located at the center of the structure. To model real-world disturbances, noise was added to the magnitude and orientation of the thrust. The noise property of the actual system will be dependent on the design of the system and environment where the robot operates. However, since the noise property of our system is not known at this stage, we chose to use a simple Gaussian model for the noise. At each time step, the mean of the noise was set to the current values of thrust magnitude and orientation angles with the standard deviations of 0.02 and 0.002, respectively. As a result, an axis of a thruster nozzle and thrust orientation are not aligned, and an error between these two vectors accumulates over time. This is a rather pessimistic open-loop control model that results in large positional errors. However, results below show that even under this assumption it is possible to meet design goals.

In our simulations, different ground conditions were considered: 1) Smooth terrain and 2) Hilly terrain. Furthermore, three hopping profiles were simulated for each terrain condition. For each hopping profile, the desired flight distance per hopping was 1) 1000 m, 2) 100 m, and 3) 10 m. The required thrusts and burn times for each hopping distance were first obtained with a particle mass system model developed in the first quarter and then applied to the robot in the NTRT environment. Specifically, burn times in the simulations were 9.3 s, 2.94 s, 0.93 s for hopping distances of 1000 m, 100 m, 10 m, respectively, with a thrust of 50 N.

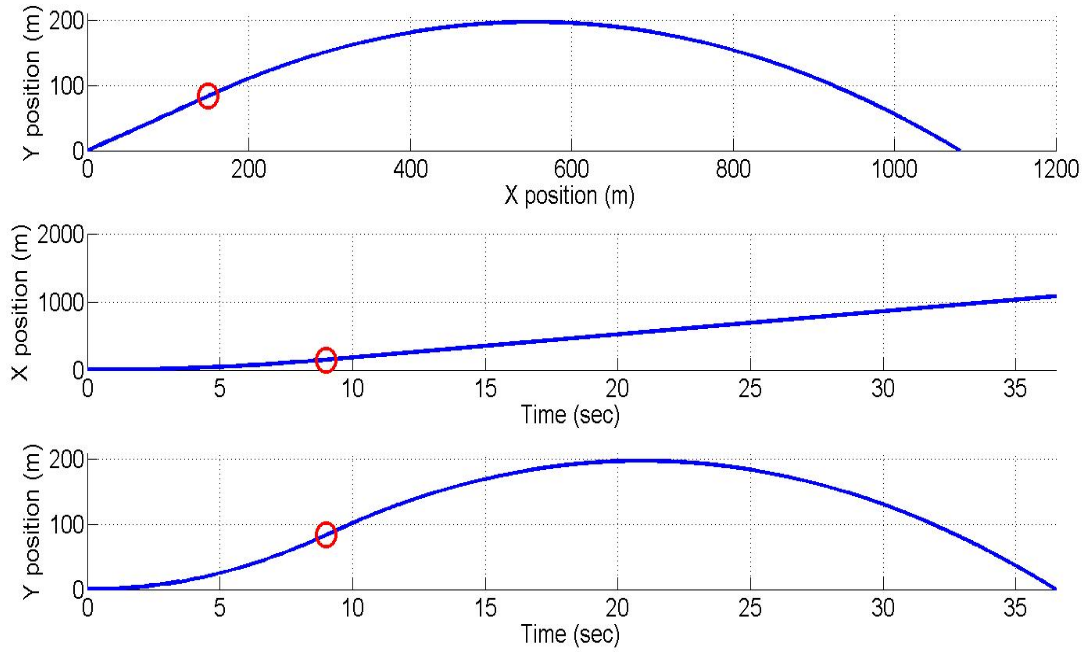

Figure 2. Trajectories of a thruster robot performing a single hop. Red circles represent end of burning period.

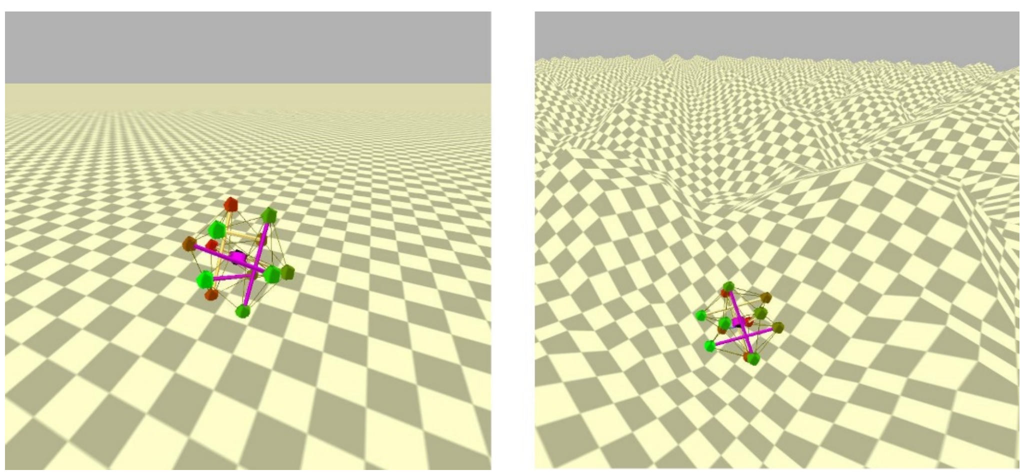

Figure 3. Thruster tensegrity robot modeled in NTRT on smooth (left) and hilly terrains (right).

3.1.3. Tasks 3, 4: Evaluate hopping profiles for smooth (Q1-Q2) and hilly (Q2-Q3) terrains

The main objective of this robot is to travel 1 km on the Moon with precision. Unlike other conventional rigid body robots, tensegrity robots are lightweight and compliant allowing them to travel long distance efficiently by hopping without damaging them from impact at landing.

Several options are present for possible hopping trajectories. The robot may travel the whole distance at a single hop or it may break its path into multiple hops. We consider the following three representative cases categorized by a desired flight distance per hop: a) 1000 m, b) 100 m, and c) 10 m. The choice of the hopping trajectory will depend on several factors, such as terrain conditions, presence of obstacles, energy expenditure, etc. In the first award year, we simulated the abovementioned hopping profiles on two different terrain conditions, namely, smooth and hilly terrains, and examined their energy expenditure, as measured by the total amount of propellant required for the robot to carry to accomplish its goal.

Some examples of the flight trajectories of the thruster robot for different hopping profiles and ground conditions are presented in Figure 4. In simulations, it is assumed that the target is located 1,000 m away from the initial position of the robot in +X direction.

By comparing plots on different rows of a column in Figure 4, one can see that, under a given terrain condition, the final location of the robot is closer to the target as the burn time per hopping (or equivalently, a nominal flight distance per hop) gets smaller because the hopping resolution increases (Figure 5 and Table II). Moreover, by comparing plots on different columns of a row in Figure 4, one can find that, for a fixed burn time, the total number of hops performed by the robot until it arrives close enough to the target is smaller when it is traveling on a smooth terrain than on a hilly terrain (Table II). This is because the distance that the robot bounces and rolls upon impact after each hop is farther on a smooth terrain as there is no obstacle blocking this secondary motion of the robot while, on a hilly terrain, the robot is easily trapped in between hills (Figure 6).

In operation of a cold gas thruster, the amount of propellant used is closely related to the total burn time. The total burn time of all cases of different hopping profiles and ground conditions are summarized in Table III. While a hopping profile with a greater burn time per hop tends to be more energy-efficient (Table III), it is less accurate in hitting the target location (Figure 5). This suggests that a hopping profile that is a mix of multiple long and short hops will not only be more energy-efficient than a single long hop strategy, but also is better able to position the robot to the target location before precision rolling begins. Furthermore, we would like to include as many long hops as possible in this hopping profile, in order to maximize the energy efficiency and minimize the amount of propellant used. However, it is not desirable to use a hopping profile with a single long hop, such as the cases shown in the top rows of Figure 4, because the robot reaches a maximum height of over 100 m with this profile and the fall from such a height may damage the robot upon landing. In our previous research before the grant, we showed that a tensegrity structure can survive from a 10 m drop under the Earth's gravity. In terms of a terminal velocity, this corresponds to a 60 m drop under the Moon's gravity, and thus we expect the robot to land safely after each hop as long as its maximum height is kept below 60 m. According to the first principle particle mass model we developed in Q1, this corresponds to a nominal burn time of 5.3 seconds and a nominal flight distance of 330 m per hop with a 50 N of thrust. Figure 7 shows examples of such hopping profiles consisting of either two or three hops. The variation of the peak heights is the consequence of noise in the thrust magnitudes and orientations we added to our simulations. With this profile, the robot was able to reach to a point 63.7 m and 95.4 m away from the target in average on smooth and hilly terrains, respectively. From this point, the robot can roll towards the target by precise rolling. These hopping profiles resulted in less total burn times than those estimated for the hopping profiles consisting of 100 m hops, see Table III.

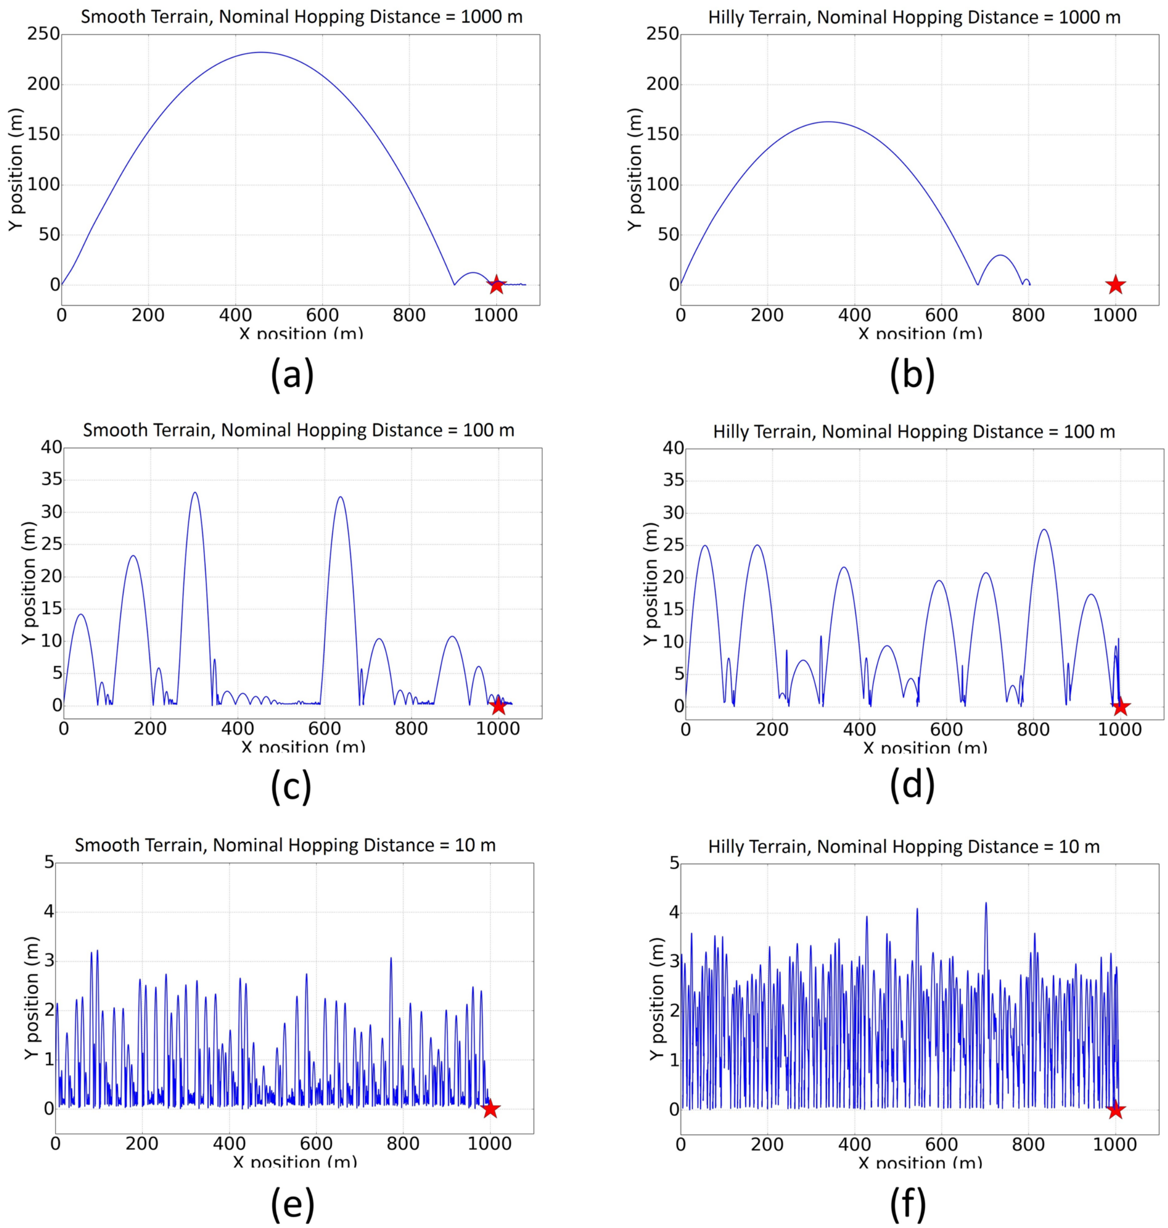

Figure 4. Example flight trajectories of a thruster robot on smooth (left column) and hilly terrains (right column). The nominal hopping distances per hop are 1000 m (top row), 100 m (middle row) and 10 m (bottom row). Red star markers represent the target location. Generally, for a given nominal hopping distance, the robot makes fewer hops on a smooth terrain than on a hilly terrain until it reaches to the target because it bounces and rolls more upon impact after each hop.

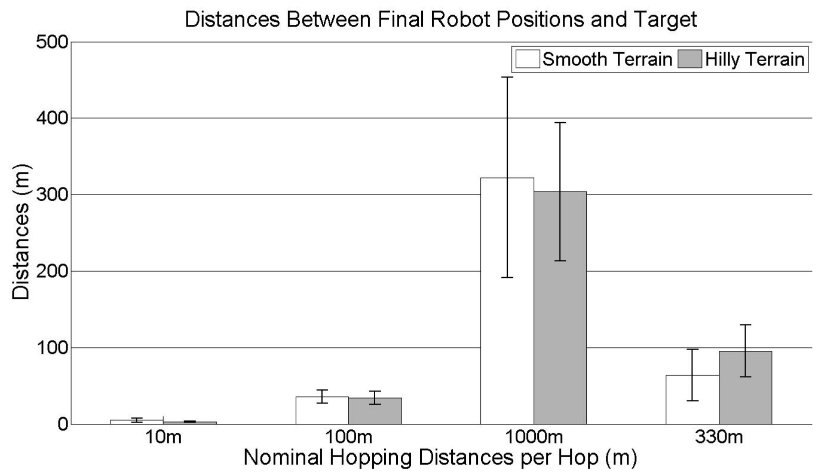

Figure 5. Average distances between final locations of the robot and the target after hopping phase is over. Simulations are run five times for each nominal hopping distances on smooth and hilly terrains assuming standard deviations of 0.02 and 0.002 for thrust magnitude and orientation angles, respectively. In both terrain conditions, the final location of the robot is closer to the target as a nominal flight distance per hop gets smaller because the hopping resolution increases, but at the cost of increased energy expenditure.

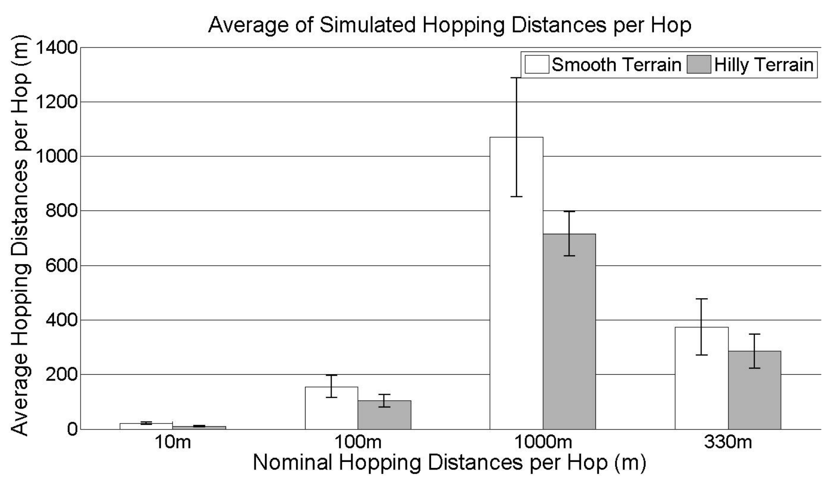

Figure 6. Average of simulated hopping distances per hop for different hopping profiles and terrain conditions. Total of five simulations are run for each case assuming standard deviations of 0.02 and 0.002 for thrust magnitude and orientation angles, respectively. For each nominal hopping distance, the robot travels farther on a smooth terrain than on a hilly terrain because it bounces and rolls more upon impact after each hop.

Table II. Averages of total number of hops performed, simulated hopping distances per hop and final distances between robot and target under different hopping profiles and terrain conditions. Total of five simulations are run for each case.

	Nominal Hopping Distances	Smooth Terrain	Hilly Terrain
Average of total number of hops performed to the target	10 m	53.6	111.4
	100 m	7.4	10
	1000 m	1	1
	330 m	2.6	3.8
Average of actual hopping distance per hop	10 m	20.3 m	9.6 m
	100 m	155.4 m	103.7 m
	1000 m	1069.5 m	715.3 m
	330 m	388.2 m	285.2 m
Average distance between final robot positions and the target after hopping is done	10 m	4.8 m	2.9 m
	100 m	35.8 m	34.0 m
	1000 m	322.4 m	303.8 m
	330 m	63.7 m	95.4 m

Table III. Average of total burn time under different terrain conditions and hopping profiles (simulations are run five times for each case).

Nominal Flight Distance per Hop	10 m	100 m	1000 m	330 m
Smooth Terrain	49.9 s (53.6 hops)	21.8 s (7.4 hops)	9.3 s (1 hop)	13.8 s (2.6 hops)
Hilly Terrain	103.6 s (111.4 hops)	29.4 s (10 hops)	9.3 s (1 hop)	20.1 s (3.8 hops)

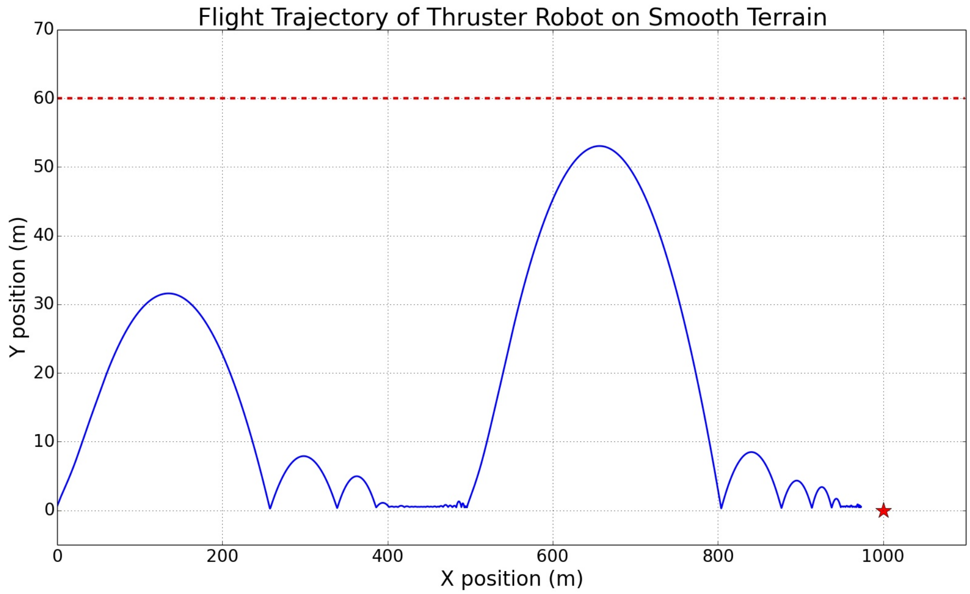

(a)

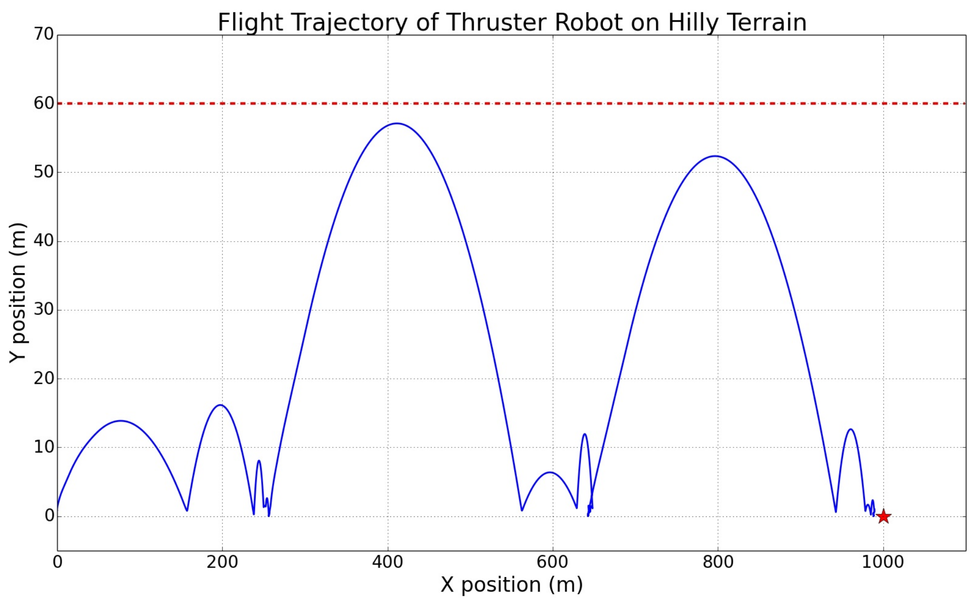

(b)

Figure 7. Example flight trajectories of a thruster robot with a hopping profile consisting of three hops with a maximum height below 60 m seen from side. The simulations are run on (a) a smooth terrain and (b) a hilly terrain. Red dashed lines indicate the 60 m height constraint. Red starts represent the target location which is 1,000 m away from the initial position of the robot. The total burn time for this hopping profile is 15.9 s. On a smooth terrain, the robot reached to the target after two hops while it did so with three hops on a hilly terrain. This is because the robot bounces and rolls farther on a smooth terrain after landing each hop. Once a hopping phase is over, the robot may move close to the target by precision rolling.

3.1.4. Tasks 5, 6: Develop control algorithm for gimbaled thruster (Q2-Q4) and tensegrity cable-based thruster (Q3-Q5) orientation control

Under the presence of noise in the thruster simulation (standard deviations of 0.02 and 0.002 for thrust magnitude and orientation angles, respectively), it was observed that the trajectories of the robot were not necessarily the most efficient. An error in the initial thrust direction could cause off-track lateral motion of the robot, wasting part of its propellant in moving towards an undesirable lateral direction. This behavior is also expected to happen in a real-world system and the need for thrust orientation control arises for energy-efficient operation of the robot. To address this issue, we have proposed to develop suitable controls for both changing thruster orientation while hopping with different types of gimbal systems.

There are many ways to implement a gimbaled thruster system. We explored three high level approaches for adjusting the thruster direction during a hopping event in the previous year.

(1) One concept is to install a two degree-of-freedom gimbal at the nozzle of the thruster, which will be referred to as a gimbaled-nozzle thruster. Gimbaled-nozzle thruster systems are well researched in rocketry and space flight for flight direction control (Figure 8). One of the designs that allow us to achieve the gimbaled-nozzle thruster is the Canfield joint system. It is compact and efficient mechanism that provides a full hemisphere of motion.

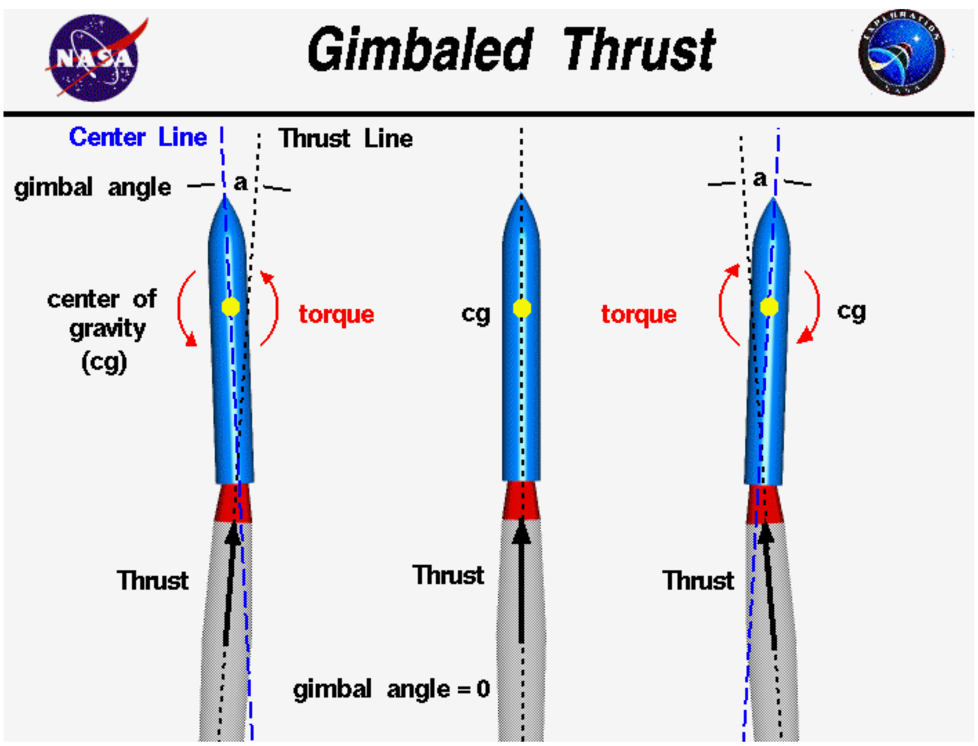

Figure 8. In a gimbaled thrust system, the exhaust from the gimbaled-nozzle is used to vector the thrust to control the flight. (Source: NASA, https://exploration.grc.nasa.gov/education/rocket/gimbaled.html/.)

(2) Another concept is to enclose the thruster inside of a 2 degree-of-freedom gimbal structure, which will be referred to as the gimbal-enclosed thruster system. From our evaluation, the gimbal-enclosed thruster system allows the thruster a much larger range of motion in comparison to the gimbaled-nozzle thruster system. We modeled and prototyped this concept using a 3D

printer as shown in Figure 9 to understand the large range of motion and visualize the payload size and spacing within the tensegrity robot. The payload will be attached on either the gimbal or to the thruster. Figure 10 shows when the robot is at a desired location for hopping but the thruster is pointing at the ground, the gimbal-enclosed thruster system can fully adjust the thruster to the correct orientation. This can be useful when the robot has to adjust the thruster orientation for hopping but it cannot do so by moving its whole body due to its environment, for example, when the robot is stuck at a crater. In Figure 11, the gimbal-enclosed thruster system is shown at the center of the tensegrity robot, and at the center of the gimbal structure a thruster and imaging system is attached. This concept allows not only the thruster but also the payload to have the advantage of the gimbal system. For example, this setup will allow the imaging equipment to rotate and sweep within the gimbal system.

(3) The third concept we explored for controlling the thruster orientation is to change the shape of the robot structure by using the shape-shifting capability of the cable actuation of the outer tensegrity structure. In the simplest form of this system, the thruster and payload would be connected to fixed length cables to the ends of the tensegrity rods shown in Figure 12 and orientation would be achieved by changing the outer structure shape. A variant of this approach is to control the length of the cables ("inner cables") that attach the thruster and the payload at the center. With this approach, the cable-actuated thruster can be oriented by controlling the lengths of inner cables connecting the thruster to the outer robot structure or those of outer cables constituting the robot structure or both. This approach has an advantage over the gimbal-enclosed thruster system in that there is no need to add a complicated gimbal system at the center of the structure and we can fully use the center space for installation of a payload shown in Figure 13. A downside of this approach is that the orientation error may be larger than either of the other two concepts being explored – the gimbaled-nozzle or the gimbal-enclosed thruster. In addition, the degrees of freedom achieved could be lower and the control and actuation complexity could be higher.

3.2. Tasks for the Previous Award Year 2

3.2.1. Task 5: Develop control algorithm for gimbaled thruster (revisited)

In Year 1, we proposed two gimbal-based systems, namely, gimbal-enclosed and gimbaled-nozzle thrusters, for the purpose of regulating a thrust vector. This year, we continued our work in this direction and made further progress by exploring more alternatives. First, among the two gimbal-based mechanisms, we concluded that the gimbaled-nozzle system would be more effective than the gimbal-enclosed system based on our trade study of strategies for thruster control. Thus, we only consider the gimbaled-nozzle system and develop its control algorithm in this section. However, reaction wheel systems are a strong contender and an evaluation of its potential is also explored. More details on the trade study is provided in Sect. 3.2.4.

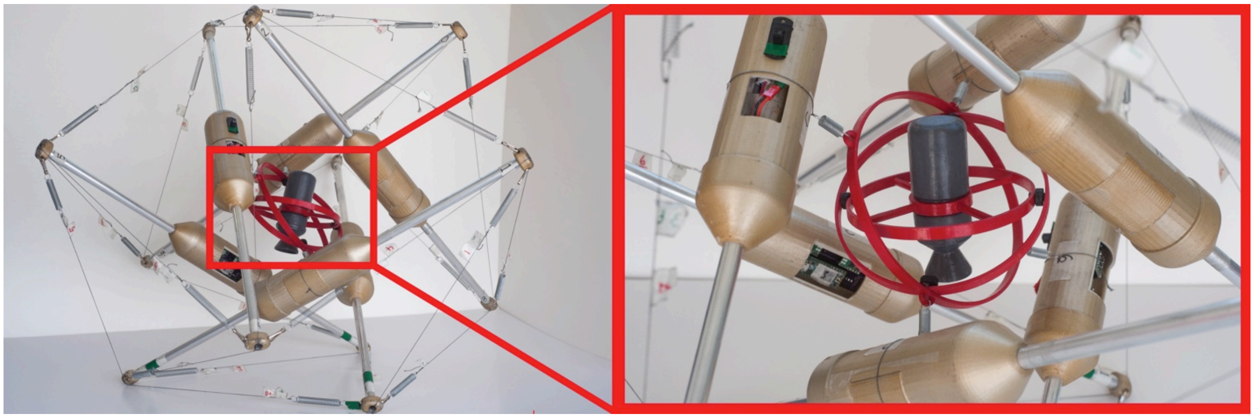

Figure 9. A mockup of a gimbal-enclosed thruster system located at the center of the robot. Red and grey parts are mockups of a gimbal and cold gas thruster, respectively. The gimbal-enclosed thruster is connected to the robot structure with springs.

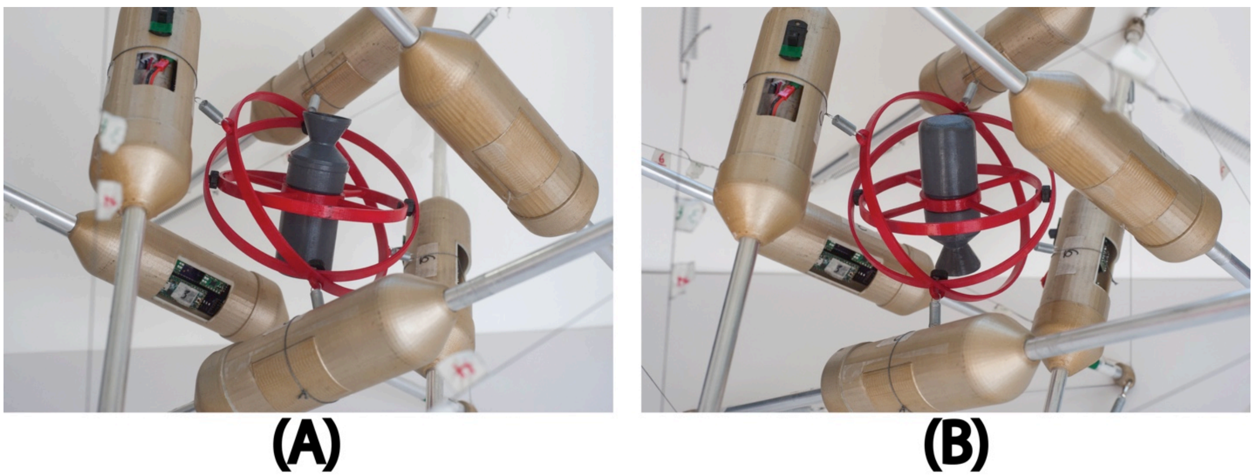

Figure 10. The large range of motion of the gimbal-enclosed thruster system allows the thruster to be adjusted to the desired orientation without moving the structure of the tensegrity robot. (A) The cold-gas thruster is in an incorrect orientation for hopping while the robot is at a desired location. (B) The cold-gas thruster can be adjusted to the correct orientation with the gimbal system without moving the robot.

Figure 11. An imaging system was attached to the mockup of a gimbal-enclosed thruster system at the center of the robot to represent the potential payload size and location. With this setup, the imaging equipment or payload can rotate and sweep within the gimbal system.

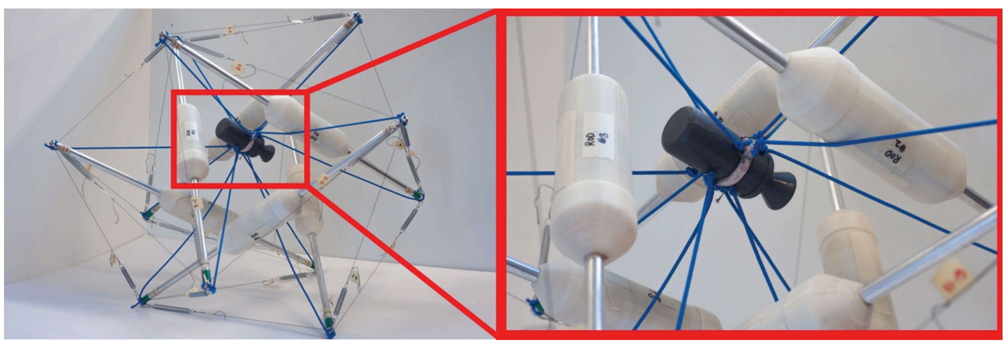

Figure 12. The cold-gas thruster system is at the center of the robot through cable connections to outer tensegrity structure.

Figure 13. (A) The cold-gas thruster system suspended at the center of the tensegrity structure. (B) In addition to the cold-gas thruster system, other payload or imaging equipment can be connected along the thruster system at the center of the robot.

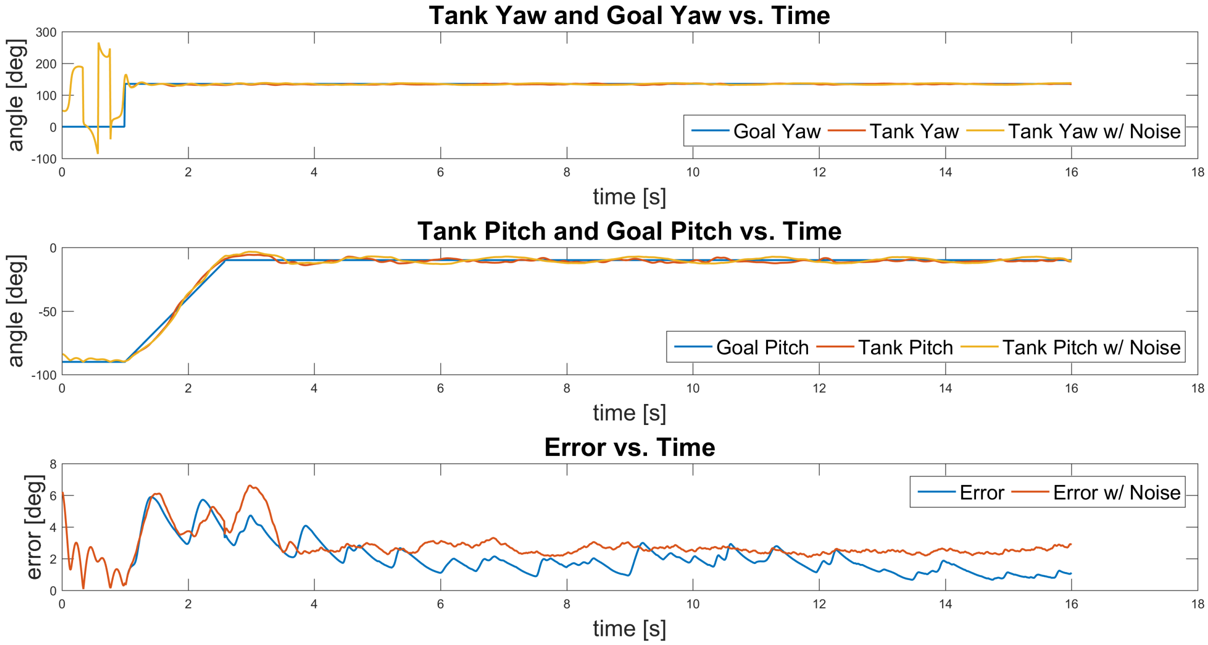

Figure 14. Simulation of gimbaled-nozzle thruster in NTRT. Thrust is activated at t = 1 and turned off at t = 16.

Thrust vector control allows for precise manipulation of the flight path by generating a control moment about the center of mass of the moving body. With rigid body structures, such as rockets, missiles, and satellites, the moment generated is accurately known. With compliant tensegrity structures, however, the precise center of mass of the structure is difficult to predict mid-flight. To test the gimbaled-nozzle method on tensegrities, a model representing the dynamics of an idealized system were used. A feedback controller was created in which the outer-shell, comprised of all six rods and 24 cables, was modeled as a lumped element. We neglected cable forces, outer-shell dynamics and payload orientation relative to the outer shell by treating the entire robot structure, the payload and outer-shell combined, as a single rigid body. Based on the simulations, this assumption is justified when cables are sufficiently stiff, payload inertial forces are dominant over outer-shell inertial forces, and external disturbances and inputs are small. Passive cable stiffness plays a large role in thrust controller stability, with very compliant systems requiring faster actuator responses (Figure 15). Note that stiffness can be dynamically tuned by actuating the cables. Given these assumptions, a linear-quadratic regulator (LQR) was employed to optimally control the system according to a quadratic cost, penalizing the payload heading error and its time derivative. Feedback gains calculated from the infinite-horizon, continuous-time LQR were used to calculate the desired control input moment to follow the reference trajectory. This controller was implemented and evaluated in the NTRT simulation environment, and Gaussian noise with a standard deviation of two degrees was added to the gimbal Euler angles to better represent realistic sensor and actuator conditions. The results (Figure 14) reinforce the idea that gimbaled-nozzle thrust vectoring is an applicable strategy for tensegrity systems given the desired conditions described above.

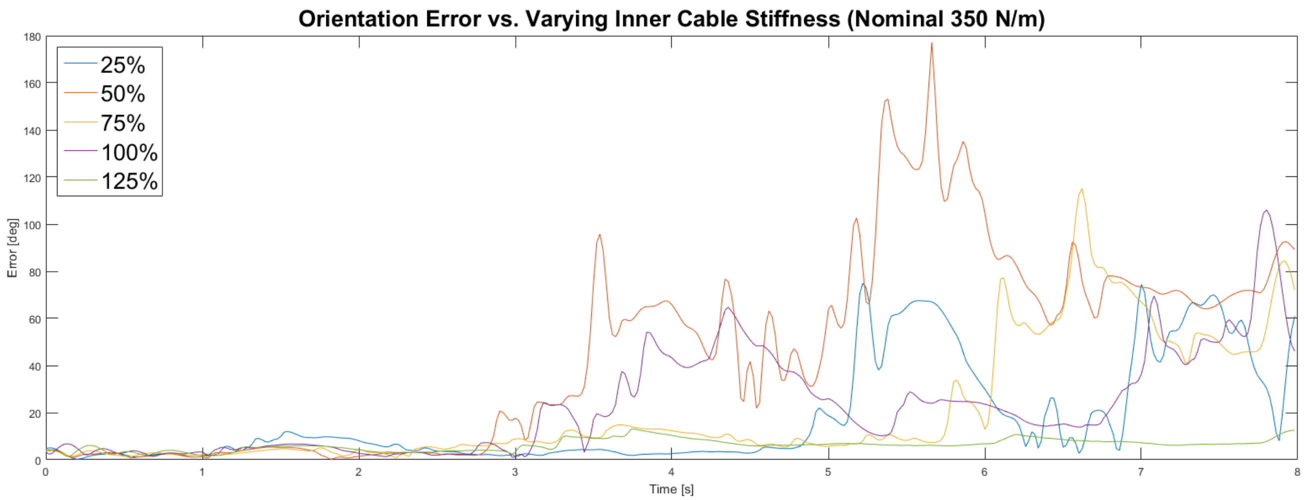

Figure 15. Thrust Control Stability with Varying Inner Cable Stiffness (Thrust for 8 seconds).

3.2.2. Task 6 (continued): Develop control algorithm for tensegrity cable-based thruster orientation (Q3-Q5)

Since tension is distributed in a highly coupled and nonlinear manner across tensegrity members, it is difficult to predict how the tension change will affect the thruster orientation. To examine the validity of this approach, a workspace analysis (i.e., range of thruster orientation angles that can be achieved using cable-actuation) is performed. To do this, a set of target rest lengths of the structure's cables are randomly sampled and applied to a tensegrity dynamics simulation[1] to find one terminal structure shape and final thruster orientation. However, the sample is discarded if the final structure shape places the center of mass outside of its ground-contacting polygon. Such a shape is expected to make the robot to perform punctuated rolling and as a result, the thruster will point towards a wrong direction even after the desired shape-shifting. This procedure is then repeated until enough samples are obtained that would represent a good estimation on the region of thrust angles that could be achieved with cable actuation. In the current analysis, gravitational effects are neglected and collision detection of structural members is not considered, as the main objective of the analysis is to formulate an initial conception on possible ranges of thrust angles.

In Figure 16, three different cases are considered: (a) only inner cables are actuated, (b) only outer cables are actuated, and (c) both inner and outer cables are actuated. The initial azimuthal and elevation angles of the thrust orientation are 0 and 90 degrees, respectively. The ranges of these angles after cable actuation are summarized in Table IV. The range is the widest when both inner and outer cables are actuated because the amount of shape-shifting is the greatest in this case. The range of thrust orientation angles achieved by inner cable actuation is wider than that achieved by outer cable actuation. However, inner cable actuation requires additional actuators and electronics, while outer cable actuation can be performed with the existing actuators used for punctuated rolling, thereby avoiding any additional components. This difference is critical when considering the 10 kg mass restriction that competes with the benefit of having a wider range of orientation angles with inner cable actuation. For this reason, control of thrust orientation solely through outer cable actuation seems the most promising for this system. A downside to this

---

[1] We used a dynamics model from: R.E. Skelton and M.C. de Oliveira, *Tensegrity systems*. Springer, 2009.

approach is that the orientation error may be larger than the other two concepts using gimbals. In addition, the achievable workspace could be smaller and the control and actuation complexity could be higher.

With this approach, when a desired thrust orientation is outside of the achievable workspace, the robot can perform punctuated rolling in order to change its pose until the desired thrust orientation falls under the region of the workspace. The robot can then change its cable tension to control the thrust orientation.

Table IV. Ranges of thrust orientation angles for three different cases in Fig. 14. Units are in degrees.

Actuation Location	Inner cables	Outer cables	All cables
Min. azimuthal	-27	-17	-22
Max. azimuthal	46	55	63
Azimuthal difference	73	72	85
Min. elevation	44	54	42
Max. elevation	120	140	147
Elevation difference	76	86	105

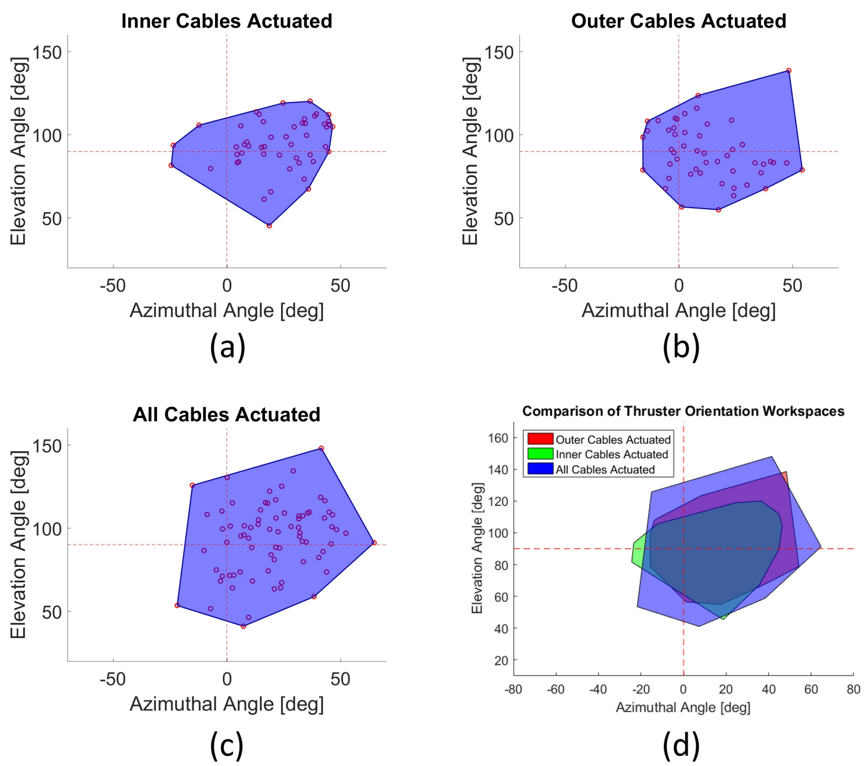

Figure 16. Expected ranges of thrust orientation angles estimated with random sampling. Initial azimuthal and elevation angles are 0 and 90 degrees, respectively. (a) Only inner cables are actuated. (b) Only outer cables are actuated. (c) Both inner and outer cables are actuated. (d) Comparison of the three workspaces.

While the analysis presented above provides estimated ranges of the thruster orientation achievable with cable actuation, it does not tell us how much tension should be provided to the cables given a desired thruster orientation. One possible way to approach this problem is to make use of the random samples already generated for the workspace analysis. That is, we find the sample that is the closest to the desired orientation in both azimuthal and elevation angles. Then we apply the set of target rest lengths that generated the sample. The final thruster orientation achieved with this method will not precisely match the desired target orientation, but the error between the two orientations will become negligible if the total number of samples is large and if the samples are distributed evenly over the workspace.

3.2.3. Alternate approach: Thrust vectoring via two reaction wheels

Reaction wheel systems have been widely used for attitude control of rigid bodies, especially with spacecraft. It is well known that at least three wheels are required in order to orient a body in an arbitrary direction. However, for the thruster orientation control, we only need two wheels as it is not necessary to regulate rotation about the nozzle axis. Such a problem is known as a spin-axis stabilization problem. In addition to the gimbal-based and cable-based approaches for orientation control, we developed a control algorithm[2] that globally and asymptotically stabilizes the thruster about an arbitrary spin axis using two reaction wheels (Figure 17). Because the reaction wheel system is fundamentally an angular momentum exchange device, the additional wheels can be designed as small and lightweight as long as actuators can supply large enough angular velocities to the wheels. The response of the device to a disturbance, however, might be slower than the gimbal-based systems because it may take some time to exchange angular momenta between the wheels and main thruster system. We believe from our preliminary analysis that this could be an effective alternate control system for this mission.

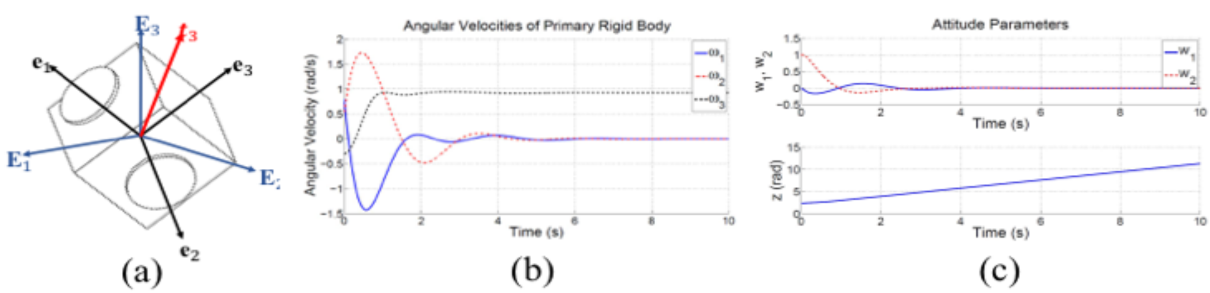

Figure 17. (a) A thruster system (package represented by a cube) with two reaction wheels attached for orientation control. (b) With the developed control algorithm, angular velocities about two body-fixed axes converge to zero. The last angular velocity converges to a nonzero constant, in general, which represents a spin motion of the body. (c) Two of the three angular position parameters (w_1 and w_2) converge to zero with the presented control algorithm. The last angular position (z) increases or decreases linearly as a result of the spin motion.

[2] For more details, refer to our recently published work: K.Kim and A.M.Agogino, *Spin-Axis Stabilization of a Rigid Body about an Arbitrary Direction using Two Reaction Wheels*, presented in the Proceedings of 55th IEEE Conference on Decision and Control, Dec. 12-14, Las Vegas, NV, USA.

3.2.4. Trade study of strategies for thruster control

A trade study was conducted in order to evaluate potential solutions for thruster attitude control, taking into account the utility and benefits/costs of each solution, including qualitative considerations such as risks of failure. Five solutions were evaluated: gimbal-enclosed thruster, gimbaled-nozzle thruster, cable-actuated morphing, reaction wheels, and rods with end thrusters. We used the criteria of mass, energy consumption, control actuation accuracy, volume, and mechanical robustness. Details are provided in Table 5 and Figure 18. Based on the quantitative analysis, the two most effective strategies for controlling the thruster are the (1) gimbaled-nozzle system and (2) reaction wheel system. Among these two control methods, we chose to proceed with the gimbaled-nozzle system because we believe it is a more widely accepted solution for the orientation control problem of the system of our scale. We have simulated the gimbaled-nozzle system in NTRT to control the thruster system on a tensegrity robot (Figure 19).

Table V. Trade study results. Alternative solutions are given scores for each criterion based on estimates using example hardware and simulation results. Total score is a weighted sum of criteria scores according to individual utility curves, and delta deviations allow for consideration of uncertainties in evaluations.

Criteria	Total Mass (kg)	Energy Consumption per Hop (W-h [current total 18 W-h])	Resolution (Degrees)	Volume (cubic inches)	Structural Robustness (max ft. drop)	Total Score
-δ	8	0.50	3	700	1	3.61
Gimbal-Enclosed Thruster	5	0.20	2	381	3	5.12
+δ	3	0.10	1	300	5	5.88
-δ	4	0.50	3	268	1	4.78
Gimbaled-Nozzle Thruster	2	0.10	2	113	2	5.88
+δ	1	0.06	1	33	5	6.73
-δ	3	0.05	20	0	5	5.15
Cable-Actuated	0	0.04	10	0	7	6.75
+δ	0	0.03	5	0	10	7.23
-δ	6	0.66	1	524	2	4.56
Gyroscopic Flywheels	4	0.12	1	180	4	5.83
+δ	3	0.08	1	113	5	6.16
-δ	10	2.00	20	900	2	0.56
Rods with Thrusters	8	1.75	10	700	5	3.27
+δ	6	1.50	5	500	10	4.75

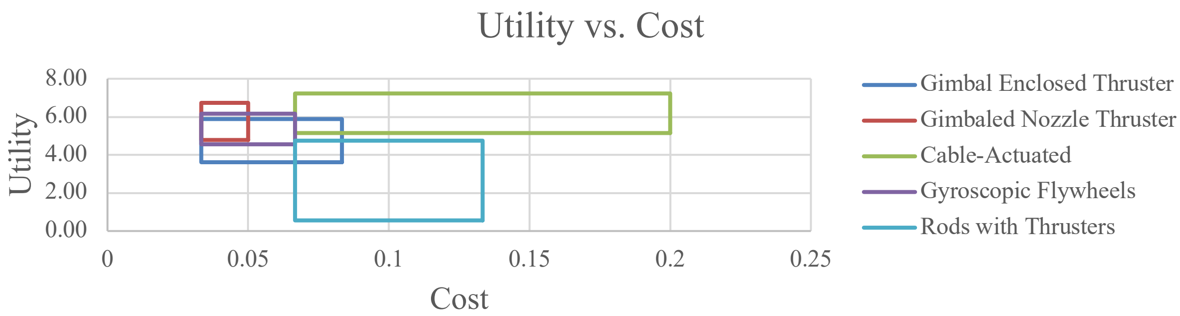

Figure 18. Visualization of the trade study of gimbal options, summarized in utility versus cost plots.

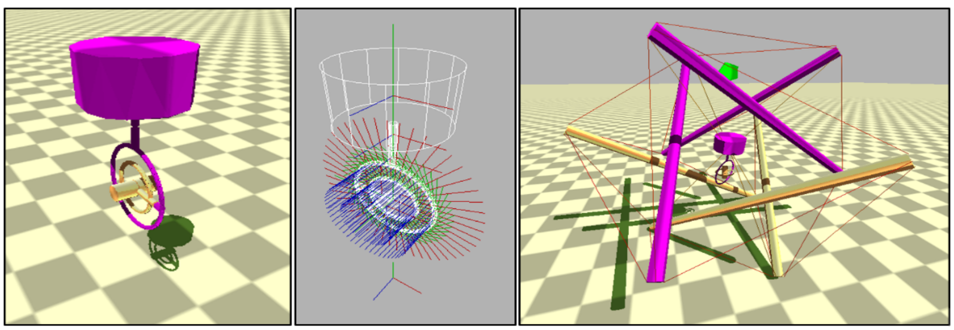

Figure 19. Payload and gimbaled-nozzle thruster model in NTRT (Left). Wireframe of NTRT model depicting rigid body reference frames (Center). Entire robot structure with thruster suspended in center and green marker representing thruster direction (Right).

3.2.5. Task 7: Develop control policy for navigation (Q4-Q5)

The primary goal of the tensegrity probe is to reach a target location that is 1 km away from its initial location on the Moon. For safe and energy-efficient navigation, the tensegrity probe may have to combine rolling and hopping when traveling on the Moon's surface, largely depending on terrain conditions of the target exploration region. In other words, the probe needs to be equipped with appropriate path planning algorithms for completion of its mission.

For a realistic analysis, we consider the map of an actual lunar surface[3] instead of a structured hilly terrain that we considered in in Year 1. The map of the mission surface is assumed to be readily available and divided into a grid, with each grid cell representing a square patch with local height information. This representation allows importing height maps directly into the simulation to test the algorithm on real data. Before the robot can plan its path and actions to reach the goal, it first localizes itself on the given map. To this end, the robot relies on a dynamically updated belief space that represents the probability of the robot being at a specific

[3] The map used in this task was obtained from USGS Astrogeology Science Center. In Q4 and Q5, we were not able to import the actual lunar map into NTRT, so we developed our own software in Java only for this task. In Q6, however, we have completed this import and the simulation in NTRT can now be run with actual lunar terrain maps. See, Sect. 3.2.7.

position. Initially, the probability is uniformly distributed across all possible positions in the map. Assuming that the robot can sense the height of the four neighbor cells – north, south, east, and west - the robot updates its belief map by examining which cells on the map have their neighbors with approximately the same heights. After having its beliefs updated, the robot rolls one step in a direction and repeats the process until it is confident enough of its position (Figure 20). In addition to finding the robot's initial position, this method can also be used to precisely locate itself on the map after hopping which may deteriorate the robot's estimation of its position due to thruster noises and secondary rolling after landing.

Once the robot's position is known, we used the A-star search algorithm[4] with the Euclidean distance to the goal as a heuristic to find the most energy efficient path (Figure 21). The cost function consists of several components, each of which defines the cost of different actions (e.g., rolling and hopping) based on the travel distance with the actions[5]. For hopping motion, an additional constraint on the height difference between the robot's initial and final positions is considered in order to prevent damaging of the robot from hard impact. At each iteration, the robot can move to any of the eight adjacent grid cells by punctuated rolling. This movement is practical only if the difference of height between the two cells is not too important. Otherwise, the robot might not be able to climb, e.g., a steep hill and the cost of taking this action will be significant. On the other hand, because going downhill requires less energy than going uphill or moving on a flat terrain, the robot will exploit this in planning its path. Thruster-based hopping allows a wider range of movement directions and distances. While more expensive than punctuated rolling, hopping is unavoidable in some cases, e.g., when the robot needs to escape from a crater with steep slopes. Although the hopping motion increases the average number of nodes reachable from each location, the A-star algorithm tries to limit the number of nodes expanded in order to reduce the computation effort.

In practice, the robot may fail when trying to execute a movement. To incorporate such failures, we define the possible outcomes of the movement along with their probabilities to happen. Since the A-star algorithm is not designed to handle random errors, the robot takes the following sequence to prevent itself from drifting too much: 1) The robot computes the best path to reach the goal from its current position. 2) The robot takes the first K actions according to its plan. Some actions may fail. 3) The robot localizes itself again after K actions. 4) The robot re-computes the best path for the remaining distance. 5) The robot lowers K to improve its accuracy near the goal. 6) The above are repeated until the robot arrives at the goal. Additionally, in terms of implementation, it may be unreasonable to assume that the robot is capable of measuring its surrounding elevation for localization. Thus, an alternative method is required. Using currently available technology, (e.g., the Lunar Reconnaissance Orbiter), we propose the use of satellite imagery to provide sparse position updates to the robot for it to calculate course corrections. Investigation into the effect of update frequency on navigation accuracy will be done in the next award year.

As the implementation of A-star path-planning in NTRT is an ongoing project, we have also developed simpler navigation policies to use in the complete mission profile simulations to

---

[4] See, for example, S. J. Russell, P. Norvig, J. F. Canny, J. M. Malik, and D. D. Edwards, *Artificial intelligence: a modern approach*. Prentice hall Upper Saddle River, 2003, vol. 2.

[5] At this moment, the cost function consists of our best estimate on the energy expenditure of hopping and rolling motions. A more thorough analysis on the energy expenditure of the robot is in progress, and the cost function will be updated per the conclusion of the analysis in the upcoming quarters.

demonstrate feasibility. This includes: 1) thrust vectoring based hopping and 2) closed-loop punctuated rolling. These will be discussed in more detail in Sect 3.2.7.

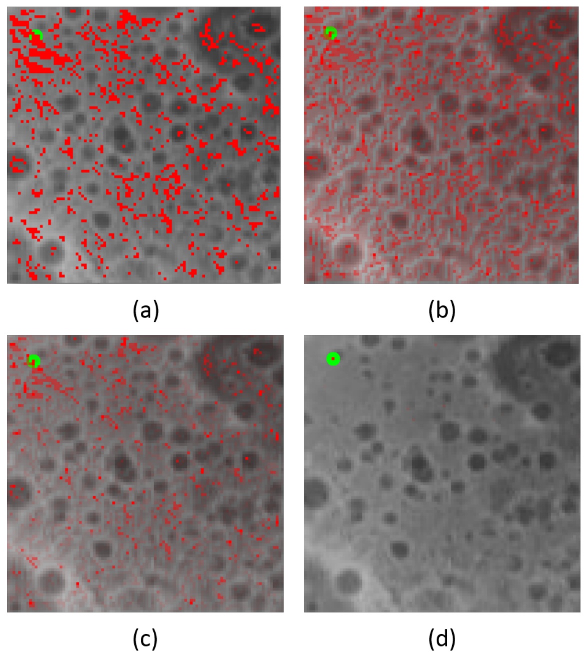

Figure 20. Evolution of the robot's belief map while localizing itself with the known lunar surface map. Red areas represent the robot's estimated positions and green dot is the robot's true location. (a) Belief map after first measurement of neighbor grid heights. (b) Belief map after performing one step of punctuated rolling. (c) Belief map after measuring new neighbor grid heights. (d) Belief map at the end of localization process. Notice that the red areas shrank into the location of the green dot. That is, the robot successfully localized itself on the lunar map. (Image: Courtesy of USGS Astrogeology Science Center).

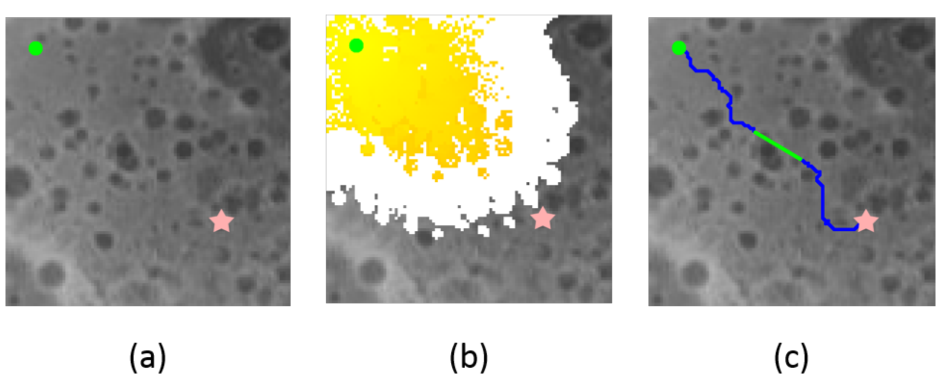

Figure 21. (a) A lunar terrain map with height information with initial and goal locations marked with green dot and pink star, respectively. (b) A-star algorithm is used to search for the best path to the goal position. Yellow nodes are those already expanded while white nodes are those to be expanded. The shade of the yellow nodes is proportional to the cost to reach the grid cell. (c) The planned path and actions. Blue and green represent paths covered with punctuated rolling and hopping, respectively. (Image: Courtesy of USGS Astrogeology Science Center).

3.2.6. Task 8: Simulate complete mission profile on smooth terrain (Q4-Q5)

Combined hopping and rolling motion of the thruster robot is simulated in NTRT on a smooth terrain (Figs. 22 and 23). Stochastic error was added in simulation for both thrust magnitude and orientation – Gaussian noise with standard deviations of 0.02 N and 0.002 radians, respectively. Our simulation showed that the robot can easily travel 1 km and arrive at the target with a small number of hops and punctuated rolling steps.

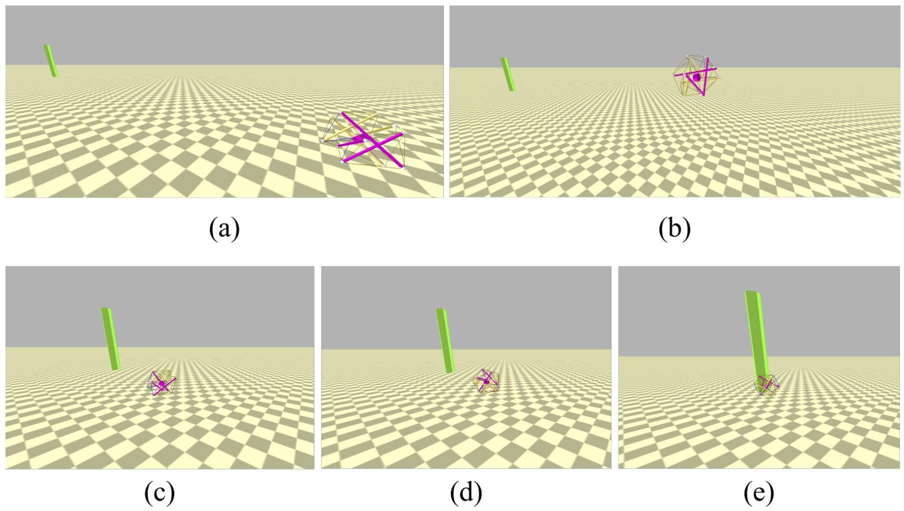

Figure 22. (a) The robot is initially distant from the target (represented by a far green rod) and is preparing for hopping. (b) The robot hops towards the target when the distance between them is large. (c) Once the robot gets close enough to the target, the robot switches to precision rolling. (d) The robot rolls toward the target. (e) The robot arrives at the target. Note that the target is placed closer than 1 km from the initial position of the robot in the above figures for the sake of better perspective.

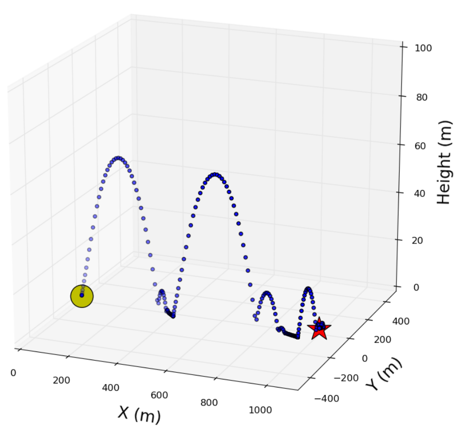

Figure 23. The simulated trajectory of the thruster tensegrity traveling 1 km to reach the target. The yellow dot is the initial position of the robot and the red star is the target location which is 1 km away from the initial robot position. During the hopping phase, the robot made two 330 m hops followed by one 100 m hop and one 10 m hop. After these hops, the robot was close enough to the target to switch to precision rolling.

3.2.7. Task 9: Simulate mission profile on hills and craters (Q5-Q6)

Once the simulation of a 1 km mission on smooth flat terrain was completed, the logical next step was to validate the proposed mission profile under more realistic simulation conditions. This includes improving model detail and accuracy, using realistic lunar terrain, and implementing more sophisticated navigation policies. This section will discuss each of these aspects in further detail.

One of the biggest areas of improvement upon the model was the implementation of a simulated gimbaled-nozzle thruster as described in Sect 3.2.1. This allowed for the inclusion of actuator constraints on the gimbal and accurate modeling of the thrust as a force *pushing* on the payload instead of *pulling* it, as was the case in previous simulations. Using a gimbaled-nozzle thruster also introduced new constraints, namely the containment of thrust directions within a closed triangle of the robot. This is necessary, as the cables and rods may interfere with or be damaged by the thruster. In order to ensure safe operation of the thruster, limits were imposed on the workspace of the gimbal.

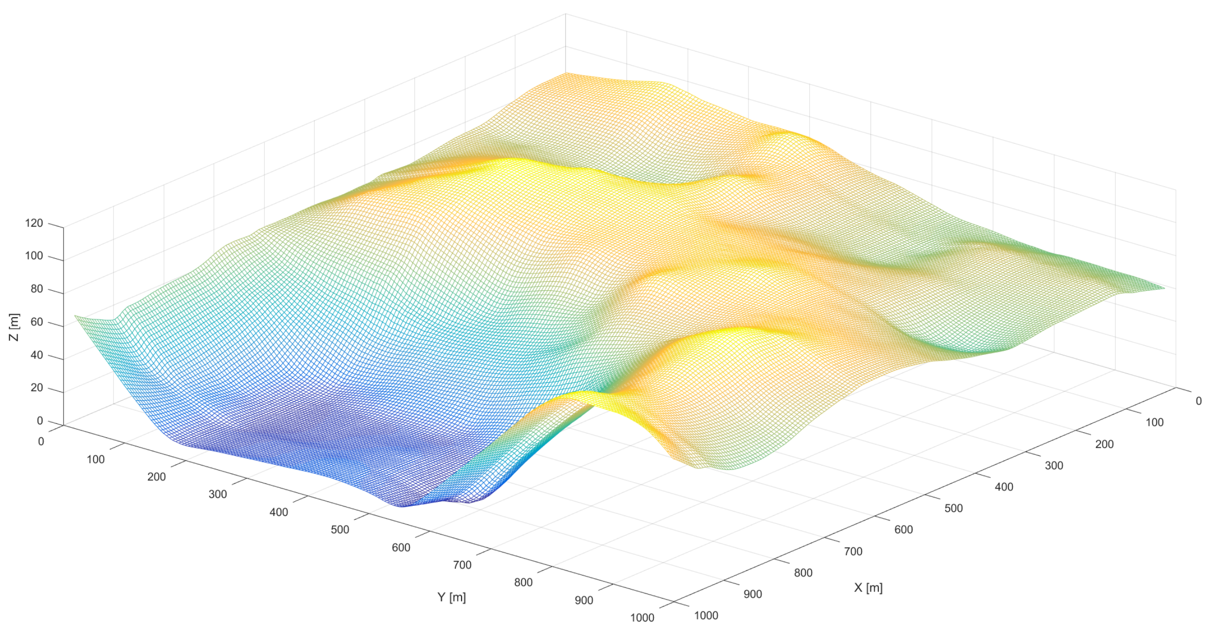

Figure 24. Lunar terrain sourced from Moon2STL.com

To address the shortcoming of unrealistic terrain simulation in NTRT, a terrain import tool was created. This tool allows for the importing of arbitrary user-generated or user-supplied terrains from STL files, which can be built using modeling software such as Blender or SolidWorks or taken from other sources such as the website Moon2STL, where accurate lunar landscapes, such as the one in Figure 24 can be found.

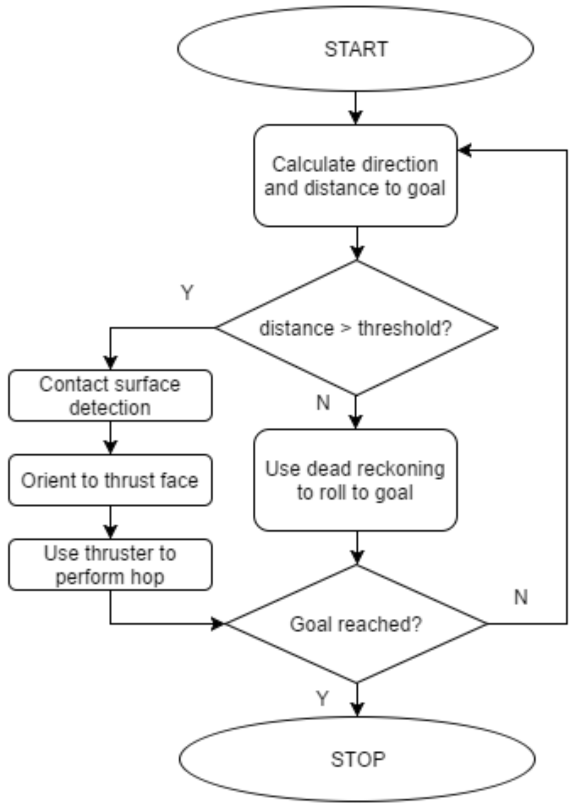

Figure 25. Proposed mission profile flow chart.

New navigation policies for hopping and punctuated rolling were developed using the decision flow chart in Figure 25. The first step in both policies is contact surface detection, which allows for the robot to sense which of its 20 surfaces is currently in contact with the ground. This is done by first calculating the normal vectors for all 20 surfaces of the robot in the body frame. Assuming that the robot's payload includes an IMU (Inertial Measurement Unit), the gravity vector in the robot's body frame can then be found as well. The dot product between the gravity vector and normal vectors will then provide a measurement of how well each face is aligned with the direction of gravity. Finally, the face with the largest dot product is chosen to be the surface in contact. As was mentioned earlier, when considering the addition of a thruster to the robot to allow for hopping motions, a major concern was that of the interference of cables and rods with the thrust output. To address this, a fixed-orientation gimbaled-nozzle thruster with a constrained workspace was used. Assuming that the robot could always be positioned to lift off from a particular launch closed triangle, the solution guarantees that no cables or rods would interfere with thruster operation or risk being damaged. Repositioning of the robot was performed using a path planner to find the optimal path from any triangle to the launch triangle. This was done using Dijkstra's algorithm on a digraph of the tensegrity robot (Figure 26). For implementation, it was assumed that all edges in the digraph are weighed equally. In addition, travel between surfaces could only occur when perpendicular to the edges of the robot. Once a path is generated, actuation is carried out according to a policy table describing the cable that must be actuated for that motion to succeed. Once the robot has reoriented itself to the launch face, hopping using thrust-vectoring described in section 3.2.1 is performed. Assuming accurate knowledge of the target destination relative to its current position, the robot intelligently decides which heading direction to follow and what duration of active thrust is needed to best approach the target destination.

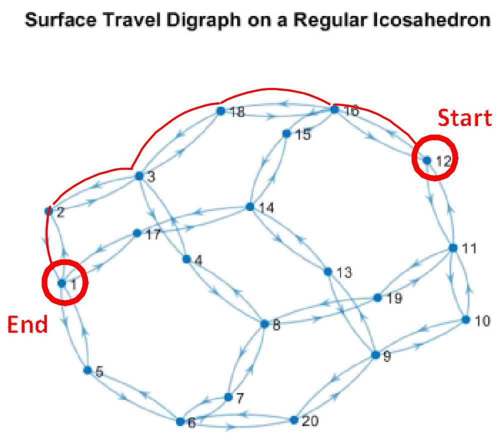

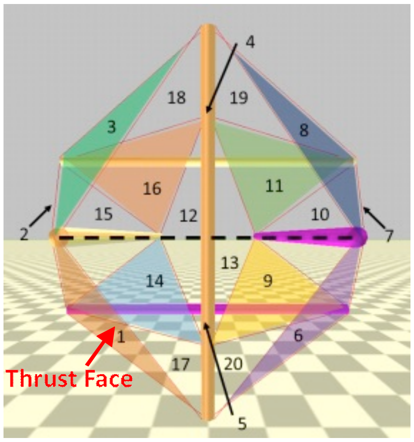

Figure 26. (Left) Surface connectivity digraph of a 6-Bar robot. (Right) Surface number convention with example thrust face highlighted

From the mission profile flow chart in Figure 25, while the robot is above a certain distance threshold to the goal position, it will continue to use hopping motions to cover large distances. However, once it lands within that threshold, the robot will change its navigation policy to that of closed-loop punctuated rolling. This is performed in a similar way to the reorientation of the robot where instead of rolling to the launch face, the robot will find a closed triangle that is most closely aligned with the direction vector towards the goal. It will then roll to that face then repeat the sequence until it is within a certain distance of the goal. This then completes the navigation mission of the robot.

Previously, we have shown results from flat or simple rolling artificial terrain. To validate our results further, we tested our methods on more realistic scenarios using our terrain import tool to import real lunar terrain (Figure 24). Realistic hopping profiles were generated using results from these simulations with gimbaled-nozzle thrust vectoring. When used in conjunction with trajectory-generated rolling, gimbaled-nozzle thrust – which includes Gaussian noise with standard deviation of .03 radians added to the gimbal angles – presents a realistic mission profile with multiple hops towards a specified goal destination. Additionally, precise control of flight trajectory allows for greater energy and propellant efficiency, as well as lower maximum heights to minimize forces on impact. Differing thrust activation periods allow for varying traversed distances, allowing the robot to land close enough to precisely roll to the target location (Figure 27).

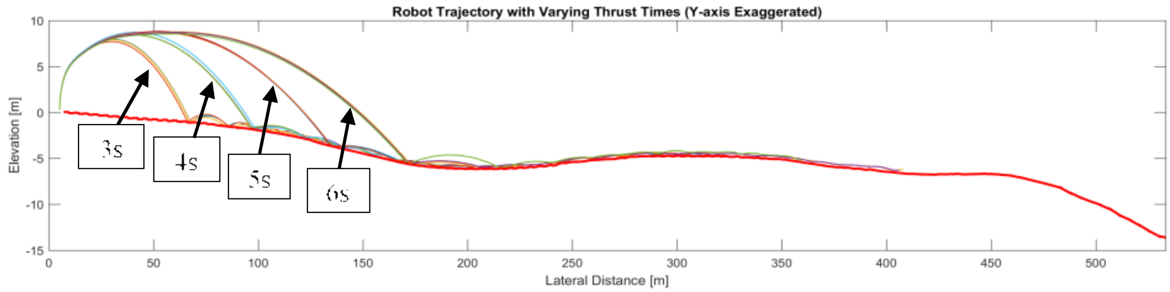

Figure 27. Robot flight trajectories with varying thrust durations (in seconds). Example lunar elevation profile is overlaid in red.

The results of the simulation on hills and craters can be seen below. With a starting position of (20,910), the robot traveled a total distance of 1062 m in 5 minutes using 2.6 kg of propellant to reach the goal location of (800, 190), which is in a large crater. From Figure 28, it can be seen that the robot performed four hops to cover the majority of distance, then used directed rolling to precisely reach the target destination. Between each hop, the robot re-evaluates its position – which may have been affected by unforeseen obstacles in the terrain, continued rolling after impact, or stochastic errors in the thrust-vectoring gimbal control – and makes corrections in the case of course deviations. This can be clearly seen in Figure 29.

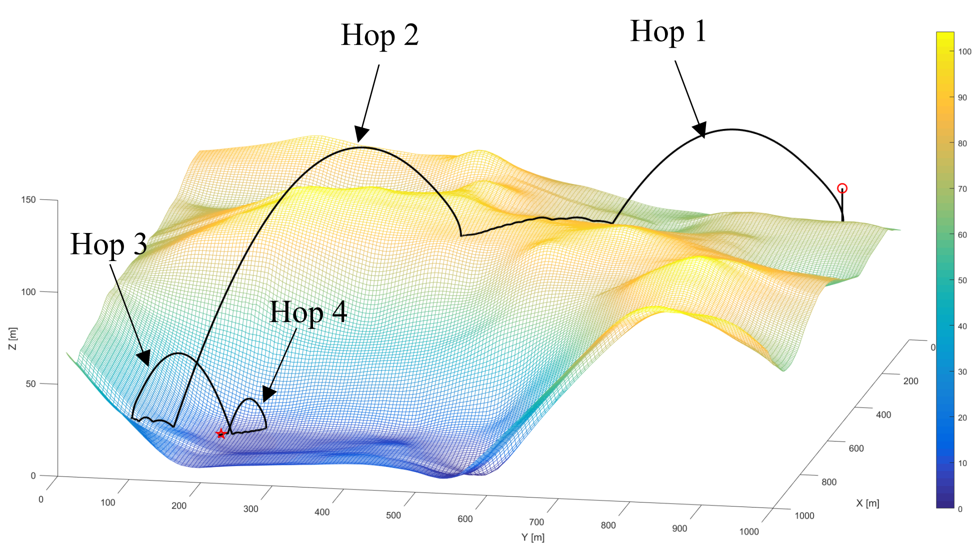

Figure 28. Robot center of mass 3D trajectory

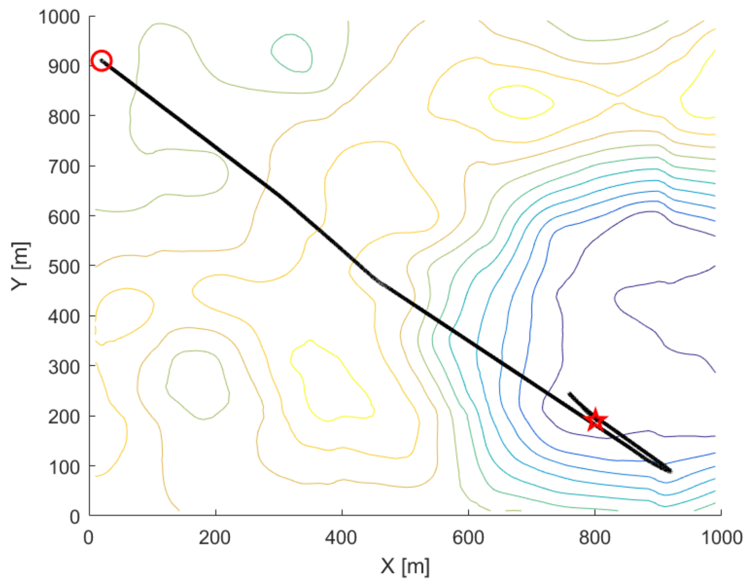

Figure 29. Top view of robot center of mass trajectory

3.2.8. Task 9: Simulation of rolling motion on an inclined surface (Q5-Q6)

One of the unique challenges that is encountered in tensegrity robotics is the development of policies for actuation. While most of the work in this area has been based on taking advantage of the deformability of the tensegrity structure, the methods that have been proposed vary in approach and complexity. These range from the relatively simple case of single-cable actuated mobility, to punctuated rolling through form-finding using dynamic relaxation, to complex dynamic gaits generated through evolutionary algorithms. In simulation these algorithms have been applied to mobility over level terrain, up inclined slopes and over rolling hills. In this project we have complemented these previous results by showing that rolling can be achieved up steep hills both in simulation and in hardware. The NASA Tensegrity Robotics Toolkit (NTRT) allows us to develop simulations on both flat and hilly terrain in order to investigate the potential of uphill locomotion using a six-bar spherical tensegrity robot on varying degrees of incline using a simple single-cable actuated punctuated rolling locomotion scheme. Using the results from simulations on a flat surface as a baseline, uphill rolling behavior will be characterized to illustrate the capabilities and limitations of this locomotion scheme.

3.2.8.1. Benchmark Rolling on Flat Surfaces

In order to provide context and baseline results for uphill rolling simulations, an NTRT simulation of punctuated rolling for a six-bar tensegrity robot with centrally located payload on flat terrain was first performed.

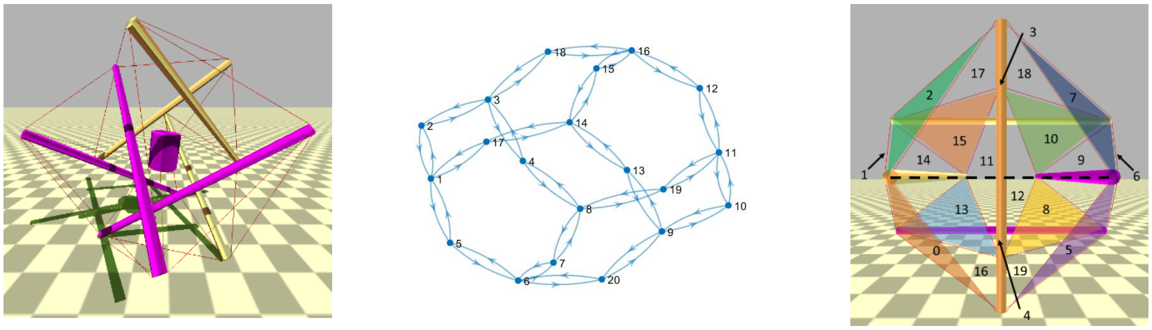

(a) (b) (c)

Figure 30. (a) The model of a 6-bar tensegrity robot with centrally located payload that is used in simulation. (b) Digraph representing surface connectivity on a 6-bar spherical tensegrity robot. (c) Surface number convention used in simulation and path generation.

In order to test the potential for uphill rolling motion, we decided to test the limits of a simple single-cable actuated punctuated rolling locomotion scheme. This means that during any forward locomotion phase, only one cable out of the 24 available is being retracted. This serves to deform the robot and move its center of mass outside of its current base triangle and thus roll, in a punctuated manner, to the next base triangle. By specifying a series of steps from one base triangle to an adjacent one, the robot is able to move in a zig-zag pattern in the desired direction. For both the flat and uphill rolling simulations, the repeating unit of the path is {15 13 0 5 7 10} where the numbers correspond to the face numbering convention specified in Figure 30(c), and the model parameters for the robot correspond to those of the SUPERball robot, which is being developed by collaborators at NASA Ames.

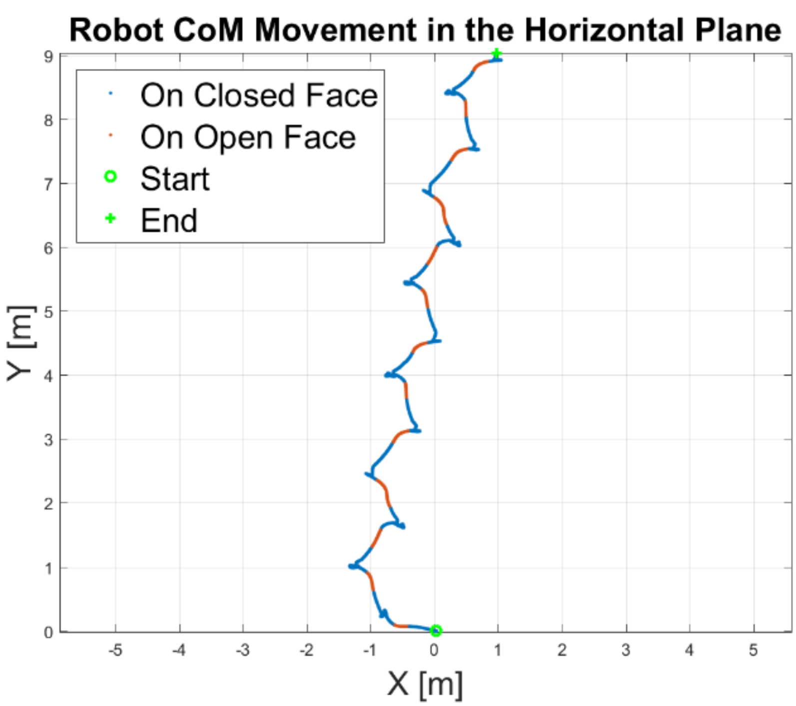

(a)

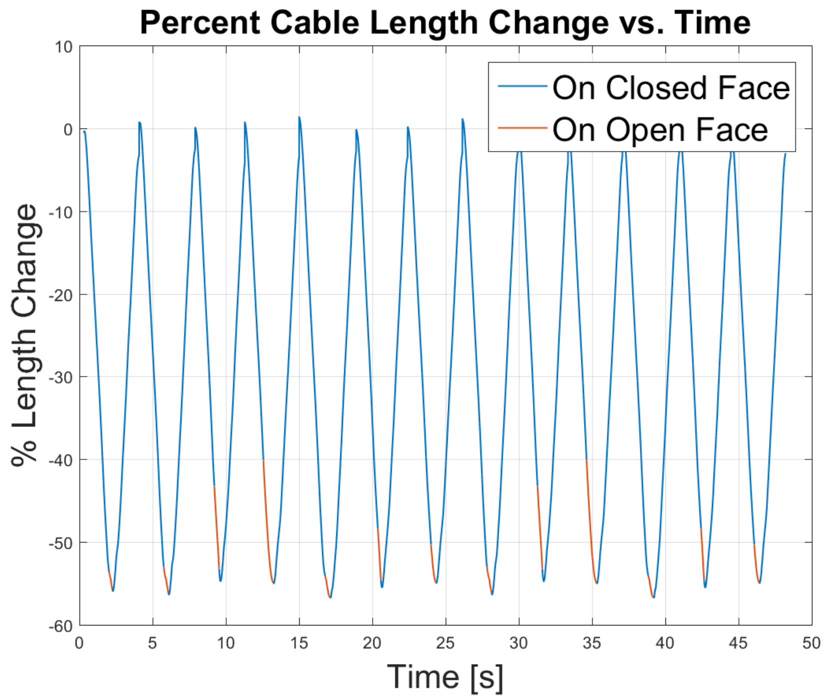

(b)

Figure 31. (a) Robot center of mass position on the horizontal X-Y plane. (b) Percent length change of the actuated cable during locomotion. (Note: closed faces correspond to triangular surfaces of the robot that are bound by cables on all three sides while open faces correspond to surfaces bound by cables on only two sides).

The movement pattern as seen in Figure 31(a) confirms the observations of single-cable actuated punctuated rolling on both the SUPERball and the UC Berkeley TT-3 robots and indicates that even with an open-loop path and simple locomotion scheme, the robot is capable of moving consistently in a desired direction on a flat surface with only one cable being actuated at a time. Furthermore, based on the cable retraction profile in Figure 31(b), it can be seen that the motions seem to occur in two sets of repeating triplets where two triplets make up one repetition of the path as specified earlier. While the symmetry of the spherical six-bar tensegrity structure suggests that each step during the punctuated rolling sequence should be identical, due to the inclusion of a cable-connected, centrally-located payload to the external structure, variance is introduced into the rolling steps. This is because the connecting cables are compliant and thus allow for relative motion between the payload and the external structure, thereby causing the overall robot center of mass to shift unpredictably during each step. This behavior, as will be seen in the next section, persists in uphill rolling and could be potentially used to augment current methods of contact surface detection.

3.2.8.2. Uphill Rolling on an Inclined Surface

In order to further evaluate the rolling performance of the six-bar tensegrity robot using the SUPERball dimensions, the rolling controller implemented on flat ground was also repeated on various inclined planes, up to 13 degrees of incline (Figure 32). Simulated sensor data was then analyzed to ascertain any significant relationship between actuation efficiency versus inclined angle.

The instant of initiation of rolling was observed for multiple steps for each angle of inclination by detecting when the central scientific payload of the robot recorded a projected velocity which exceeded a designated threshold. This threshold value was selected low enough to detect the initial moment of rolling as early as possible for each step but also greater than transient non-zero linear velocities from the central payload due to oscillations arising from natural compliance in the system, even when the robot is at rest. A significantly large velocity magnitude signified that the robot was in motion due to an unstable configuration and the cable actuation retraction length at each time of the initial rolling behavior was recorded.

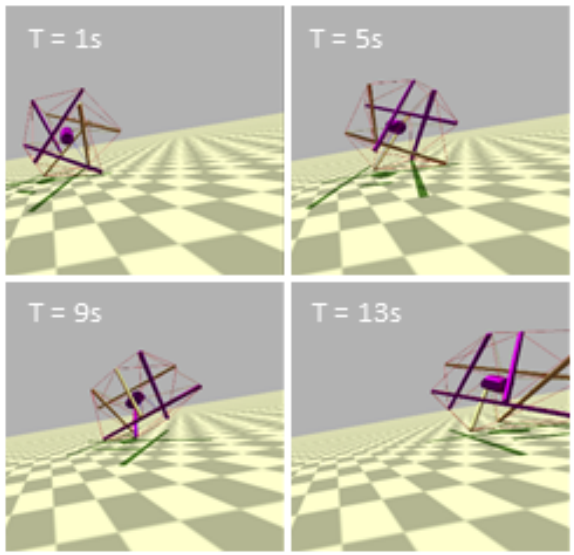

Figure. 32. 6-bar tensegrity robot rolling up a 10° incline with single-cable actuation

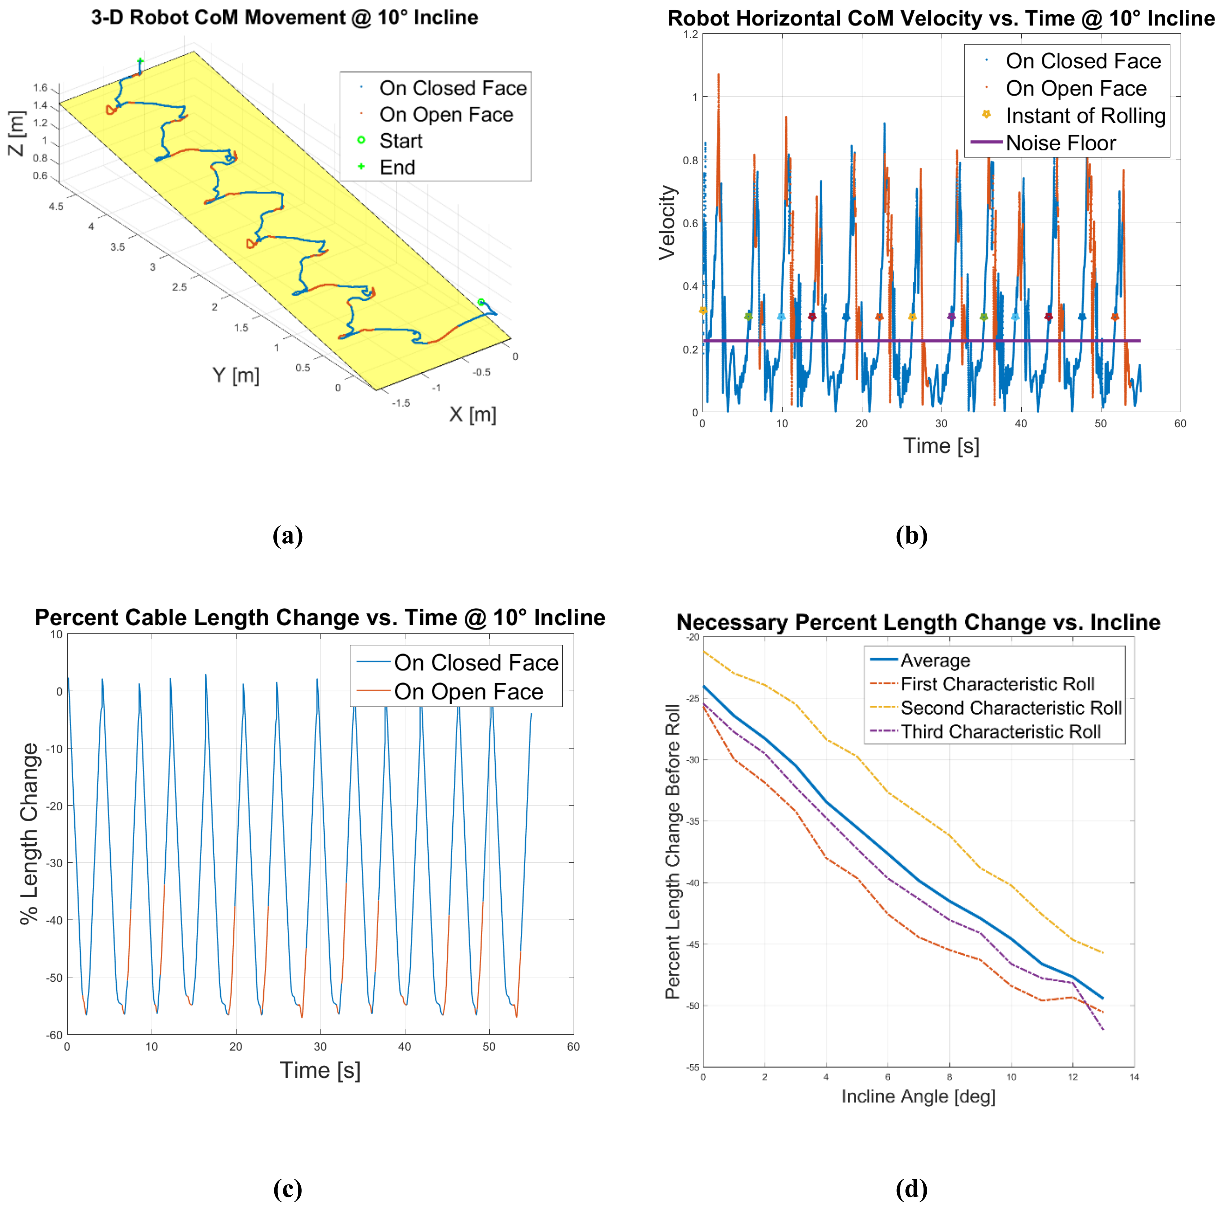

Figure 33. (a) Trajectory of robot center of mass position. (b) Robot center of mass velocity in the plane of incline. Stars indicate the point where the robot reaches its tipping velocity. Percent length change is recorded at the corresponding times. (c) Percent cable length change of the active cable during locomotion. (d) Percent cable length change required for tipping for the three characteristic rolls in each repeating triplet.

From the analysis results (Figure 33), a clear relationship between necessary cable retraction for a single step versus incline angle is apparent, with greater angles of inclination correlating to larger necessary percent retraction of the initial cable length before rolling behavior begins as seen in Figure 33(d). Interestingly, depending on the specific cable being actuated, the inclined angle has varying effect. As was mentioned earlier, the repeating unit of six steps in one direction can be separated into two groups of three "characteristic rolls" due to symmetry of the robot structure, with each group forming a repeated pattern of necessary cable retraction lengths before rolling. Although the extent to which the incline angle affects each cable varies from step to step, the average percent length change before rolling is initiated follows the same general linear trend. From this, it is clear that climbing steeper hills leads to greater power consumption for the robot, motivating energy costs which are now more definitively quantifiable and clearly dependent on angle of inclination.

3.2.9. Task 10: Manufacture and assemble robot ball hardware. (Q6-Q8)

We built a new version of the rapidly prototyped tensegrity robot, TT-4 (fourth generation prototype), which features more robust hardware architecture and a rod length of one meter. This allowed us to experiment with the robot while carrying a larger payload at the center. In parallel with TT-4 development, we have developed a new tensegrity prototyping platform using modular, elastic lattices.

We have experienced the difficulty of assembling various tensegrity structures since the members are not in balanced compression and tension until the structure is fully assembled. In the intermediary steps of assembly, forces are unevenly distributed thus the structure is difficult to constrain. Also, it is easy to make mistakes during the assembly, such as connecting the wrong tension and compression members.

The idea for an elastic lattice came from examining an assembled six-bar tensegrity structure and conceptualizing how the tension members could be deconstructed into elastic modules. We prototyped each elastic modules by laser cutting sheets of silicone rubber into modular elastic lattices shown in Figure 34. A six-bar tensegrity structure can be constructed using eight triangular elastic lattices and six rods. Figure 35 illustrates the step by step sequence of constructing a six-bar tensegrity structure using the new modular, elastic prototyping platform. We decided to use thin-walled aluminum rods as the compression elements in our static tensegrity prototype and use 3D printed endcaps as the connection between the modular elastic lattice and the aluminum rods.

The result is a tensegrity structure that can be built in a few minutes by a single person. With this novel prototyping platform, we were able to build the $TT\text{-}4_{mini}$ shown in Figure 36. It is a small version of TT-4 with 25 cm rod length and six actuators (controllable via smart-devices), that can

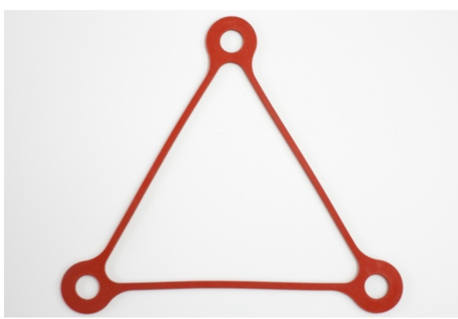

Figure 34. Modular elastic lattice prototype made with 60A durometer rubber.

allow us to rapidly experiment with different actuation schemes and smaller test environments.

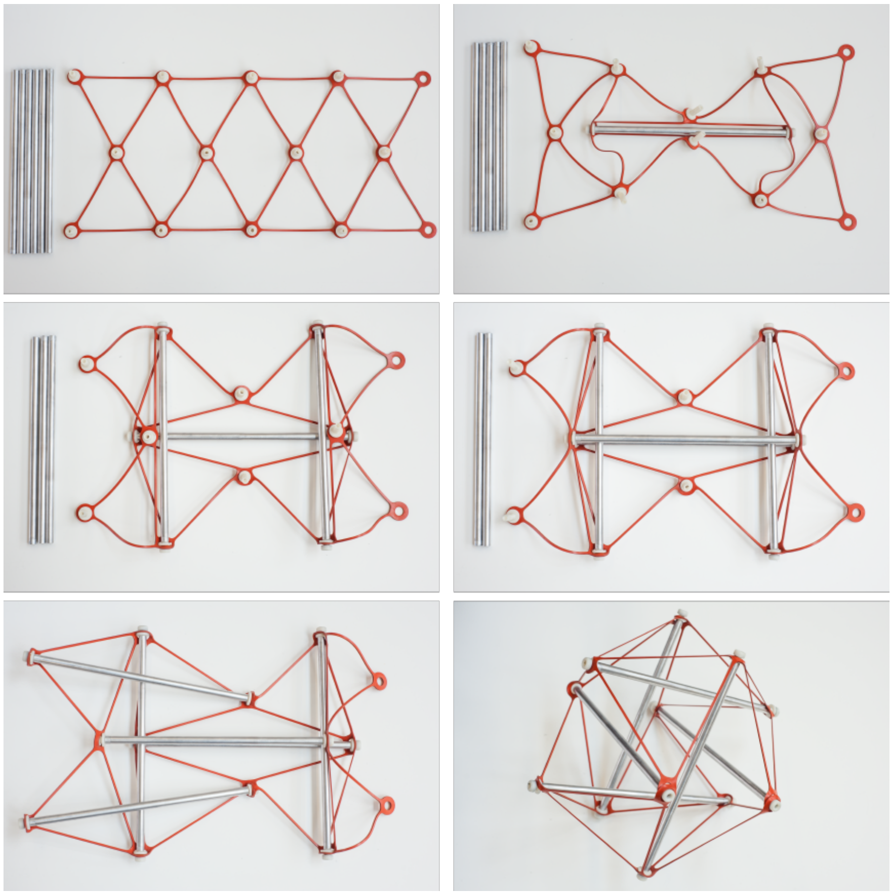

Figure 35. Step-by-step assembly sequence of a 6-bar tensegrity static model.

Figure 36. TT-4$_{mini}$ Prototype.

We also further refined our TT-4_{mini} hardware prototype for various rolling experiments. The two key experiments to test with the TT-4_{mini} were rolling on a level surface and uphill. As there have been no previous hardware tests on uphill climbing using spherical tensegrity robots, achieving uphill climbing was critical to show proof-of-concept for tensegrity robots' ability to navigate on hills on the Moon terrain. The development of policies for actuation is one of our unique challenges. A simple actuation scheme was chosen to implement on TT-4_{mini} for these experiments, in which only one actuator was activated at a time. With single-cable actuation, we can understand the base line mobility of a tensegrity robot. In our previous work and simulation (Figure 31), we have shown the single-actuation scheme can result in forward locomotion on flat ground with six-bar tensegrity robot (Figure 37).

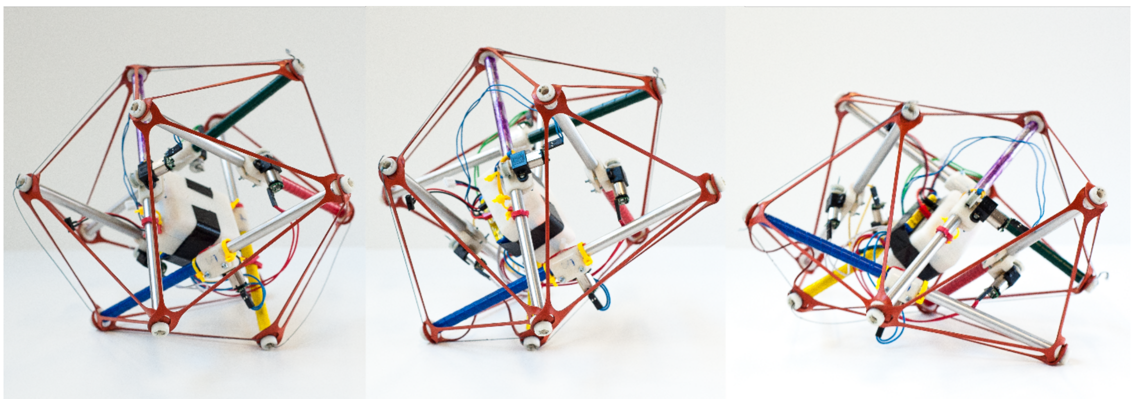

Figure 37. TT-4_{mini} prototype rolling on a flat surface with single actuation.

It was found in Section 3.2.8.2 that a spherical tensegrity robot could climb up an incline of approximately 10 degrees in simulation using the single-cable actuation policy shown in Figure 33. TT-4_{mini} prototype's first experiment was to perform punctuated rolling on flat surfaces, accomplished through shifting its center of mass by deforming the base triangle with a single cable contraction. With the single-cable actuation policy, the robot reliably performed punctuated rolling in a straight line on a level ground, as shown in Figure 37.

In order to test uphill climbing, we constructed an adjustable testing platform that allows the incline surface to be changed to the desired angle. We ran several trials in which we incrementally increased the incline angle after the TT-4_{mini} was able to perform a complete six-step rolling sequence at the set incline. We were successful in performing uphill climbing up to 13 degrees with a single actuation policy. Figure 38 shows the TT-4_{mini} climbing uphill. This is first time an untethered spherical tensegrity robot has performed uphill climbing through hardware experiments.

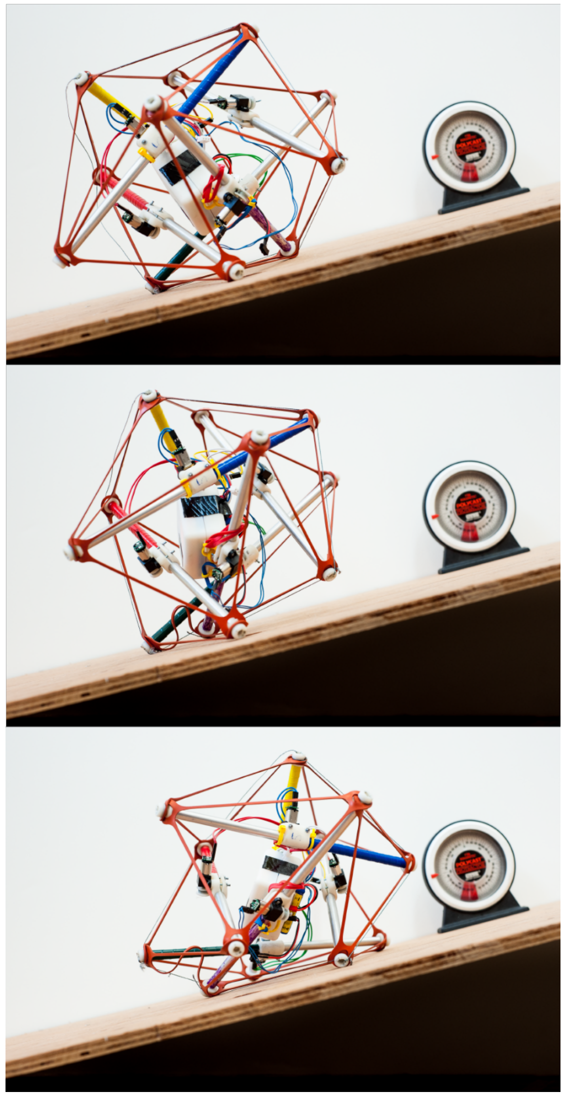

Figure 38. TT-4$_{mini}$ prototype climbing up a 13-degree incline surface with single actuation.

3.3. Tasks for the Final Award Year 3 and No-Cost Extension Year 4

3.3.1. Task 10 (continued): Manufacture and assemble robot ball hardware. (Q10-Q15)

3.3.1.1. TT-4 hardware testing

Single cable actuation was successfully tested on the TT-4 prototype, resulting in the first punctuated rolling motion of our largest, fully-actuated tensegrity robot up to that point. However, throughout the testing process, many issues and shortcomings in the design of TT-4 were uncovered. In order to address these issues and drastically improve the robustness of the robot, a new version, TT-5, was designed with similar size and form as TT-4. In addition to a revamped power monitoring and distribution system, it features an inductance sensor to measure spring tensions in real time. The sensor can be used in robot form-finding, and it provides necessary feedback for the advanced control algorithms that were investigated in Task 12. In addition, new motors have been selected which provide additional torque - this will accommodate the larger tensions that may be experienced within the system due to the actuation of multiple cables at the same time. Furthermore, we developed a ROS (Robot Operating System) network to streamline controller deployment to simulation and to the robot hardware. This addressed the need to move the system onto an established and widely-used robotics platform, further improving system robustness.

3.3.1.2. TT-5 Hardware Assembly

Assembly and integration of the electrical hardware and mechanical components for the TT-5 prototype of the hardware was completed in Q12. We first evaluated the simpler single-cable actuation policy to test the robot before moving on to more advanced multi-cable actuation policies for more robust locomotion. As described above, the TT-5 has a similar form to TT-4 but features more robust electronics and the capability to include tension sensing using induction sensors, allowing for more feedback for advanced control algorithms being investigated in Task 12. As well, improvements in wireless communication between the six microcontrollers on each rod of the robot and the master controller allowed for higher frequency of sensor readings and motor commands which help with more effective feedback control.

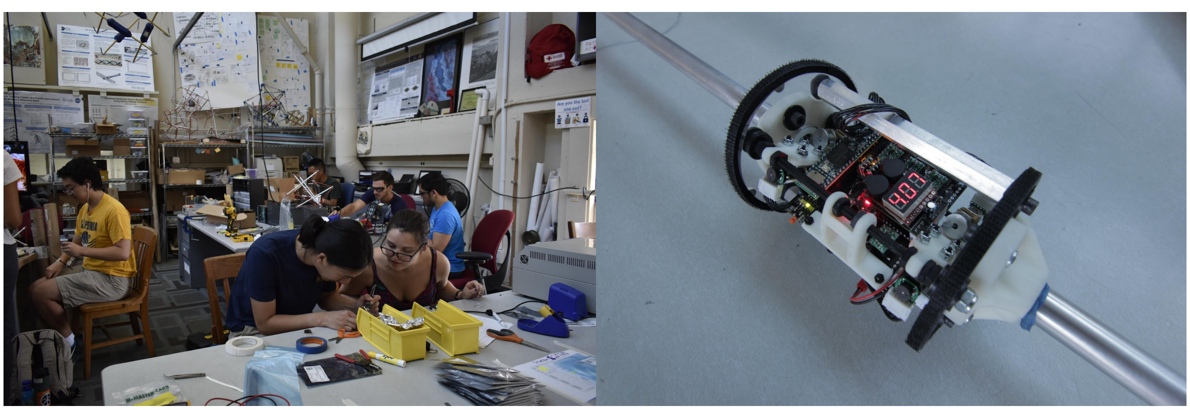

Figure 39. First TT-5 assembly by ESI team and completed capsule electronics

3.3.1.3. TT5-Meso Hardware Assembly

Construction of the TT5-meso prototype was completed in Year 4. It features a lighter weight structure that can be carried by a drone (see Fig. 40), in addition to TPU 3D-printed end caps and strain-stiffening latex tubing for better impact deformation characteristics. Custom-designed modular motor gearbox assemblies also allow for varying the number of actuators (e.g., adaptable from 6- to 12-motor actuation) to easily test different control policies rapidly-developed in software. A central payload containing the custom modular electronics designed in the lab is attached using 12 passive elastic elements connected to each of the rod ends and is used to control all actuation on the robot.

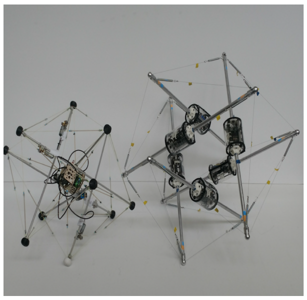

Figure 40. Completed TT-5 meso (left) next to the larger TT-5 (right)

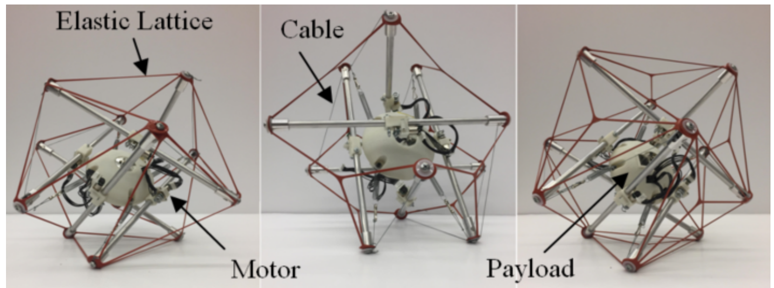

Figure 41. TT5-mini with different lattice topologies. From left to right: delta, Y-star, and hybrid lattice structures

3.3.1.4. Final Passive and Active Tensegrity Structure Assembly (Squishy Robotics)

The most significant technical contributions in Q15 (no-cost extension year) was towards the design and prototyping of two new robot systems in collaboration with *Squishy Robotics, Inc.*, a spin-off of our ESI research to commercial tensegrity robots for disaster response on Earth. ESI funds were used to partially fund MEng and doctoral research at UC Berkeley. *Squishy Robotics* provided access to their new hardware for testing. Due to our previous difficulty in reconciling the stiffness of the tensegrity structure with the power needed to actuate the cables, the new robot was divided into an "active" system focused on locomotion and control policy testing, and a "passive" system for drop test experiments, with the goal of better understanding their independent behaviors and ultimately merging the two prototypes in the future. A new control board was developed for the active system, and features a wide array of scientific sensors, long-range radio communication, and versatile motor driver capabilities. The passive system went through multiple iterations of its structural design after repeated drop tests and can withstand falls of up to 400 ft (122 m), the highest drop legally allowable from a drone per FAA regulation. Previously, the highest successful drop was only five meters, using UCB prototypes shown in Figure 42.

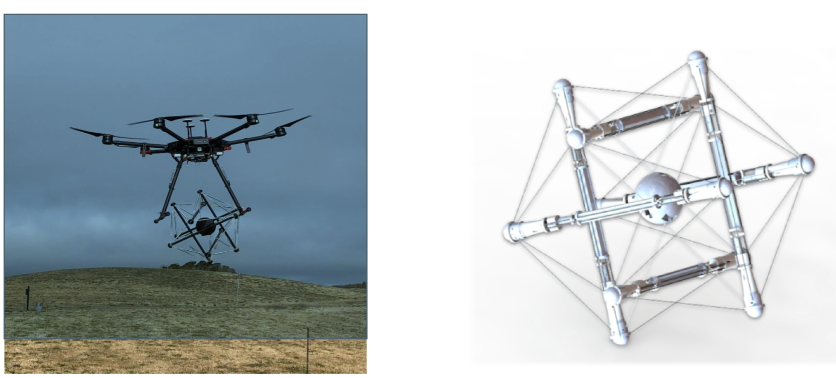

Figure 42. *Left*: Squishy Robotics' passive system prototype dropped from a drone. *Right:* 3D CAD Model of the active system

Both the Squishy Robotics' passive and active systems feature a new light-weight structure using carbon fiber tubes. The passive system features stiffer spring connections using a new cable-routing method that places the springs and nylon cables in parallel, effectively increasing the stiffness of the elastic elements, as well as an updated design of TPU (thermoplastic polyurethane) 3D-printed endcaps. The payload is also printed using this flexible TPU material to provide further protection during impact (see Fig. 42). To avoid the stress concentrations that previous iterations created by having an electronics capsule in the center of each rod, the control board's form factor was modified to fit entirely inside the tube itself for the active system. As before, spooled cables controlled by the onboard motors exit out of the tube endcaps and connect with the rest of the actuated system.

3.3.2. Task 11: Test control algorithms on tensegrity ball hardware. (Q9-Q15)

In Q9, progress was made on the control of 6-bar tensegrity robots. Specifically, the following three control algorithms and actuation schemes were tested either in simulation, on hardware, or both, (1) Two-cable rolling on inclined surfaces, (2) Dynamic rolling using Model Predictive Control (MPC), and (3) Deep reinforcement learning for tensegrity locomotion. In addition, there has been ongoing work with implementing the A* path-planning algorithm both in NTRT and in hardware experiments. This will be combined with the aforementioned control strategies to create a full mission profile using robust path-planning, rolling, and hopping methods.

1. **Two-cable rolling on inclined surfaces**

Having reached the limits of inclined locomotion for the single-cable actuation policy (13-degree) in our previous report, the following actuation policies were explored based on the robot's ability to actuate multiple cables simultaneously or in alternating order:

- **Simultaneous actuation policy**: Similar to single-cable actuation, except the next cable contracts *while* the current releases, allowing for more steps to be made in less time.
- **Alternating actuation policy**: To preserve a low center of gravity during uphill rolling, the next cable is contracted *before* the current is released.

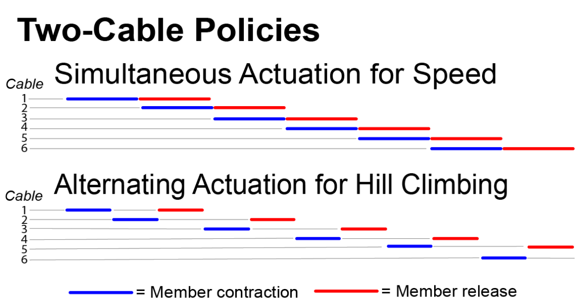

Figure 43. Visualizations of the two-cable actuation policies. Each row corresponds to one cable, and each policy can be repeated indefinitely.

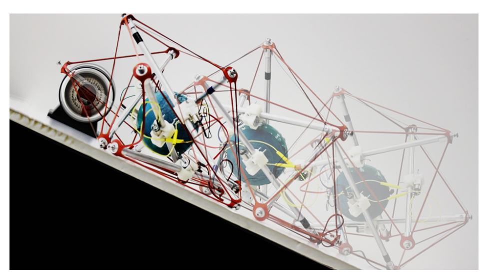

Figure 44. TT-4mini prototype performing punctuated uphill rolling on an inclined surface of 24-degree.

The two-cable actuation policies, as described above, were implemented and tested in NTRT. These simulations demonstrated vast improvements in incline locomotion stability as well as average speed, with the robot able to navigate inclines up to 26° using alternating two-cable actuation and 24° using simultaneous two-cable actuation. The significant performance improvements achieved with the two-cable policies are a result of the center of mass being consistently lower throughout the actuation sequence of the robot. Both two-cable policies maintain at least one cable in contraction at all times, thus keeping the robot in a forward leaning neutral stance with four points of contact with the ground.

The TT-4_{mini} was able to leverage alternating two-cable actuation to reliably climb a 24-degree (44.5% grade) incline, far outperforming the robot's previous capability of 13-degree (23.1% grade) climbing via single-cable actuation. In addition, there was also a significant improvement in average velocity. The traditional, single-cable actuation policy traveled a distance of 3ft on a 10° incline with an average velocity of 4.1 cm/s. However, when performing two-cable simultaneous relaxation and contraction, the robot could travel the same distance with a 10-trial average velocity of 6.18 cm/s, achieving an increase of nearly 48% over the single-cable baseline.

2. Dynamic rolling using MPC

Motivated by the greater capabilities afforded by two-cable actuation policies for uphill rolling, one approach for advancing the actuation policy for locomotion - from “punctuated rolling” behavior to a continuous dynamic rolling behavior - is using model predictive control (MPC) in combination with more robust contact surface detection (CSD) through a neural network. Our updated contact surface detection allows the robot to determine its orientation during dynamic movements using accelerometer and cable tension inputs for inclines up to 40 degrees (previous best was ~24 degrees). Model-based optimal control enables the generation of actuation policies which are locally optimal according to a user-defined objective function (e.g. speed, energy-efficiency, etc.). This approach uses a linearized model of the robot to solve a constrained optimal control problem to find reference trajectories of cable restlength inputs that move the robot in an energy-efficient manner to initiate a rolling behavior. Robust CSD using neural nets trained on simulation data is then used to blend the individual “local” one-step policies generated using MPC, depending on the current pose of the robot. To address the long solve-times of the optimization problems, in Q10 neural nets trained using supervised learning on the offline-generated trajectories have been utilized to implement real-time control of the robot in simulation. Using this method, multi-cable actuation policies for rolling motions with the robot have been successfully simulated in NTRT simulations. Preliminary simulation results of MPC for rolling locomotion of a six-bar spherical tensegrity using physical parameters of the new TT-5 hardware show excellent results with rolling speeds of up to 30 cm/s when motors are able to actuate cables at a linear velocity of 5 cm/s. Furthermore, simulation results show that if we can improve motor speeds such that cable lengths can change at 15 cm/s, rolling velocities of ~70 cm/s are attainable. Finally, although MPC has many advantages for optimal control of the robot, computational time is a significant barrier to real-time online feedback control. To that end, we incorporate deep models such as neural networks trained through imitation learning on the MPC policy to provide online control in hardware.

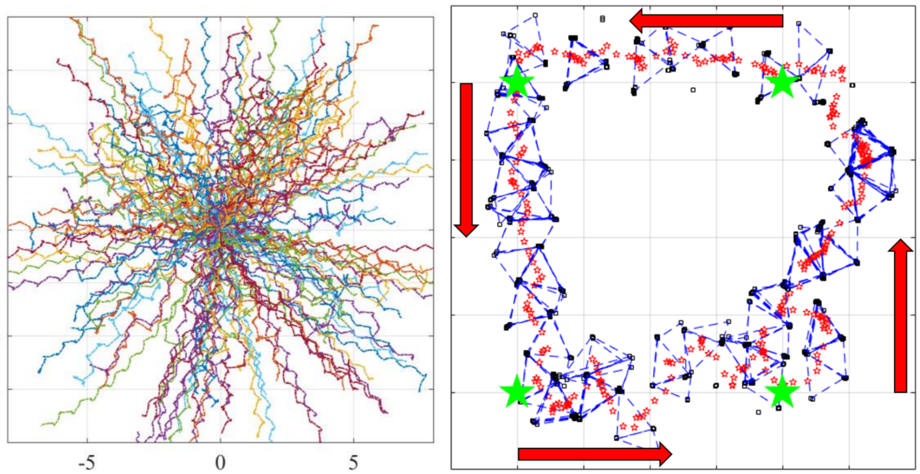

The plots above present some of the results of using MPC in combination with supervised deep learning. Using the 150+ generated optimal state-action trajectories using MPC in randomly sampled directions (above, left), we trained a contextual neural network policy with a directional contextual input to control directed-rolling locomotion of the robot. An example squared-shaped rolling trajectory for the robot is depicted (above, right), which shows the simulated robot's traversed path in blue, through sequential supporting base polygons that come in contact with the ground.

In Q14, we extended some of the progress on rolling trajectory locomotion control through model predictive control and supervised learning. The simulation conditions considered a cost function that favored speed in the horizontal direction, with a parameterized penalty for deviation from the zero y-axes. We ran several scenario simulations with the geometry of the TT-5 prototype, using a random starting face and 10% perturbation in initial rest length. Fig. 45 shows one example result with over 150 runs demonstrating that precision of rolling trajectories with the six-bar spherical tensegrity topology with 24 cable actuation is inherently limited by the geometry of the structure itself, as evidenced from the zig-zag rolling path. In comparing simulations with different assumptions, the control strategy realizes a range of trade-offs for optimal behavior, depending on the parameters in the cost/reward functions used. Using this framework, we generate and evaluate control policies for a variety of different actuation schemes (e.g., 6-cable or 12-cable actuation) over different physical prototype geometries.

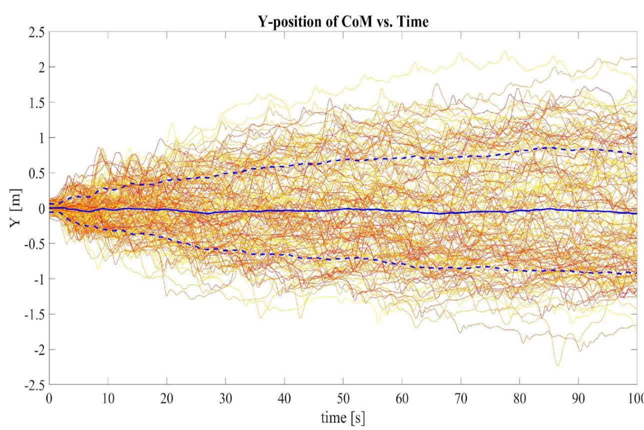

Figure 45. Simulation results from rolling trajectories generated using model predictive control. Y-axis represents deviations from the desired straight-line path.

In Q15, one paper was presented on the algorithms and simulations described in this section: *"Multi-Cable Rolling Locomotion with Spherical Tensegrities using Model Predictive Control and Deep Learning", (B. Cera and A.M. Agogino), Proceedings of the 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Oct. 1-5, 2018."*

3. Deep reinforcement learning for tensegrity locomotion

By adopting a framework called *guided policy search* for tensegrity locomotion, the tensegrity robot can learn locomotion skills autonomously by imitating optimal controllers. Policy algorithms search over a parameter space to generate an optimal control law under specified cost functions. Supervised learning can be used to optimize complicated policies by leveraging the data from either a human or computational "teacher". Guided policy search uses optimal controllers as "teachers" to train a neural network. Rather than the standard optimal control techniques, where an optimal control law is generated from a given initial state to the end of the horizon; guided policy search seeks optimal solutions only in some regions of the whole state trajectory where simple linear models can be applied, and utilizes supervised learning to mimic the underlying pattern from the data generated by these optimal controllers. This algorithm has been successfully demonstrated in NTRT, where the robot learns locomotion policies to maximize rolling speed.

In Q10, we applied the guided policy search method to the TT-4 robot within the simulation platform NTRT. This technique allows the robot to autonomously acquire locomotion skills with only a few trajectory rollouts. Results show that the TT-4 robot was able to maintain a stable speed of 0.225 m/s on flat terrains. We also built a software platform based on Robot Operating System (ROS) such that control schemes can be deployed in simulation and to hardware in a seamless manner. Taking advantage of such a flexible software structure, we can obtain actuation policies in simulation and adapt them to hardware testing in a highly efficient process.

In Q11, we also demonstrated that in the domain of tensegrity robotics, it is possible to efficiently learn end-to-end policies even in very low-dimensional observation spaces. From our results, we further conjecture that neural network policies are unique in their expressive power to transform components of these extremely low-dimensional feature representations into higher-dimensional actions without having to resort to modeling the dynamics of the system. Furthermore, we validated their expressive power by showing that deep reinforcement learning outperforms other tensegrity locomotion control policies in simulation in limited-sensory-input environments. We showed preliminary results with one such model on rough terrains characterized by highly discontinuous dynamics.

4. A* path-planning implementation and experimental validation

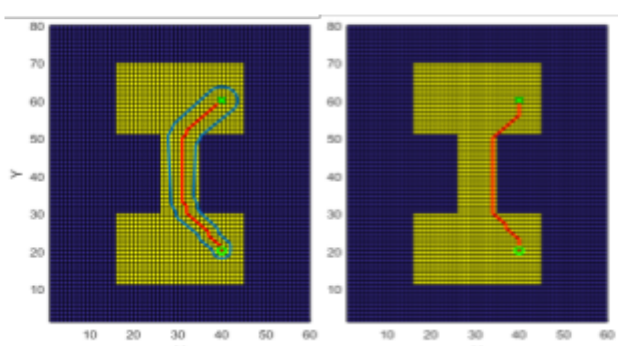

Figure 46. (Left) Path-planning with growing uncertainty. (Right) Path-planning with no uncertainty. Start node is denoted by the green circle, the goal node is denoted by the green 'X.'

Implementing an A* path-planner on a system that has two methods of locomotion has produced novel results which have already been discussed in previous quarters. However, due to the complex restitution dynamics of the robot as well as unpredictability in terrain, there exists a high level of uncertainty especially in the thruster-based hopping motions of the robot that have yet to be addressed. In Q10, progress was made in quantifying the uncertainties associated with the bouncing behavior of the robot. This was done through modeling the robot's stance dynamics and performing data-based parameter estimation. The variation in model parameters is then quantified in order to create a probabilistic model of the robot's bouncing behavior and to provide an estimation of where the robot is likely to come to rest. This model can be incorporated into the path-planning process where risk assessment can be performed based on belief of the robot's location and the surrounding terrain/hazards. Preliminary results where only rolling is considered have shown promising results where the path-planner is able to avoid potentially hazardous terrain features in order find a more conservative route, thereby increasing its chances of successful navigation. This is illustrated in a simple example above (Figure 46) where the path-planner chooses to cross a bridge through the middle, albeit at a penalty to path cost rather than hug the edge which, due to uncertainty in the robot's motions, could lead to a disastrous fall.

The main contribution in Q12 was in the modeling of uncertain restitution behavior in path-planning. In addition, we completed an in-depth analysis of what is learned by *Mirror Descent Guided Policy Search* using neural networks to reveal a reduced-order model for spherical tensegrity locomotion.

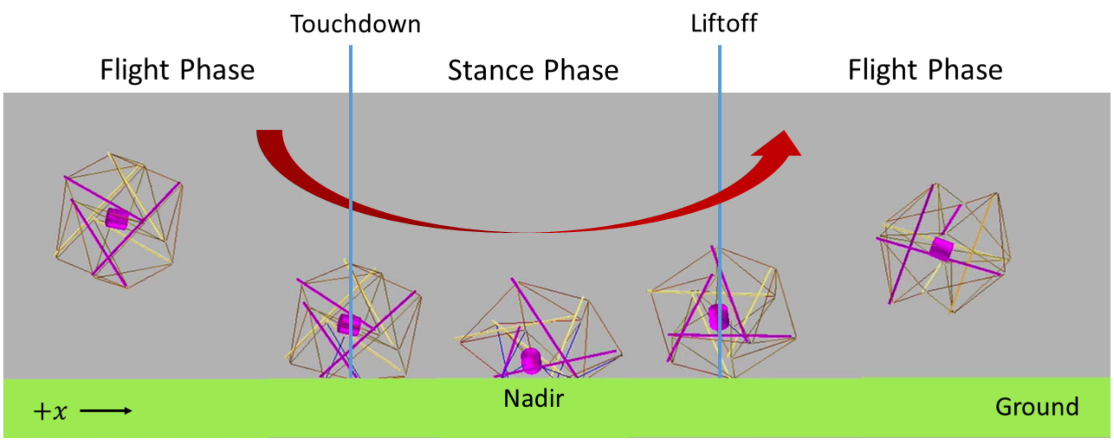

Figure 47: Restitution behavior of a 6-bar tensegrity robot.

In order to safely use the thruster based hopping motions that were investigated in earlier work, the restitution behavior that occurs when the robot contacts the ground must be accounted for. While expressive rigid body models exist, modeling such behavior in rapid rollouts is difficult due to the high dimensionality of these models. What is proposed then is a simplified hybrid model which treats the restitution behavior of a 6-bar tensegrity robot as a stochastic process from which landing zones can be predicted. Specifically, there are two phases (Figure 47). The first phase is the flight phase which is modeled using projectile motion. This phase then transitions into the stance phase at touchdown, which is modeled as a Gaussian Process (GP). The GP model using impact rollouts from NTRT to estimate the touchdown to liftoff relationship and performed well as seen in Figure 47.

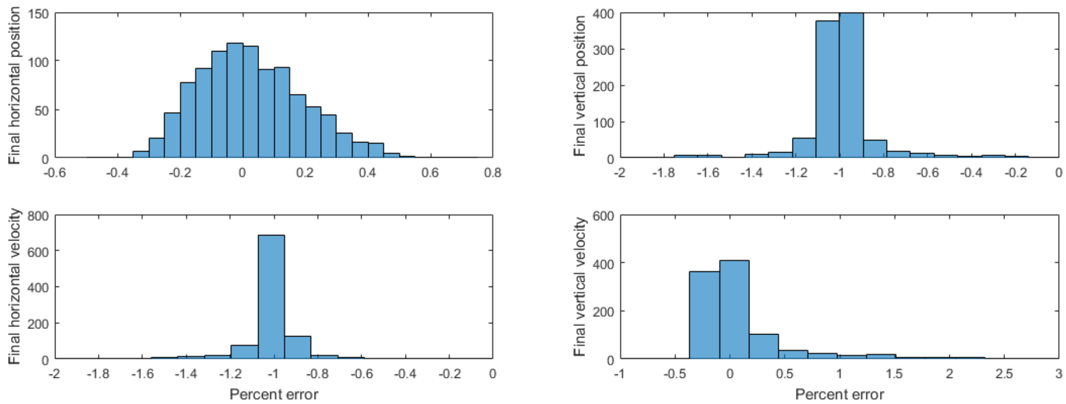

Figure 48: Percent error between predicted and actual final states using the GP stance mode.

Ultimately, this model was incorporated into an A* path-planner to achieve risk-aware and conservative paths. This was done by weighing the cost of the landing zones by probability mass in order to for the path-planner to avoid treacherous terrain. An example of this can be seen in Figure 49.

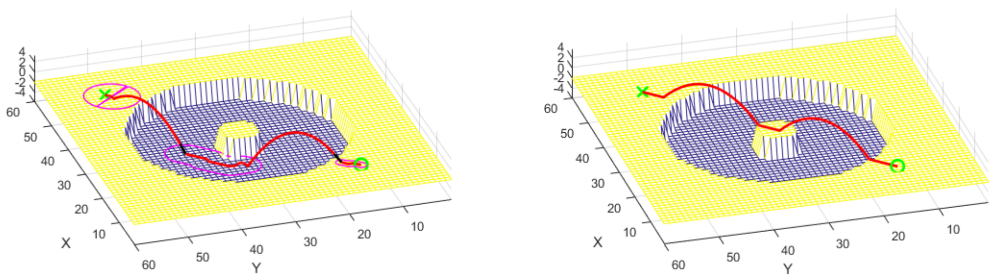

Figure 49: Comparison of solution paths with robot restitution and uncertainty (left) and without (right).

3.3.3. Task 12: Test payload protection under mission impact profiles. (Q9-Q15)

In order to investigate the impact profiles and payload protection characteristics of tensegrity robots, in Q9 multiple testbeds and prototypes were developed in order to perform a variety of characterization experiments. These include, (1) TT-4 impact vertical drop test, (2) TT-4 mini launch test, (3) Payload connection characterization under various thrust profiles, and (4) Impact characteristics of 12-bar tensegrity structures.

3.3.3.1. TT-4$_{impact}$, TT-5$_{impact}$, TT-5$_{meso}$ vertical drop test

A full-scale model of TT-4 (1 m rod length, 5.5 kg mass without payload) was developed in Q9 to measure the forces and acceleration that the robot experiences under external impacts. It was dropped from a height of 3 m, with two IMUs placed on the payload and at the center of one rod for collecting acceleration data. During impact, the maximum acceleration of the payload was found to be 12 g, while for the center of the rod it was 52 g, indicating strong impact absorption characteristics in the robot.

A new drop test circuit board was designed in Q10 that enabled us to log IMU data locally onto an SD card, rather than wirelessly transferring it to a computer. This allowed us to gather much more accurate data than before: when we were wirelessly streaming the measurements, we were limited by the low frequency of the transmissions and likely missed key data points during the impact's short timeframe. Furthermore, the drop test model was augmented to allow for easy adjustment of cable tensions as well as payload mass.

In continuation with the plan set forth in Q10, a drop test plan for the tensegrity robot was delineated in Q11 to help understand robot survivability under given planetary conditions. To this end, many drop tests were completed with the *TT-4$_{impact}$* while varying physical parameters of the robot such as spring stiffness, cable tension, and robot orientation. To conduct the drop tests under realistic conditions, we are also built a medium-sized robot named *TT-5$_{impact}$*, which helped evaluate impact effects on electronics, and we designed a new drop test mechanism by adapting a drone or an aerial vehicle. A lightweight and reliable dropping mechanism was designed that allowed drones to travel large distances with repeatable and consistent drop test results. Furthermore, we examined the stress-strain relationship between the individual members of the robot structure once the tensegrity has impacted the surface. A failure analysis was conducted by analyzing existing data and understanding the root cause of failure of every component of the robot.

To further evaluate the payload protection capabilities of the tensegrity structures, we completed several rigorous drop test studies in Q14 on both the *TT-5$_{meso}$* (see Fig. 40) as well as the *TT-5$_{mini}$* and its variety of different elastic lattice designs (see Fig. 41) to see how different design choices affected the impact-resilience of the robots.

In summary, from the *TT-5$_{meso}$* experiments, we found that the payload came into contact with the ground at relatively high velocities, and that orientation of the tensegrity had a large influence on how well the central payload was protected, as expected. A key takeaway from the experimental results is that while tensegrities do offer significant benefits towards impact protection of a payload, incorporating a more holistic approach utilizing additional methods of impact-resistance (e.g., sacrificial components which fracture or energy-absorbing dampers) would be a promising direction.

Rigorous drop test experiments from 1 m and 5 m drop heights were also completed using the *TT-5$_{mini}$* and its different elastic lattice structures. In modifying how the rods were interconnected by the lattices, we altered the possible load paths for local collision forces upon the tensegrities' impacts with the ground. Drop test results shown in the table below summarize our findings for drops of two levels of stiffness for each lattice topology compared to a base case of dropping a solid block of an equivalent mass with no inherent impact-dampening characteristics. Note that the less stiff topologies were ignored for the 5 m drop because the payload consistently impacted the ground from this higher drop height. This is a result of both the larger deformations of the lower stiffness lattice as well as the relative size of the payload in comparison to the tensegrity structure. While we expected that higher structural stiffness would result in higher magnitude accelerations, what we observed was that, for this experiment, higher stiffness elastic lattices resulted in relatively lower accelerations, potentially due to the payload contacting the rods if the structure was not sufficiently stiff enough. Furthermore, landing on an open triangular face of the tensegrity's outer-shell showed better payload impact characteristics than landing on a closed triangular face due to the increased stiffness in the overall system, resulting in the payload not contacting the ground. In all cases, the tensegrity systems were shown to reduce the maximum acceleration experienced during a landing event, though the lattices alone may not be sufficient as a structural solution.

Topology	1 m Drop (G's)	5 m Drop (G's)
Base case: solid block	114.9	228.0
Y-star, closed face	82.8	-
Y-star, open face	42.9	-
Delta, closed face	59.0	-
Delta, open face	46.5	-
Hybrid, closed face	57.6	-
Hybrid, open face	40.9	-
Delta (double stiffness), closed face	47.6	175.1
Delta (double stiffness), open face	39.0	131.0
Hybrid (double stiffness), closed face	45.9	194.7
Hybrid (double stiffness), open face	38.1	157.8

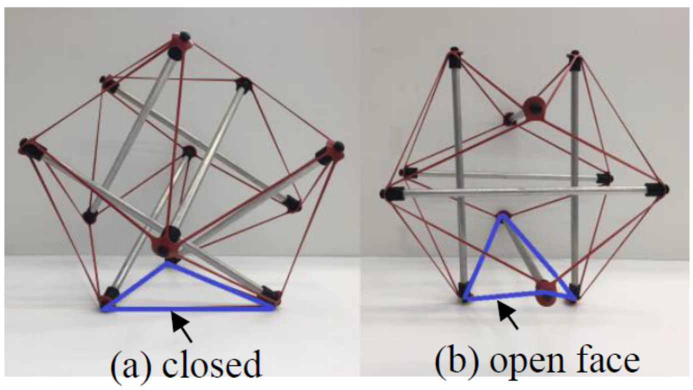

Figure 50: Tensegrity robot landing onto base triangle (a) and onto bars (b).

To further evaluate the payload protection capabilities of the tensegrity structures, in Q15 we completed drop test studies on 19 passive system iterations to see how different design choices addressed potential failure modes and improved its impact-resilience. In the table below, the stiffness refers to the tensile stiffness of the elastic elements in the cable direction, not the overall effective stiffness of the structure.

Model	UCB TT-5 Passive	UCB TT-Mini Passive	UCB TT-5 Meso Active	Squishy 15 Passive
Weight (w/o payload)	5.5 kg	0.39 kg	2 kg	0.62 kg
Stiffness	513 N/m	329 N/m	271 N/m	508 N/m
Rod Length	1 m	0.305 m	0.714cm	0.50 m

The initial Squishy Robotics prototypes used a stiffer version of the silicon rubber lattice structure reported in Q14 (see Section 3.2.9), with improved payload protection and impact absorbing endcaps. The tests were evaluated from drop videos for failure points and causes. Drops originally started with 10 m and moved to 20 m after the rods proved to strengthen the structure. After multiple 20 m drops, the elastic lattices made of silicone rubber were determined to be the weak point, as the strain rate was too large (only survived 3-5 drops before failure). Subsequent prototypes switched to steel springs for the increased stiffness as well as damage resistance. Despite this, we were still able to observe the springs' end hooks deforming during impact (noticeable deformation after ~10 drops), which lead to accumulated damage and unreliable performance in the repeated hops and landings found in the robot's ideal mission scenario.

Overall, the newer passive system prototypes performed much better in the drop tests than the previous versions of the robot did, due in part to the stronger carbon fiber tubes, stiffer spring connections, and the robust endcap and cable connector designs. For higher drop heights there were far fewer "catastrophic" failures of cables snapping, endcaps breaking, or rods buckling, and the payload itself rarely experienced any direct impact from contact with the ground or one of the structure's rods.

ANSYS 15.0 was used to simulate the TT-5 tensegrity structure under different impact conditions. The necessary constants and variables used in the simulation include: structural properties (lengths, cross-sectional areas, Young's modulus, densities, and pre-stress), motion conditions (initial orientation and velocity), and gravitational acceleration. The information of the test model is shown in the table below, and the FEM model is shown in Figure 51.

Component	Length (m)	Young's Modulus (GPa)	Density (kg/m^3)	Outer Diameter (m)	Internal Diameter (m)	Pre-stress (N)
Rods	1	4.00E+11	1.75E+03	0.01905	0.01656	-
Cables	0.6124	2.03E+09	3.00E+02	0.0003	-	38

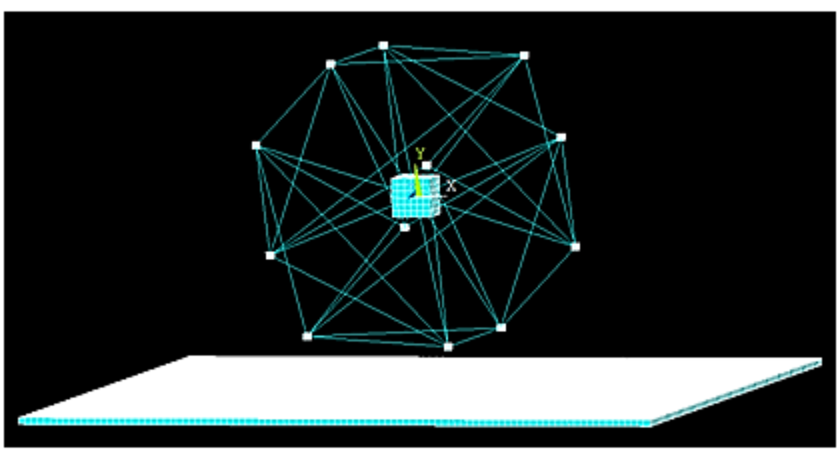

Figure 51: FEM Model.

Assuming an impact velocity of 15 m/s for a payload weight of 500g, the simulation results of the *LS-dyna* module in *ANSYS* includes: impact form, inner force of each element, velocity, and acceleration of each element. The whole motion of the 6-bar tensegrity during the impact is shown in Figure 52 below and the acceleration of the payload over time is shown in Figure 53.

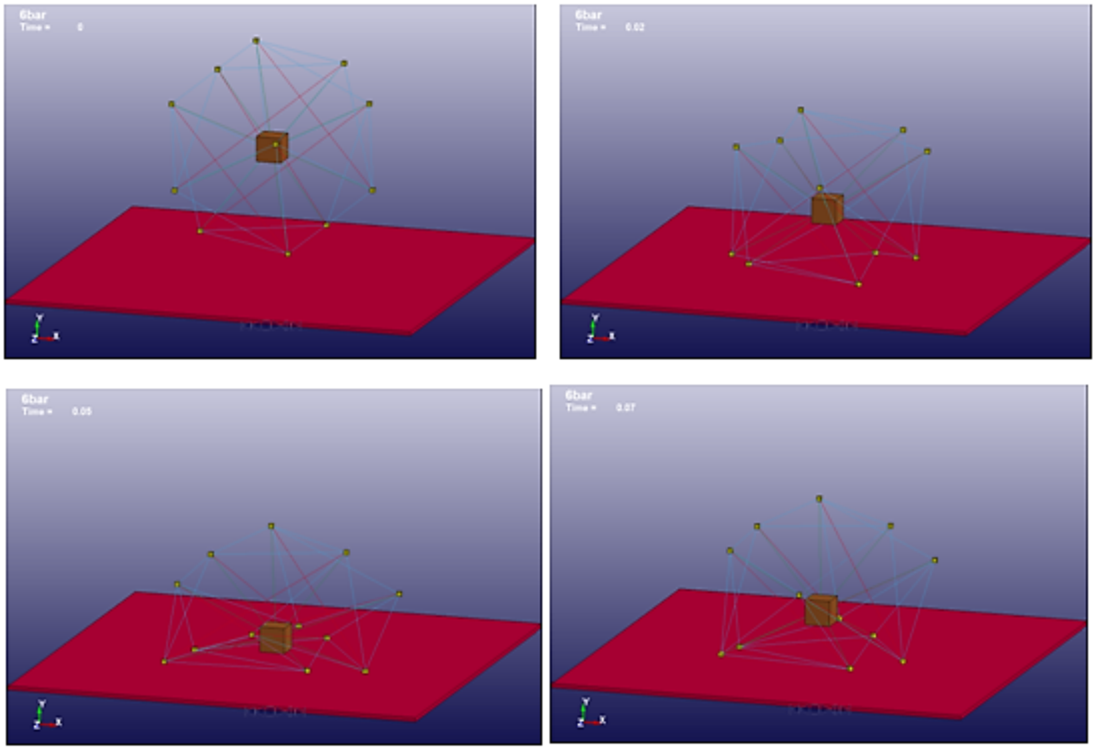

Figure 52: Impact Motion.

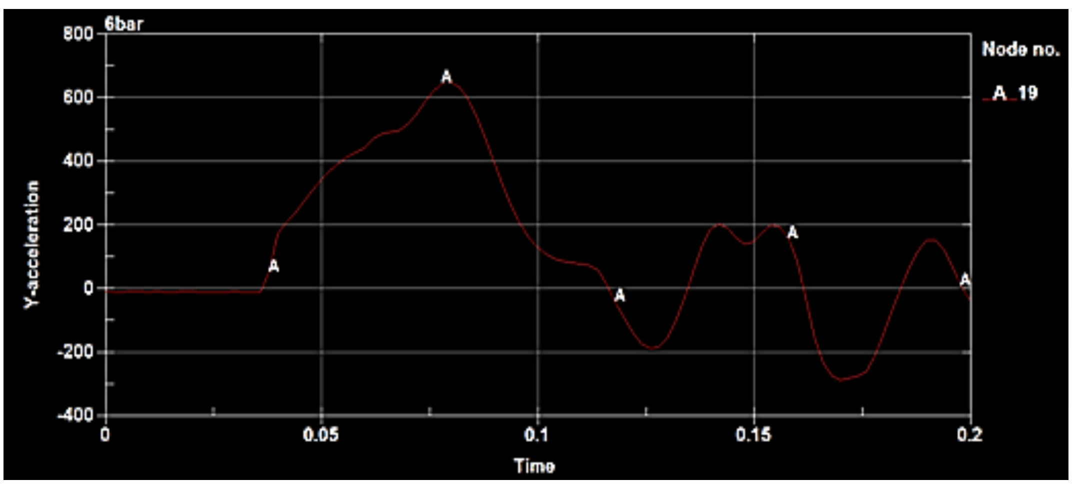

Figure 53: Payload Acceleration.

The maximum acceleration of the payload is nearly 70G. This is lower than our experimental results, indicating a mismatch between our model assumptions and our measured data. In future, we will need to adjust some parameters, such the coefficients of friction, to make the simulation more reliable and accurate. A future goal will be the optimization of the simulation model so that the topology and pre-stress form of the 6-bar tensegrity can be rapidly designed to meet different impact scenario requirements.

3.3.3.2. *Payload connection characterization under various thrust profiles*

Testing to characterize the effect of different thrust profiles on payload connections was performed. A testbed, which consists of a gimbaled thruster prototype centrally suspended to a static 6-bar structure, was built to observe the effects of thrust (see Figure 54). The objective was to characterize how different payload attachment schemes and spring stiffness reacted with respect to different thrust magnitudes and gimbal motion profiles. The dynamic reaction of the gimbaled thruster system is observed through potentiometers measuring tension via linear deflection of the payload attachment springs, acceleration and orientation data from an IMU mounted on the payload, as well as footage from a high-speed camera. Preliminary tests were conducted to verify the functionality of the testbed.

With the construction and initial testing of the gimbaled-thruster testbed completed in Q9, we moved forward with performing the planned experiments which would improve understanding on how the centrally suspended payload interacted with the outer tensegrity shell under different thrust profiles. Three experiments were performed with a vertical thrust vector, an angled thrust vector, and a constant frequency oscillating thrust vector respectively. In all three experiments a 30 N thruster was used. The results show that the thrust force was being transmitted effectively, through tension in the connecting cables, to the outer shell. In addition, given a sufficiently rigid outer shell, it appears that the system was able to achieve equilibrium states that matched the thrust profile of the thruster used. While the work done so far is not yet sufficient to characterize the behavior of a centrally-suspended thruster, it serves as validation of the proof of concept of attaching a centrally-suspended cold-gas thruster.

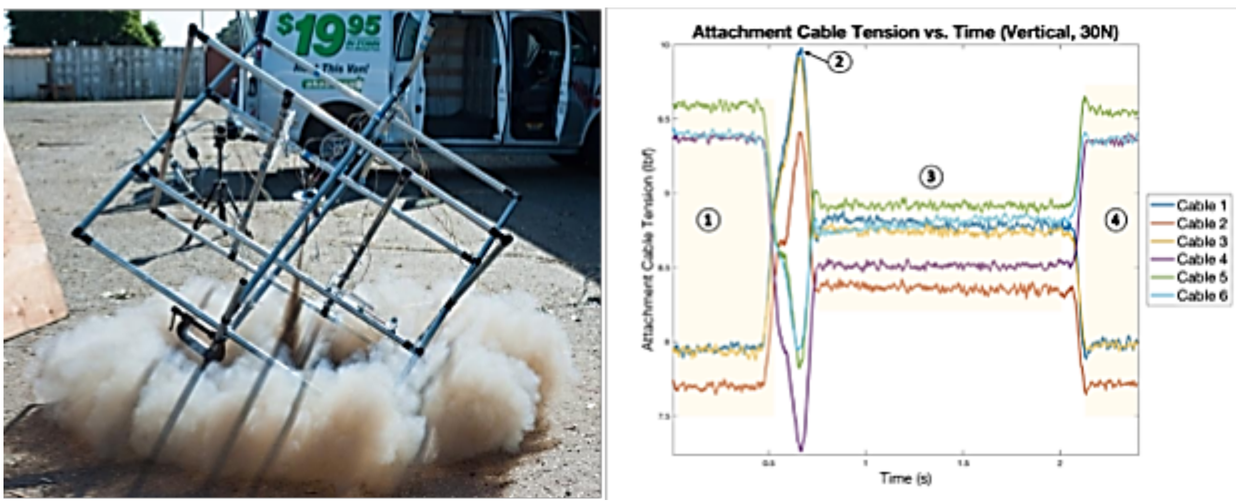

Figure 54: (Left) Experiment in progress. (Right) Data from vertical thrust experiment illustrating the four phases of a thrust event. 1: Original equilibrium. 2: Peak thrust. 3: Constant thrust. 4: Return to equilibrium

3.4. Additional Research: Tensegrity Robots Based on 12-Bar Structures

So far, all our work has been based on a six-bar tensegrity structure. However, we have identified several advantages for tensegrity structures with a higher number of bars and thought it would be useful to explore these as well.

We explored the advantages in terms of locomotion, payload, and impact capability of tensegrity structures with more than six rods. We selected tensegrity structures with 12 rods, which form the next-largest symmetric structures, to explore the advantages that these new topologies may offer.

Unlike a six-bar tensegrity structure, multiple configurations are available for 12-bar tensegrity structures. The Class I configurations that we have investigated are shown in Figure 55. These structures are named *cube*, *octahedron*, *double-six*, and *rhombicuboctahedron*. Cube and octahedron are so named for the shapes from which the rods of the structures evolve. The double-six is so named because it is like the six-bar structure but with another bar added in parallel to each of the six bars. The rhombicuboctahedron is so named for the shape of its exterior lattice.

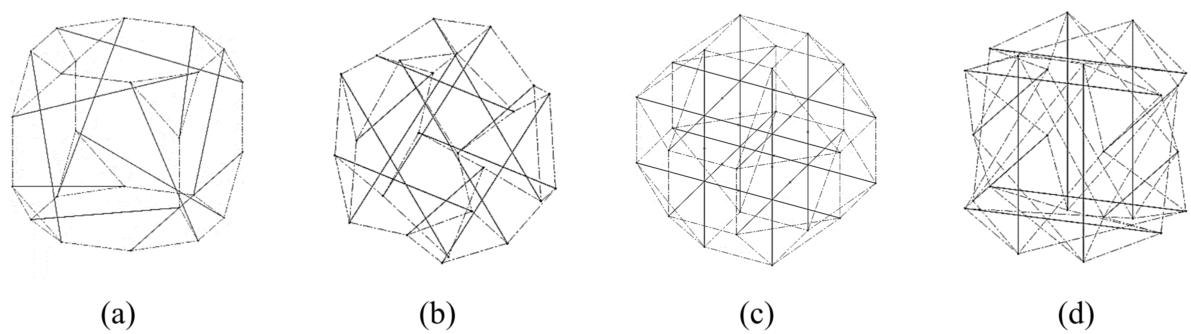

Figure 55. Dimetric views of four Class I configurations of 12-bar tensegrity structures: (a) cube, (b) octahedron, (c) double-six, and (d) rhombicuboctahedron.

The rhombicuboctahedron is symmetric about three orthogonal planes and its outer surface consists of 18 rectangles/squares and eight triangles. Thanks to this geometric features, a robot of this structure can roll in a straight line, as shown in Figure 56. Thus, the momentum loss arising from change of moving direction does not occur, and this robot has a potential in achieving high-speed dynamic rolling.

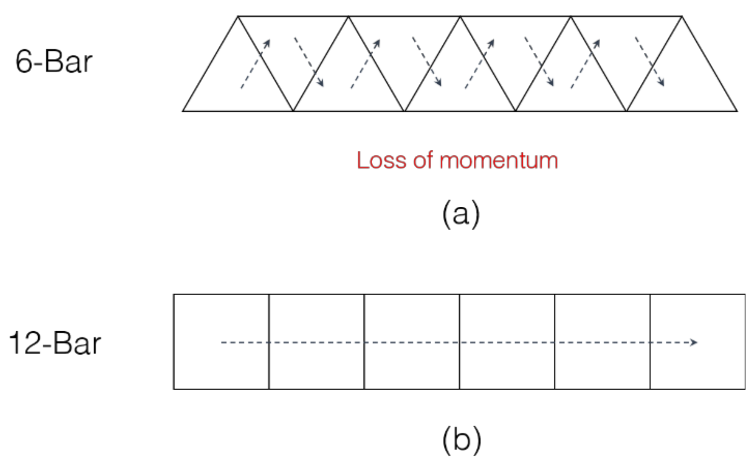

Figure 56. The paths taken by 6-bar and rhombicuboctahedron 12-bar tensegrity robots when moving forward. (a) Because the outer surface of a 6-bar robot consists only of triangles, it moves in a zigzag way, resulting in the loss of momentum. (b) The outer surface of a rhombicuboctahedron 12-bar robot mainly consists of rectangles/squares, so it takes a straight path when moving forward. This prevents the momentum loss arising from change of moving direction.

Supported by the ESI grant, we built a prototype robot, as shown in Figure 57, that we can use for testing in hardware. The overall diameter of the robot is approximately the same as its rod length of 45 cm. The total weight of the robot is 1.78 kg. Located at the center of each rod is a bundle of electronics including distributed controller, wireless radio, battery, etc. There are total of 16 motors installed close to the ends of selected rods and they control cable lengths by spooling. This number of motors is chosen for the initial development of the robot because it is sufficient to realize rolling in a straight line of one chosen direction.

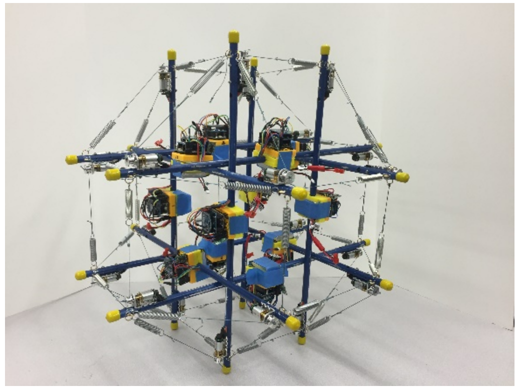

Figure 57. A rapidly prototyped tensegrity robot based on a 12-bar tensegrity structure. The robot has 48 cables. Unlike 6-bar tensegrity robots, this robot consists of 18 rectangles/squares and eight triangles, enabling dynamic rolling. It has distributed controllers on the rods and they communicate wirelessly.

Indeed, our simulations show that the robot should be able to roll with the current number of actuators. The expected deformations are shown in Figure 58. Another advantage of having a rhombicuboctahedron-like shape is that dynamics of the robot can be fully described in a two-dimensional space instead of three, if the deformation is guaranteed to be symmetric about a vertical mid-plane that divides the structure into two symmetric parts (Figure 59). This is a very useful feature to have because the state space dimension can be greatly reduced. A slender rod has 5-DOF in 3D space, and thus a 12-bar tensegrity robot has total of 60-DOF, which makes any controller design on this robot highly challenging. On the other hand, if we assume the symmetric deformation, eight rods that are parallel to the symmetry plane have only 3-DOF per rod and the other four rods that are perpendicular to the plane have only 2-DOF per rod. Furthermore, because the two rods that are mirrored about the plane always have the same motion, the total degrees of freedom for the eight parallel rods are halved to 12-DOF. As a result, the 12-bar tensegrity robot exploiting the symmetry has a total of 20-DOF. This is only one third of the original system's 60-DOF and is even smaller than the total degrees of freedom of a 6-rod tensegrity structure that has 30-DOF. By exploiting this reduction in DOF, dynamics of the 12-rod robot can be compactly written, which facilitates the robot's locomotion controller design.

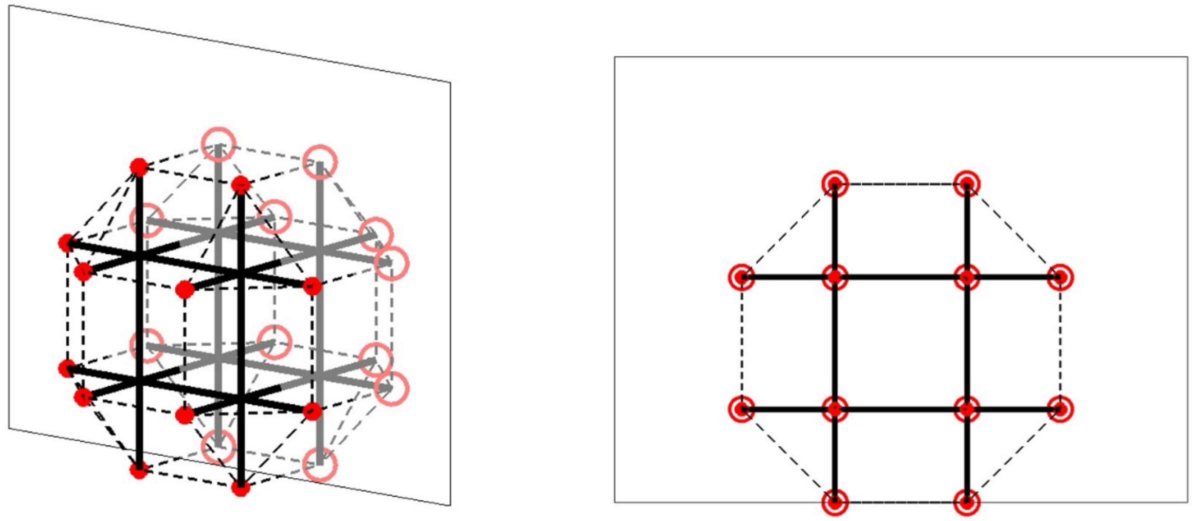

Figure 58. If the deformation of a 12-bar tensegrity robot is symmetric about the mid-plane, then the members of the robot can be projected onto this plane and dynamics of the robot can be fully described on this plane. Notice that for the projected system (right figure), there are four rods that have 3-DOF each and another four rods (shown as points) that have 2-DOF each. In total, the projected system has 20-DOF.

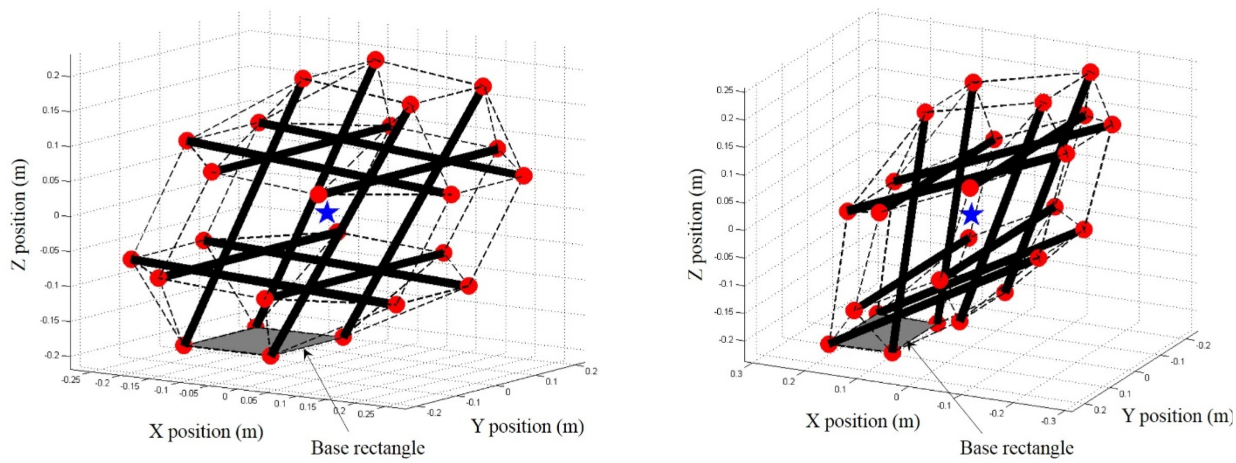

Figure 59. 12-bar tensegrity robot is expected to perform punctuated rolling if deformed as shown in these simulated figures. The ground projection of the center of mass (denoted as a blue star) is outside of the base rectangle, and unbalancing torque is generated by the gravity.

Other structural prototypes of the cube, octahedron, and double-six were made from wooden dowels and rubber bands, as shown in Figure 61. This work was supported by the ESI grant in summer 2016. It was revealed through these prototypes that the intersecting cables of the double-six complicated assembly and hindered shape-shifting when tested by hand. In a full robot, the intersecting cables could cause fraying and snagging. We decided to narrow our focus to the cube and octahedron, which offered promising properties without these limitations.

We made another, more robust set of structural prototypes of these two forms, as shown in Figure 60. These prototypes were constructed using sheets of silicon rubber that were laser cut to create the elastic members. We connected these pieces to create lattice shells. We then attached hollow aluminum rods to the interior of each lattice shell to erect the tensegrity structures. These structures are quick to assemble, easy to adjust, and mechanically robust.

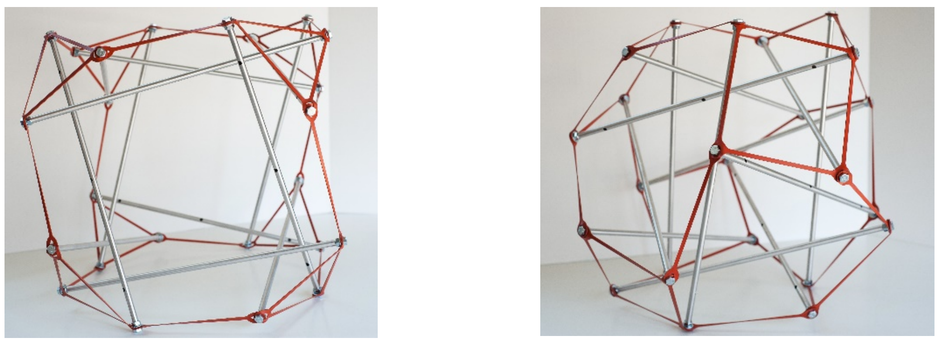

(a) (b)

Figure 60. Structural **prototypes of the (a) cube and (b) octahedron.**

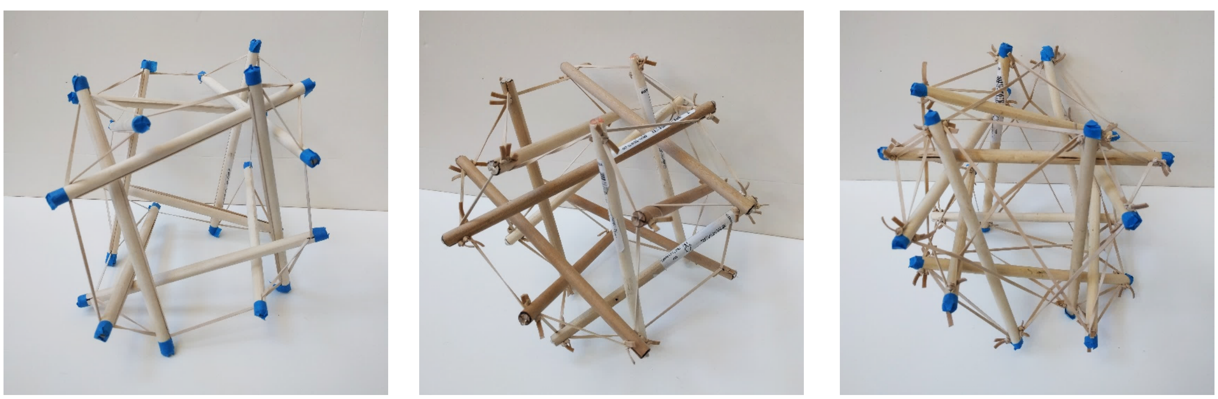

(a) (b) (c)

Figure 61. Initial structural prototypes of the (a) cube, (b) octahedron, and (c) double-six.

Drop tests were conducted for two forms of 12-bar tensegrity structures, the cube and octahedron. Drop test prototypes were made and payloads mounted in the center of each structure. An IMU was attached to each payload, and high-speed cameras were used to record the drop tests. The parameters of greatest interest to an evaluation of these two structures were maximum magnitude of acceleration of the payload and payload safety (i.e., whether the payload impacted the ground). Drop tests were conducted at heights of 3, 4, and 5 feet, measured from the top of each structure (each structure is approximately 1.5 feet in diameter). The payload mass, rod mass, and drop orientation were varied, as summarized in the table below. For each configuration, the test was repeated five times. The maximum accelerations observed ranged from 3.25-7.79g for the cube and from 3.54-7.85g for the octahedron.

Variable	Possible Values
Drop height	3 ft., 4 ft., 5 ft.
Landing orientation	Compliant face, stiff face, intermediate position
Payload mass	342 g, 549 g
Structural mass (not including payload)	542 g, 1165 g

3.5. Summary

In summary, we completed or went beyond all proposed tasks with the ESI grant. The research included simulating and testing hardware concepts for conditions needed for hopping profiles on both smooth and hilly terrain. We expanded the visualization capabilities of the NTRT (NASA Tensegrity Robotics Toolkit). We completed the development and testing of five 6-bar prototypes in hardware, as well as three 12-bar concepts. The TT-4_{mini} became the first untethered spherical tensegrity robot to be able to successful "walk" on an uphill slope.

We developed and analyzed different control strategies to allow the robot to achieve dynamic rolling, resulting in more robust and faster locomotion. We evaluated impact, payload interaction, controls capabilities and payload protection characteristics of the robot with a horizontal launcher, a gimbaled thruster testbed, slope testbed and with vertical drop testing. We made significant progress in achieving TRL 3 to show proof-of-concept with mission profiles tested in simulation with hardware validation.

Finally, in our no-cost extension year, we worked with Squishy Robotics, Inc. – a commercial spin-off of the ESI research – to use their hardware for further refinement and testing resulting in the ability withstand falls of up to 400 ft (122 m), the highest drop legally allowable from a drone. This research is on track to reach levels of TRL 4 or 5 through the partnership with Squishy Robotics, Inc.

The ESI grant was highly leveraged by student teams at UC Berkeley. Two Ph.D. dissertations were completed, and three more are in progress. Twenty-six Masters degrees were completed. Twenty-six undergraduate researchers were engaged, along with two visiting students and two high school students. Thirty-four peer-reviewed publications, theses and presentations were completed with this research. Demonstrations were provided to NASA administrators, researchers as well as the general public in K-12 outreach opportunities. The research was featured on the Discovery Channel, KQED, news articles and various radio programs. Three patents were filed. More details are provided in the impact spreadsheet.

4. NASA COLLABORATION

During most of the years of the ESI grant, the UC Berkeley team was divided into two teams who met separately weekly: (1) controls and simulation and (2) mechatronics design. We held joint weekly meetings for both teams with PI Alice Agogino and frequently with Co-PI Adrian Agogino. Members of the UC Berkeley team work closely with Vytas SunSpiral and Adrian Agogino at NASA Ames.

One Ph.D. student Kyunam Kim and one MS student Kyle Zampaglione worked as interns at NASA Ames during the summer of 2015 and both gave presentations. During the Fall 2016 semester, a new team was formed to explore the potential of 12-bar tensegrity structures with graduate student Mallory Daly, who was awarded an NSTRF grant in Fall 2016.

We coordinated with Terry Fong at high-level events, such as the demos to Dr. Dava Newman[6]. On June 8, 2017, the ESI team was visited by the Associate Administrator Steve Jurczyk of the NASA Space Technology Mission Directorate, where we gave a summary of our research and provided demos of all of our tensegrity robots[7].

After Vytas SunSpiral left for a position with industry, we worked with Terry Fong, Massimo Vespignani and Jonathan Bruce at NASA Ames.

5. ANNUAL TECHNICAL SEMINARS

Award Year	Date	Seminar POC	Seminar Topic and Comments
1	8/5 /2015, 9/1/2015, 9/17/2015	Adrian Agogino, Vytas Sunspiral, NASA Ames	A Structured Linear Actuator for Tensegrity Robots (Kyle Zampaglione) and Precision Hopping and Rolling Robotic Surface Probe Based on Tensegrity Structures (Kyunam Kim)
2	7/14/2016	Vytas Sunspiral, NASA Ames	Precision Hopping/Rolling Robotic Surface Probe Based on Tensegrity Structures (Alice M. Agogino, Adrian K. Agogino with graduate students.)
3	7/25/2017	Massimo Vespignani, JPL	Precision Hopping/Rolling Robotic Surface Probe Based on Tensegrity Structures (Alice M. Agogino with graduate students Drew Sabelhaus, Brian Cera & Edward Zhu.) Also presented two posters during poster session.
4	7/17/2018	Terry Fong, NASA Ames	Multi-Cable Rolling Locomotion with Spherical Tensegrities using Model Predictive Control and Machine Learning (Brian Cera and Alice Agogino)

[6] BEST Lab News Blog: BEST Lab Demos Tensegrity Robots to NASA Deputy Administrator, July 17, 2015, http://best.berkeley.edu/2015/07/17/hello-world-2/.

[7] BEST Lab News Blog: https://best.berkeley.edu/2017/06/08/demo-for-associate-administrator-steve-jurczyk-nasa-space-technology-mission-directorate/

6. POSTDOCTORAL RESEARCHER(S) / STUDENT(S)

The table below summarized the doctoral, masters and undergraduate students who have worked on the ESI grant.

Assistance Type	Number	Roles / Comments
Doctoral Students (Year 1 & 2)	7	Kyunam Kim[8], Lee-Huang Chen[7], Mallory Daly[7,9], Angelo Brian Cera[7], Edward Liu Zhu, Alan Zhang and Jianlan Luo: Worked on both simulation and hardware for ESI.
Doctoral Students (Years 3 & 4)	2	Angelo Brian Cera Daly[7] and Jianlan Luo[7]: Worked on both simulation and hardware for ESI.
MS & MEng (Years 1 & 2)	14	Kyle Zampaglione[7] DNA actuator. Alex Lim, Azhar Khaderi, Deegan Peadar, Xiang Li: Dynamics analysis and mechatronic design. Jeff Ware, Julien Despois, Vincent Viola, Marcel Pozo, Anupama Madiyan, Yang Zheng, Yinglong Li, Borui Xia: Dynamics analysis and mechatronic design.
MS & MEng (Years 3 & 4)	22	Carrina Dong, Henry He Huang, Eric Jiang, Anosh Sethna, Michael Wu, Nikki Chen, Marshall Hoaglan, Zhong Jin, Juan Ordonez, Zareen Cheema, Stuart Sonatina, Jeff Ying, Jovin Fu, Mrunal Sarvaiya, Tianyi Chen, Joshua Peterson, Mason Friedberg, Prasad Hemant Gaikwad, Yee Lin,. Jovin Foo, Charlotte Chapellier, Zining Wang. MEng students working on dynamic systems and testing.
Undergraduates (Years 1& 2)	20	Raymond Ennis, Kimberley Fountain, Kevin Li, Wesley Wang, Ellande Tang, Jeremy Wan, Richard House, Ellande Tang, Ankita Joshi, Saunon Malekshahi, Lua Varner, Hunter Garnier, Faraz Ghahani, Grant Emmendorfer, Mari Verdugo, Kevin Li, Abhishyant Khare, Cameron Bauer, Wesley Wang, Jae Young Bin: Supported mechanical design and testing.
Undergraduates (Years 3 & 4)	6	Antonia Bronars, Andrew Plewe, Grant Emmendorfer, Ryan Cosner, Aliya Kusumo, Anthony Thompson: Provided support as undergraduate researchers on dynamic systems and testing.
Visiting Students	2	Osvaldo Romero, Yuen Wun Chau: Worked on K-12 kit.
K-12 Students	2	Eirren Viray, Sebastian Anwar: Worked on K-12 kit.

[8] Has received partial funding by ESI grant.

[9] Has also received an NSTRF fellowship.