NASA/TM–20210013185

Feasibility Analysis of Unmanned Aerial Vehicle Based on Tensegrity Structure

Bjorn Johnson
University of California, San Diego

Dr. Adrian Agogino
NASA Ames Research Center

March 2021

NASA STI Program... in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI Program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collection of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI Program, see the following:

- Access the NASA STI program home page at `http://www.sti.nasa.gov`

- E-mail your question to `help@sti.nasa.gov`

- Phone the NASA STI Information Desk at 757-864-9658

- Write to:
  NASA STI Information Desk
  Mail Stop 148
  NASA Langley Research Center
  Hampton, VA 23681-2199

NASA/TM–20210013185

Feasibility Analysis of Unmanned Aerial Vehicle Based on Tensegrity Structure

Bjorn Johnson
University of California, San Diego

Dr. Adrian Agogino
NASA Ames Research Center

National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035

March 2021

Available from:

NASA STI Program / Mail Stop 148
NASA Langley Research Center
Hampton, VA 23681-2199
Fax: 757-864-6500

Abstract

For future unmanned aerial vehicles to be ubiquitous, they need to be safe to other aircraft and people on the ground even in the event of unintended collisions. Unfortunately, this is difficult for current rigid designs as a crash of all but the smallest aircraft will have a significant impact on its target. As an alternative, this work explores using a soft tensegrity structure (interconnected cables and rods) as a lightweight, robust chassis for a UAS and analyzes its ability to achieve stable, controllable flight and mitigate damage during a collision. These structures are able to undergo severe deformations without permanent damage. Experiments and analysis are performed using a simple physics simulation environment. We test multiple configurations of the vehicle to understand impacts of changing various physical characteristics, and take limitations such as modern propeller thrust coefficients and motor speeds into account to ensure our results are in line with those of a real world vehicle. We find that such a vehicle has considerable real world potential, and that a physical prototype could theoretically be constructed and flown with current technology.

1 Introduction

Recent research has shown that tensegrity systems hold great potential in a variety of diverse applications including terrestrial locomotion [1–4] and wearable exosuits [5]. Their inherent properties such as low mass, elasticity, and robustness are incredibly promising for small scale aerial vehicles, inspiring us to further investigate the feasibility of a vehicle based on such a structure. A vehicle which effectively takes advantage of these properties would be valuable for payload protection, ease of repair, and compact storage configuration. In addition, with rapidly increasing numbers of small vehicles in our airspace, taking potential crashes into consideration becomes important. An increasing amount of research has been performed to mitigate this issue [6, 7]. A largely compliant aerial vehicle would be both more robust and less dangerous to other entities in the event of a collision. This paper explores a number of variations of such a vehicle, analyzing each for stability and added benefits over traditional multirotor vehicles of similar proportions. Multiple physically distinct configurations are simulated and observed, and our analytical criteria includes flight stability, controllability, responsiveness and degree of deformation in collision. Each variation of the vehicle is flown through a number of simple trajectories and intentionally crashed. Its performance is then assessed on the above criteria and compared to similar configurations to develop an understanding of both the general behavior of such a system and the impact of various changes in its design.

2 Background

Tensegrity (tensional integrity) structures consist of components under compression, balanced by a network of components under tension. They are typically constructed out of a system of stiff rods with no rigid joints, connected by springs and cables to hold the system in a desired geometric formation. Historically, tensegrity structures have been used in artistic sculptures, toys, and brain teasers because of their visually appealing and perplexing appearances, but have not been widely explored for vehicle design or actuated systems. Recent efforts, however, have shown that tensegrity structures can be successfully used in a number of considerably distinct tasks, such as in a wearable exoskeleton where tensional elements aid muscles to actuate limbs for rehabilitation purposes. More popularly, they have been used to build roughly spherical locomotive systems where actuation shortens some dimensions relative to others, allowing the entire system to roll indefinitely along the ground without the need for wheels or legs. In addition to ground vehicles research has been performed on multirotors with protective tensegrity shells [7]. For instance in [7], they develop a rigid quad rotor that is rigidly attached to two struts of a 6-rod tensegrity structure. Also a team at Innopolis University has developed a quad rotor with a tensegrity shell (named “Tensodrone”) where each pair of rotors is attached independently to two rods.

Based on the success of such systems, we further examine the application of an aerial multirotor vehicle, with a goal of damage mitigation in the event of a collision. We expected to see success in this effort due to the customizability inherent in tensegrity systems, a highly favorable size to weight ratio, and most importantly a large degree of elasticity. We aimed to assess the potential benefits of such a system in comparison to the drawbacks of moving away from a rigid frame, which we expected to include large internal vibrations and difficulty controlling the vehicle. Through our simulation efforts, we found that these systems are actually free of significant detrimental internal vibrations, robust to changes in their construction, capable of absorbing impacts effectively, and controllable with simple classical methods, albeit with a few caveats in state estimation.

3 Methods

We sought to explore the feasibility of a quadrotor vehicle based on a tensegrity structure, aiming to determine whether such a system is possible to build with current technology and whether it provides enough potential benefits to warrant further exploration. Finding that it was both possible and potentially beneficial, we next sought to analyze how robust the vehicle would be to changes in its general configuration, and provide estimates for the optimal properties and dimensions of a prototype. We built a vehicle simulation using a simple physics framework and tested a number of system configurations, focusing our attention mainly on stability of the vehicle in flight, deformations caused by aggressive maneuvers, and degree of frame compression upon collision to determine success. Large deformations during flight were found to be highly undesirable, as the orientation of the propellers with respect to each other can change significantly, altering the dynamics of the system considerably. Large deformations during collisions, however, proved much more desirable, as they prolong and soften the impact, resulting in smaller forces and thus less damage to both the vehicle and the object it collides with. All of our experimentation was performed on a six-rod tensegrity ball structure, composed of three sets of parallel rods, where each set is perpendicular to the other sets. Rods are

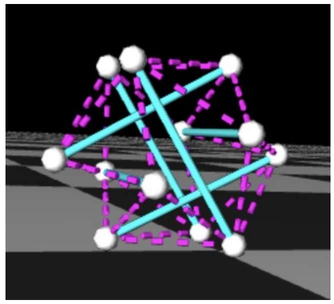

Figure 1: Six rod tensegrity structure

connected to each other using cables between their endpoints, in a predetermined pattern. This particular structure was chosen for its symmetrical construction, known robustness, and convenience of mounting propellers. Fig 1 shows the basic structure.

3.1 Exploration of Vehicle Configurations

In order to effectively explore the feasible space of vehicle configurations and properly analyze performance of the system in general, a small number of tunable parameters were chosen as the basis of a number of experimental trials. These parameters were selected based on their potential to impact the dynamics of the system, altering its observable behavior in flight and collision. For each of these parameters, a representative sampling of their possible values was selected for testing, and an exhaustive search was performed by testing all possible combinations of these values, resulting in a large number of distinct vehicles. In particular, we experimented by varying:

- Cable stiffness
- Vehicle size and mass
- Propeller thrust to torque ratio

Altering the stiffness of the cables has a massive and relatively intuitive impact on the way the structure handles the forces present in flight and collision. Stiffer cables make for a less rigid frame, which flexes less in flight and collapses less on collision. We aimed to find a reasonable balance between these two traits, searching for a cable stiffness that both appropriately exploited the novelty of the flexible frame for lessening impact forces upon collision and minimized deformations during flight. We used cables acting as simple springs, and chose 4 uniformly distributed values for spring constants to cover the possible range of usable cables. Values were selected such that the structure ranged from being compliant enough with the most flexible to fold almost completely flat when at rest on the ground, to being just rigid enough with the stiffest that there was near negligible deformation of the frame in any scenario, across both flight and collision.

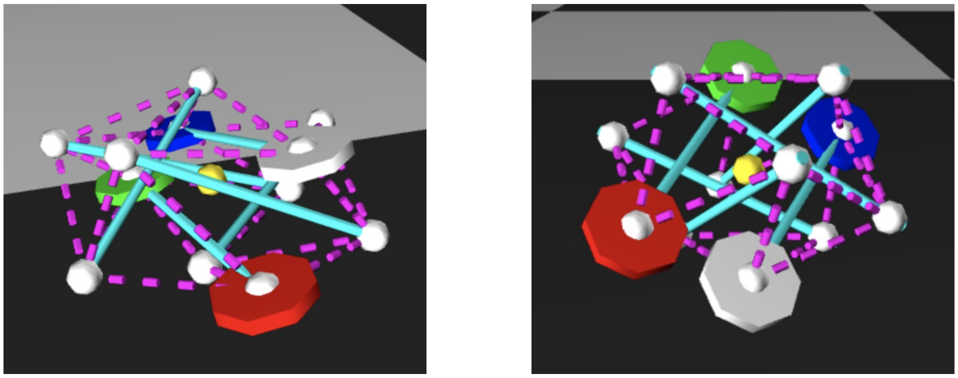

(a) Extreme (most flexible cables) (b) Moderate (stiffer cables)

Figure 2: Examples of resting deformation at different cable stiffnesses

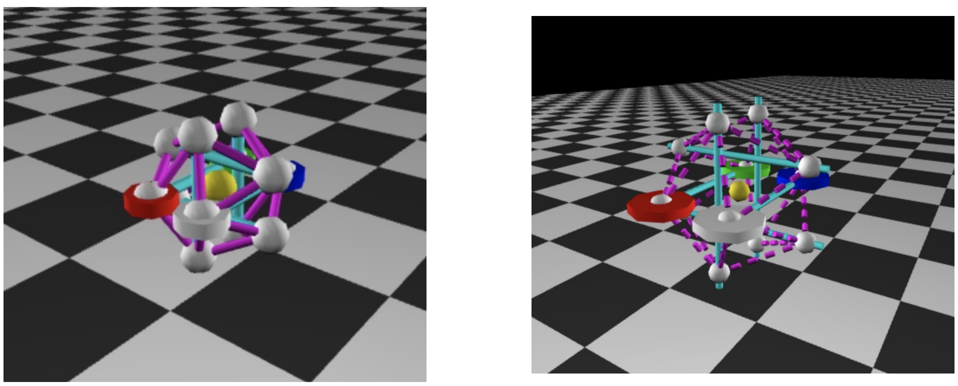

(a) Small configuration 0.5m rods (b) Large configuration 1m rods

Figure 3: The two size configurations tested

For overall size and mass, we chose to experiment with two discrete sizes, assuming that behavior did not change dramatically at intermediate sizes with otherwise identical characteristics. In initial experiments, very small configurations needed very stiff cables to prevent rotors colliding with each other and the frame, to the point where they behaved quite similarly to a similarly sized rigid quadrotor, defeating the purpose of using the more complicated chassis. With this in mind we selected two relatively large overall sizes in comparison to common quadrotors, with rod length and thus roughly diameters of 0.5m for the small and 1m for the large configuration. Based on estimates of the masses of typical tensegrity frames and common masses of other necessary quadrotor components such as motors, batteries, and propellers, we assigned these configurations overall masses of 1.5 and 3kg, respectively.

The other important parameter we experimented with was the thrust to torque ratio, a constant intrinsic to the propeller. This quantity was more difficult to identify as important, but was found to have a huge impact on the system and thus warrants a more in depth discussion. We initially selected the thrust and torque coefficients of the propeller as parameters to explore independently, but experimentation quickly showed that varying the coefficient magnitudes didn't always result in intuitive or particularly useful results. While propeller speeds were often changing a lot in response to changes, the system was not always showing a proportionate change in behavior. Upon analyzing the equations governing quadrotor flight dynamics more thoroughly, it was discovered that the actual values of the coefficients for these properties were irrelevant, and that only the ratio between them had any observable impact on the simulated system. The motors can spin arbitrarily faster or slower to make up for decreases in the performance of the propeller, resulting in identical forces being applied to the system when both coefficients were modified by the same factor. In a real world setting the propellers cannot take on any arbitrary speed to generate the exact same forces, but the principle is the same in that motors can spin faster to account for slightly weaker overall propellers without drastically changing system dynamics.

We enumerate the relevant equations for this behavior below and show that the key takeaway is both thrust and torque generated by typical hobbyist propellers can be modelled reasonably accurately by a simple quadratic relationship. Thrust and torque each scale directly with the squared propeller speed, governed independently by coefficients intrinsic to the propeller's physical characteristics. Examining this relationship, we see that changing the values K_f and K_t by the same factor can be compensated for by a corresponding change in ω^2, so the system is capable of producing the exact same output thrust and torque with a number of distinct propellers.

$$
\begin{aligned}
F =& \text{Thrust generated in the direction of flight} \\
T =& \text{Torque generated around center axis} \\
\omega =& \text{Angular velocity of propeller} \\
F =& K_f \omega^2 \\
T =& K_t \omega^2
\end{aligned}
$$

In light of these findings, we keep one constant fixed at a reasonable value and vary the other to change the behavior of the system. This is reasonable because we don't particularly care if our propellers are twice as strong if they produce the same behavior. Performing experiments in this manner, we confirm that changing the ratio between K_f and K_t causes profound differences in system dynamics, directly influencing the amount of frame deformation during flight. These differences are particularly apparent while driving rotation about the vertical axis of the vehicle.

We perform an analysis of driving yaw while maintaining a stable hover to clearly show these impacts on the system and describe a direct relation between vehicle behavior and thrust to torque ratio. This hovering and yawing behavior is quite common in multirotor flight, especially in investigating the behavior of a system, and involves maintaining a fixed net thrust while modifying motor speeds to generate a desired net torque. Assuming a relatively modest desired net torque, which is

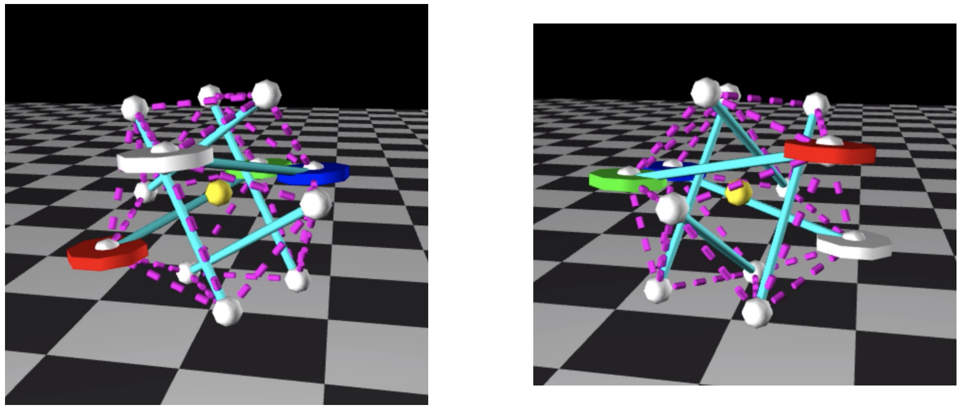

Figure 4: Two examples of extreme deformation in flight

all that is necessary in the hovering scenario, only a relatively modest change in individual propeller torques from the typical hover state is needed. The controller can therefore change the speed of each motor by a modest amount, generating the newly desired individual and net torques. For small values of the thrust to torque ratio, this generates a relatively modest change in individual thrusts, since change in thrust is proportional and similar in magnitude to change in torque at the propeller level. Therefore for a small thrust to torque ratio, upward thrust at each propeller is only changed by a similar modest amount.

Assuming an identical K_t and a much larger K_f, i.e. a much higher thrust to torque ratio, a very similar change in the speed of each motor is required to generate the desired change in torque at each propeller, but with a significantly increased value for K_f, we see that this similar change in rotor speed generates a much larger change in upward thrust at each propeller. The practical result of this observation in a quadrotor is that adjacent corners of the vehicle experience significantly different upward thrusts, resulting in a bending moment on the frame. This is true in all quadrotors, but in a rigid vehicle the frame can absorb the torsional force with very small corresponding deformation, effectively rendering it a non-issue. Due to the highly compliant inherent nature of a tensegrity system, however, we see significant amounts of deformation in the frame under similar bending moments, which as explained previously can be highly undesirable. This ratio is thus incredibly important to the success of our vehicle both in simulated experiments and a potential physical prototype.

Since this ratio was found to be crucially important to the behavior of our vehicle, we researched a large number of hobby propellers and their associated coefficients to assemble a range of realistic ratios for use in our experiments. We pulled dynamometer data from [8] and performed a simple analysis to find the coefficients that best represent thrust and torque as a factor of squared angular velocity for a number of common propellers. We also extensively researched the issue in current literature, and found that the usual type of propellers in multirotors have a relatively small range of possible values for this ratio compared to the range of magnitudes in the constants themselves. Based on our findings, common hobbyist propellers tend to have an coefficient ratio K_f/K_t in the range of roughly 30 to 65. For the sake of both coverage of the typical range and consideration of extreme cases, we tested our simulation using propellers with thrust to torque ratios of 25, 50, and 75.

In total, we enumerated 3 different highly impactful properties of the vehicle to experiment with in simulation, and a small number of carefully chosen representative values for each. This made for a total of 24 distinct vehicle configurations, where we tested each by manual flight through a representative trajectory, exploring the limits of the vehicle under aggressive inputs. Each distinct configuration was analyzed in flight as well as collision into both the ground plane and a vertical wall, and was evaluated using a number of simple criteria to allow comparison across configurations. We looked mainly at the degree of deformation experienced by the system during both flight and collision, as well as overall stability and responsiveness of the vehicle to input commands.

3.2 Simulation Details

We performed our simulation using the Cannon.js framework [9], which implements simple physics integration and constraints for a number of solid bodies. The biggest challenge in building the simulation was designing a controller to fly the vehicle, as a standard quadrotor controller relies on knowing the position and orientation of the system, and given the flexible and distributed nature of our vehicle this posed a problem. In particular, we found it necessary estimate the roll, pitch, yaw, and corresponding angular velocities of the vehicle to within a reasonable degree of the ground truth in order to achieve stable flight. With proper estimates, a simple PD controller with the same assumptions as those of a rigid quadcopter was shown to perform well in teleoperated flight, but calculating these estimates was nontrivial. Standard quadcopters use a central inertial measurement unit, and in the simulated case this can be emulated by observing the exact orientation of a body mounted to the center of the frame, optionally adding noise to increase robustness. With no central component of the tensegrity frame to observe, we were forced to rely on the connection nodes to calculate the desired properties, as these nodes had observable positions.

We found that using the four nodes with propellers was sufficient, and once a set of body fixed coordinate axes was found some simple trigonometry was enough to back out the necessary measurements. Generating the necessary coordinate frame requires building and utilizing a number of support vectors, the geometric process is visualized in Figure 5. It begins with calculating a midpoint between the front two rotors as well as the rear. We then defined the x axis of the vehicle as the unit vector pointing from the rear midpoint to the front. Next we calculated the vector directions from the rear midpoint to each of the two front rotors, and took the normalized cross product between them to generate the vehicle's z axis. A representation for the y axis can then be found by taking the cross product of the unit x and z vectors, completing the coordinate frame. Armed with a reasonable estimate for the body frame, it becomes possible calculate the roll, pitch, and yaw of the vehicle by finding the rotations about each of the x, y, and z axes which would transform a world fixed frame into our body frame. As it happens this calculation

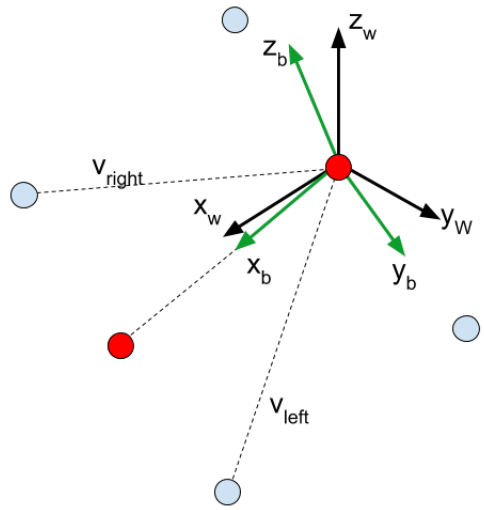

Figure 5: Visualization of the orientation calculation process

is relatively easy using inverse tangent and sine functions, but is not particularly stable and breaks down in many common configurations of a multirotor.

If any of the rotations are particularly large, the quality of the other estimates deteriorates, since the body frame axes leaves the quadrants of 3D space where our simple calculation is valid. This issue is present when any of the angles are particularly large, but in a realistic flight setting is mostly a problem for large yaw angles. Large pitch and roll angles render the controller unstable anyway, as the vehicle becomes oriented nearly horizontally and is inherently more difficult to control.

For any arbitrary yaw angle, however, the vehicle can hover equally well from a physics standpoint and thus yaw should have no impact on our controller. As long as neither the pitch or roll angle exceed 90 degrees, which is the case for nearly all reasonable configurations, calculating the yaw angle about a world frame based on the body fixed x axis is robust, so we use this observation to improve our calculations of roll and pitch. This can be done by using the calculated yaw at the current timestep to find a coordinate frame which represents exactly the world fixed frame, only rotated about the z axis by our current yaw. This makes the calculation significantly more robust, as the yaw angle between our new world frame and body frame is zero, allowing the pitch and roll to be calculated without issue.

We find this frame by projecting the x axis of our vehicle onto the horizontal 2D plane and normalizing to unit length, calling the result $x_{adjusted}$, which represents the x axis of the desired coordinate frame. We now take the cross product $x_{adjusted}$ with the world fixed unit z vector, resulting in the axis we'll refer to as $y_{adjusted}$. This gives us a full coordinate frame defined by $x_{adjusted}$, $y_{adjusted}$, and the world fixed unit z vector, where the axes of our body frame fall into much better behaved quadrants. The same simple trigonometry as before can now be used to find the pitch and roll angles of the vehicle with respect to the adjusted world frame, producing equally viable estimates of pitch and roll which are longer disturbed by rotations of the vehicle about the vertical axis.

Once we have an estimation of the vehicle's orientation at our current timestep, calculating the angular velocities about each of the body axes becomes near trivial, we simply take the difference between the previous timestep's angles and the current timestep's, and scale by the unit of time separating our timesteps, which is a property of the simulation. Control of this system can be performed identically to the process used for a rigid quadrotor, by calculating the error between the current angles and velocities and those set by our trajectory planner (or keyboard inputs in our experiments) and applying PD control to find control inputs in the form of desired torque about each of the vehicle's body axes, as well as net thrust of the propellers. This is a very common choice of control inputs in current literature, as it is near globally stable and can be used to effectively drive the quadrotor to almost any desired state. For a typical fixed frame quadrotor using the control policy described, the relationship between squared motor speeds, body torques, and net thrust is linear and can be represented in a simple constant state transition matrix. We assume the same constant state transition matrix as a fixed frame quadrotor for our system, an assumption that is roughly accurate while frame deformation is small, and was found experimentally to hold for a well designed system in nearly all cases except for a crash.

In the event of a crash, collision effects dominate and inaccuracies in control input are no longer particularly important, so we can discard them anyway. This state transition matrix is conveniently invertible, and thus finding the necessary motor inputs given desired control input is trivial with a simple matrix multiplication. The simulation then uses these motor inputs with the previously mentioned equations calculating thrust and torque as a factor of propeller speed to find net forces on the propellers, and applies these forces to effectively close the control loop. We have successfully created a system for finding the current vehicle state, calculating desired control inputs, converting to necessary motor signals, and applying the new motor states to update the subsequent vehicle state. In later sections we will elaborate on the parallels between this process and that of a theoretical prototype vehicle, diving into their similarities and necessary differences, focusing largely on how this workflow can be modified to function in the real world.

3.3 Transitioning From Simulation to Reality

Our efforts have successfully shown that such a flexible aerial system is highly likely to be physically feasible, but this finding alone is not enough to dive into building a prototype without further investigation. While the simulation environment was relatively robust, we expect a physical system to be significantly less so, and for this reason need to carefully consider the failure modes of our vehicle. We already discussed a number of consequences of changing physical characteristics of the model and ways to alleviate these problems, so we now focus our study on the difference between the simulation environment and a real world prototype vehicle.

In simulation, vehicle weight was roughly evenly distributed between the connecting nodes and the propellers, concentrating the mass near the outer limits of the vehicle. A physical prototype would also have a number of electrical components, particularly dense batteries, nearer to the center of the vehicle, meaning that our simulated vehicles have slightly higher inertia than similarly built real world vehicles. Thus the simulated vehicles likely experience slightly more deformations during aggressive flight, which is actually somewhat convenient, as it shows us that a realistic system would probably experience less of the critical problems we observed. It is likely we would see the platform exhibit considerably more stable flight than our simulation, especially if the real system were lighter overall. This also gives us some margin for error in the opposite direction, if a real vehicle were more massive than our simulation, this extra mass would likely be concentrated toward the center, and thus the same ideas apply. In effect this means we are confident that a physical prototype would experience similar or smaller deformations in flight than our simulations, which based on our experiments was one of the biggest risks of utilizing a tensegrity frame. This is the most profound physical discrepancy between the simulation and reality, and as explained it functions in our favor.

In addition, discrepancies will arise in our control policy, particularly in the estimation of vehicle pose. In typical rigid systems an inertial measurement unit would be mounted to the center of the frame, and could be used relatively easily to estimate the orientation of the vehicle. For a tensegrity based vehicle, mounting the IMU poses a problem, as there is no fixed center point of the frame, and even if the sensor were in the center, it would not be aligned with the rotors in the event of a large deformation. In simulation we avoided addressing this issue by simply calculating the orientation based on the known positions of the rotors, but in the real world this information does not exist. We can theoretically apply the exact same control policy we implemented in simulation, but we need to find a different way to estimate the orientation of the vehicle. We consider two solutions to this problem, and evaluate the potential benefits and issues faced by each.

The first solution directly tackles the problem of lacking a mounting point by introducing a central platform, connected to the outer frame of the vehicle with similar elastic cables as those constructing the frame. This creates a fully suspended central location where we can mount batteries, IMUs, processing units, and any other arbitrary hardware which might be necessary for a multirotor system. Such a solution is very appealing in that it provides a mount point for a significant amount of electrical hardware which might otherwise be difficult to deal with. It is also very appealing from a pose estimation standpoint, as the IMU is in the center of the structure and can thus directly generate estimates of position and orientation from its own measurements without any complex analysis or sensor synthesis. On top of this, it places all of the most expensive components on the vehicle as far from the outer edges as possible, effectively protecting the payload of the vehicle to the best of its ability in the event of a crash, one of the properties most desirable in these systems.

A number of pitfalls must be considered, however: largely the tension introduced by such a system and the noise this would introduce into the IMU measurements. Suspending a large amount of mass such as this in the center of the vehicle would require the supporting cables to be rather stiff, in turn pulling the frame inward more than previously necessary. It's possible this would not be a problem in a well designed system, but it introduces a new failure mode where the propellers draw together and potentially collide with each other or the frame. In addition, mounting the IMU to a platform suspended by elastic cables would inevitably introduce a significant amount of noise into its measurements, as the IMU could shift a small amount in any arbitrary direction without the frame moving correspondingly, and might introduce oscillations of the sensor. This has huge negative impact on the system, with lower quality data the reconstructed orientation estimate is necessarily of lower quality, which could lead to unstable or more erratic controller inputs.

The second solution would be slightly more difficult from a software perspective, but much more physically robust, and would involve mounting batteries, IMUs, and essentially as much necessary hardware as possible to the rods connecting the propellers. This would avoid concentrating mass in the center of the vehicle, likely making it more responsive to control inputs and experiencing a more uniform distribution of forces in a crash. In this case a smaller valuable payload could still be suspended in the center for maximum protection, but as much mass as possible would be distributed to other parts of the frame. This system would require at least two IMU sensors, one per horizontal rod, and synthesize the outputs from these sensors to form a unified estimate of the vehicle's orientation. This would be a more involved process for calculating the necessary estimate, but would likely be significantly more robust, which is incredibly important with a vehicle as ill behaved as a flexible multirotor. With a fusion of two sensors instead of one, their individual noise measurements matter less, and the resulting estimate would have lower variance. In addition, these sensors would be mounted to the same rods as the propellers, ensuring that they could not move, twist, or oscillate unless the vehicle as a whole was experiencing the same disturbances, eliminating a massive source of possible error.

4 Results

We aimed to assess the potential benefits of a UAS based on a tensegrity structure in comparison to the difficulties and drawbacks of moving away from a rigid frame. Experiments were conducted on a number of system configurations, and interesting flight characteristics of each were noted and qualitatively ranked. We focused on the stability of the vehicle, deformations, and compression upon collision to assess the quality and feasibility of each configuration.

Contrary to our expectations, we found that internal vibrations were not largely present, and did not present a problem for controlling the vehicles in our experiments. Instead, deformation of the vehicle during aggressive maneuvers appeared to be a much more considerable issue, resulting in significantly less stability. In addition, this poses a massive problem for pose estimation and control, as the accuracy of both vehicle orientation calculations and the assumptions of the internal physics model degrade when the frame flexes.

While deformation in flight did prove to be an issue, we found a large number of configurations which successfully mitigated this issue while also maintaining many of the desirable properties of tensegrity structures in this application. Through our simulation efforts as a whole, these systems proved to be relatively robust to changes in their construction, even those with considerable impacts on the system dynamics. The best of our tested configurations were also highly capable of absorbing impacts effectively, and very easily controllable with classical linear methods given a reasonable orientation estimate. These results are incredibly promising, as they show that a physical prototype of a UAS based on such a tensegrity structure is feasible with current technologies, and is quite likely be successful with minimal design iterations.

4.1 Optimal Vehicle Configuration

Experimenting with a large number of noticeably discrete combinations of important variables (n=24), we were able to find some common trends with favorably meaningful consequences on proper characteristics of a prototype. The simplest relationship was unsurprisingly found between cable stiffness and vehicle deformations. In general, stiffer configurations were found to be more stable and experienced less deformation in flight, characteristics which were highly desirable for controllability and aggressive flight. At the same time, however, stiffer configurations compressed significantly less on impact, effectively distributing less of the forces throughout the frame of the vehicle and incurring more jarring forces, severely reducing the benefits of using a flexible tensegrity based frame.

The relationship between propeller thrust/torque ratio and handling of the vehicle was also found to be very important, with low ratio configurations seeing more stability and much less deformation in flight than their high ratio counterparts. Increasing the stiffness of the cables was effective in reducing harmful deformations and increasing the stability of high ratio configurations, but again at the cost of reducing the benefits of the flexible frame. Overall size and mass of the vehicle likewise had a considerable effect on performance, to the point where each could be optimized, but was still observably distinct enough to consider them as completely separate cases. Smaller configurations tended to be more successful with stiff cables, while larger configurations saw similar success with much less internal stiffness. Our results showed that smaller configurations needed stiffer cables in order to prevent rotors from colliding with other frame elements, and also to absorb impacts slightly better. The former result is simple enough to see on its own, but the second was somewhat less predictable under our initial assumption that more flexible models would on the whole see better collision results. It became apparent that if the cables in the smaller system were less stiff there was not enough travel in the cables to absorb an impact, and the system was more liable to fold nearly flat on impact, with many frame elements colliding with the impact surface and each individually experiencing significant jarring forces. With stiffer cables, however, the strength of the cables was enough to prevent other frame elements from colliding with the impact surface, such that only the few points of contact would collide, more effectively distributing the forces and extending the duration of impact.

In the larger configuration, however, travel distance was less of an issue since most frame elements were further removed from the impact surface, allowing a more flexible frame to deform significantly on impact, extending duration of the collision to absorb and distribute more force, in turn protecting potential payloads much better than the small version. There was a trade-off, however, as our experiments showed that the larger, more flexible configuration was significantly less robust to changes in the propeller thrust to torque ratio. The large configuration tended to suffer from in flight deformation, since propellers are located at the end of the rods the larger system has a significantly larger moment arm, which in combination with more flexible cables exacerbated the impacts of higher thrust to torque propeller ratios. This resulted in concerning amounts of deformation with the test ratio of 75, such that in order to make the system seem safely controllable the cable stiffness needed to be greatly increased, again counteracting the potential benefits of the tensegrity based frame.

For both the small and large configurations, however, it was possible to find at least one overall configuration which seemed to optimally strike a balance between maximizing the benefits of the tensegrity structure while also minimizing its drawbacks. These findings give us valuable insights into the potential construction of a prototype vehicle, as it seems there exists a near optimal configuration of the vehicle at each of these possible sizes. At the small configuration with 0.5m rods, an ideal vehicle would have fairly stiff cables, such that the system experiences very mild deformations at rest on the ground, and minimal deformations in the air. This smaller, stiffer vehicle was found to be relatively robust to changes in the thrust to torque ratio, such that it was possible to fly without cause for major concern using all three possible propeller configurations, although lower ratios were still slightly more promising than larger. At the larger configuration, with 1m rods, an ideal vehicle would have considerably less stiff cables, such that it experiences noticeably larger amounts of deformation at rest on the ground, at the additional cost of slightly larger deformations in flight. If a larger, more flexible vehicle were constructed, and based on our results there is a considerable benefit to building such a vehicle, propellers would likely need to have a thrust to torque ratio in the 25-50 range in order to achieve controllable flight without harmful deformations, where the lower this ratio can be driven the better. Overall, it seems that a larger system more effectively takes advantage of the tensegrity properties of the novel frame, showing more interesting behavior and ultimately holding more value in terms of follow up research. Such results send a very promising message about the feasibility of our proposed system, we have arrived at a reasonable initial estimate for the necessary parameters in each of these possible vehicles, and from analyzing the patterns present in our simulations attempts have also found a reasonable set of rules for tweaking a prototype vehicle based on experimental observations to alleviate common issues.

5 Conclusion

Our results show that in simulation a soft quadrotor can be made with a tensegrity structure as a chassis. While small quadrotors require the tensegrity to be highly tensioned to keep the rotors from colliding, larger designs can use less tensioned tensegrity structures, giving them more compliance in the event of a crash. In addition we have provided analysis of the most important physical design characteristics as well as tradeoffs between putting one IMU sensor in the center of the structure as opposed to multiple IMUs positioned on the rods.

References

1. Buckminster Fuller. Tensegrity. *Portfolio and Art News Annual*, 4:112–127, 1961.

2. Atil Iscen, Adrian Agogino, Vytas SunSpiral, and Kagan Tumer. Robust distributed control of rolling tensegrity robot. In *The Autonomous Robots and Multirobot Systems (ARMS) workshop at AAMAS 2013*, 2013.

3. Vytas SunSpiral, Adrian Agogino, and David Atkinson. Super ball bot - structures for planetary landing andexploration, 2015.

4. Andrew P Sabelhaus, Jonathan Bruce, Ken Caluwaerts, Yangxin Chen, Dizhou Lu, Yuejia Liu, Adrian K Agogino, Vytas SunSpiral, and Alice M Agogino. Hardware Design and Testing of SUPERball, a Modular Tensegrity Robot. In *Proceedings of the 6th World Conference of the International Association for Structural Control and Monitoring (6WCSCM)*, 2014.

5. Steven Lessard, Pattawong Pansodtee, Ash Robbins, Leya Breanna Baltaxe-Admony, James M Trombadore, Mircea Teodorescu, Adrian Agogino, and Sri Kurniawan. Crux: a compliant robotic upper-extremity exosuit for lightweight, portable, multi-dof muscular augmentation, 2017.

6. C. J. Salaan, K. Tadakuma, Y. Okada, Y. Sakai, K. Ohno, and S. Tadokoro. Development and experimental validation of aerial vehicle with passive rotating shell on each rotor. *IEEE Robotics and Automation Letters*, 4(3):2568–2575, 2019.

7. Jiaming Zha, Xiangyu Wu, Joseph Kroeger, Natalia Perez, and Mark W. Mueller. A collision-resilient aerial vehicle with icosahedron tensegrity structure. In *IROS 2020*, March 2020.

8. Nitya Ravi, Neil Kaushikkar, Xander Zell, Felix Chen, Julia Clark, Alexandra Saccente, Gregor Limstrom, Sierra Dabby, Ryan Ehlmann, Chloe Greenstein, Bjorn Johnson, Michael Nguyen, Philip Quijano, Ajay Sunkara, Annie Vu, Leonardo Bonanno, Nick B. Cramer, and Parimal H. Kopardekar. *Weighted Relative Performance Ranking Function for Optimizing Reconfigurable Drones.*

9. CannonPhysicsEngine. https://schteppe.github.io/cannon.js/. 2016.

REPORT DOCUMENTATION PAGE						*Form Approved* *OMB No. 0704–0188*

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE *(DD-MM-YYYY)*
01-03-2021

2. REPORT TYPE
Technical Memorandum

3. DATES COVERED *(From - To)*

4. TITLE AND SUBTITLE
Feasibility Analysis of Unmanned Aerial Vehicle Based on Tensegrity Structure

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Bjorn Johnson, Adrian Agogino

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
NASA Ames Research Center
Moffett Field, California 94035

8. PERFORMING ORGANIZATION REPORT NUMBER
L–

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSOR/MONITOR'S ACRONYM(S)
NASA

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
NASA/TM–20210013185

12. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified-Unlimited
Subject Category 63
Availability: NASA STI Program (757) 864-9658

13. SUPPLEMENTARY NOTES
An electronic version can be found at http://ntrs.nasa.gov.

14. ABSTRACT
For future unmanned aerial vehicles to be ubiquitous, they need to be safe to other aircraft and people on the ground even in the event of unintended collisions. Unfortunately, this is difficult for current rigid designs as a crash of all but the smallest aircraft will have a significant impact on its target. As an alternative, this work explores using a soft tensegrity structure (interconnected cables and rods) as a lightweight, robust chassis for a UAS and analyzes its ability to achieve stable, controllable flight and mitigate damage during a collision. These structures are able to undergo severe deformations without permanent damage. Experiments and analysis are performed using a simple physics simulation environment. We test multiple configurations of the vehicle to understand impacts of changing various physical characteristics, and take limitations such as modern propeller thrust coefficients and motor speeds into account to ensure our results are in line with those of a real world vehicle. We find that such a vehicle has considerable real world potential, and that a physical prototype could theoretically be constructed and flown with current technology.

15. SUBJECT TERMS
robotics, tensegrity, flexible, multirotor, UAS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			STI Information Desk (help@sti.nasa.gov)
U	U	U	UU		19b. TELEPHONE NUMBER *(Include area code)* (757) 864-9658

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18