Based on: *The Zero Pressure Conundrum*

Or **What Happens Next for 30-50% More Life On Your Superpressure Balloons**

Loon LLC Artifact

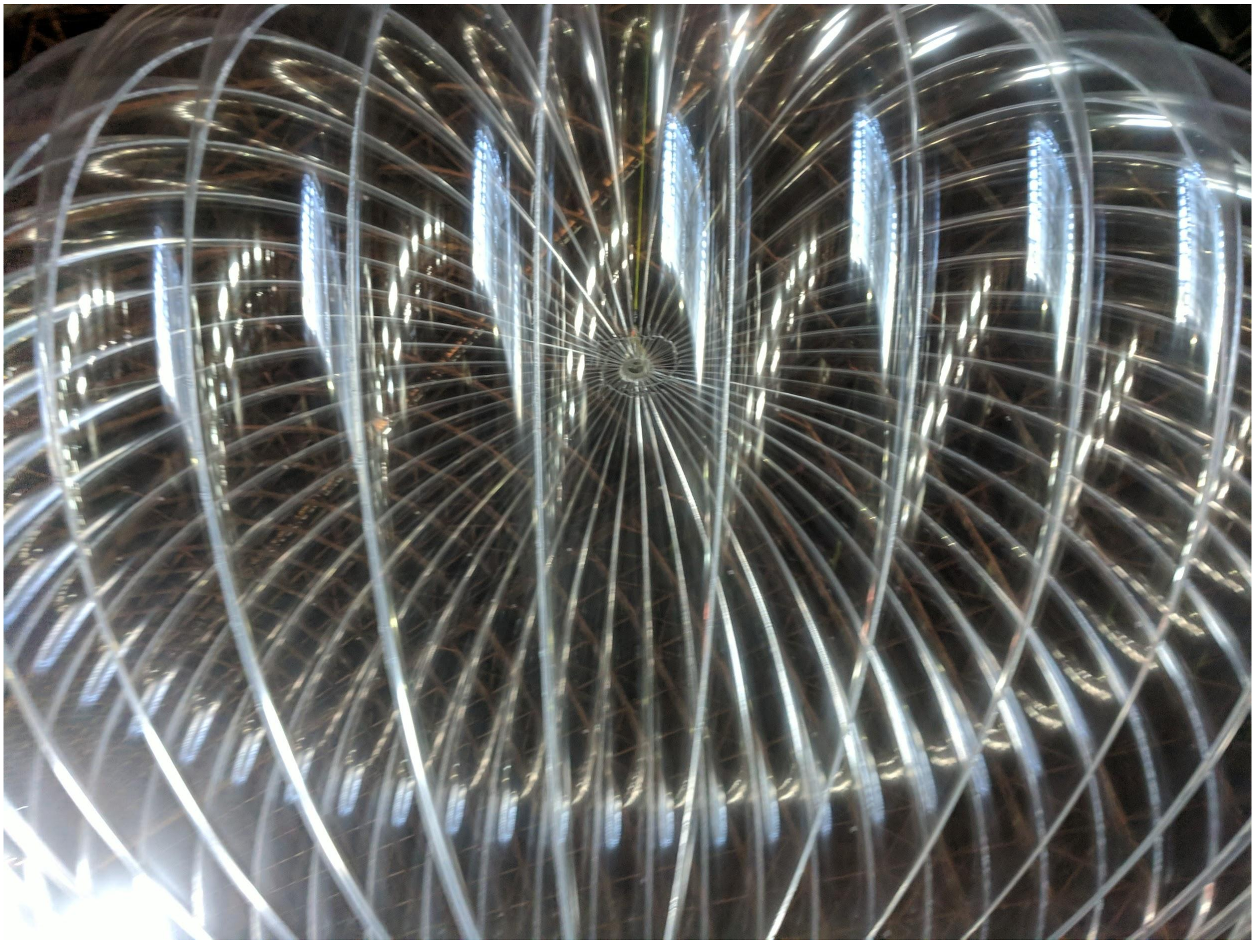

09/14/2017

Last year, we implemented ballast to allow us to turn our “bad” balloons into “ok” balloons with zero pressure life extension and to improve overall balloon health by lowering the pressurization rate as the balloon entered float. Since then we have encountered some new problems and edge cases. We also discovered that the required design limits on our envelopes as built are less than we had initially thought. The need to fly our balloons on thin margins against all of these edge cases is complicated enough that we will want to automate ballast drops according to principles expounded upon below:

If I were to try to explain with only two bullets, I would say:

- Not every zero pressure is a sign of end of life. No matter how healthy, a cold enough night combined with a low enough flight altitude will always cause a zero pressure event.
- Allowing the balloon to zero pressure *a little bit* and recover ends up giving an unexpected number of extra flight days because our balloons typically experience very cold nights on an infrequent basis. This technique is viable for reverse ballonet due to a reduced amount of tilt but yy (Yin-Yang ballonet architecture) requires more care to avoid low pressures at low altitudes.

What situations are bad for balloons?

Superpressure / Supertemperature (Upper Design Limit)

Previously, our design limit was a function of max operating pressure and we would ascend when we flew too close to this value. This may not be the correct solution in all cases and the superpressure limit alone is insufficient to describe how close the film is to its limit.

Our envelope film becomes stronger at lower temperatures, so flying lower (where it is colder) will actually make the balloon system stronger. This helps to compensate for the added pressure from a higher number of mols of gas expanding into a fixed volume[1].

---

[1] This is far too basic an explanation to be useful for anything other than this document.

A distribution of proximity to a film stress design limit (x-axis) vs altitude for Prod OLs:

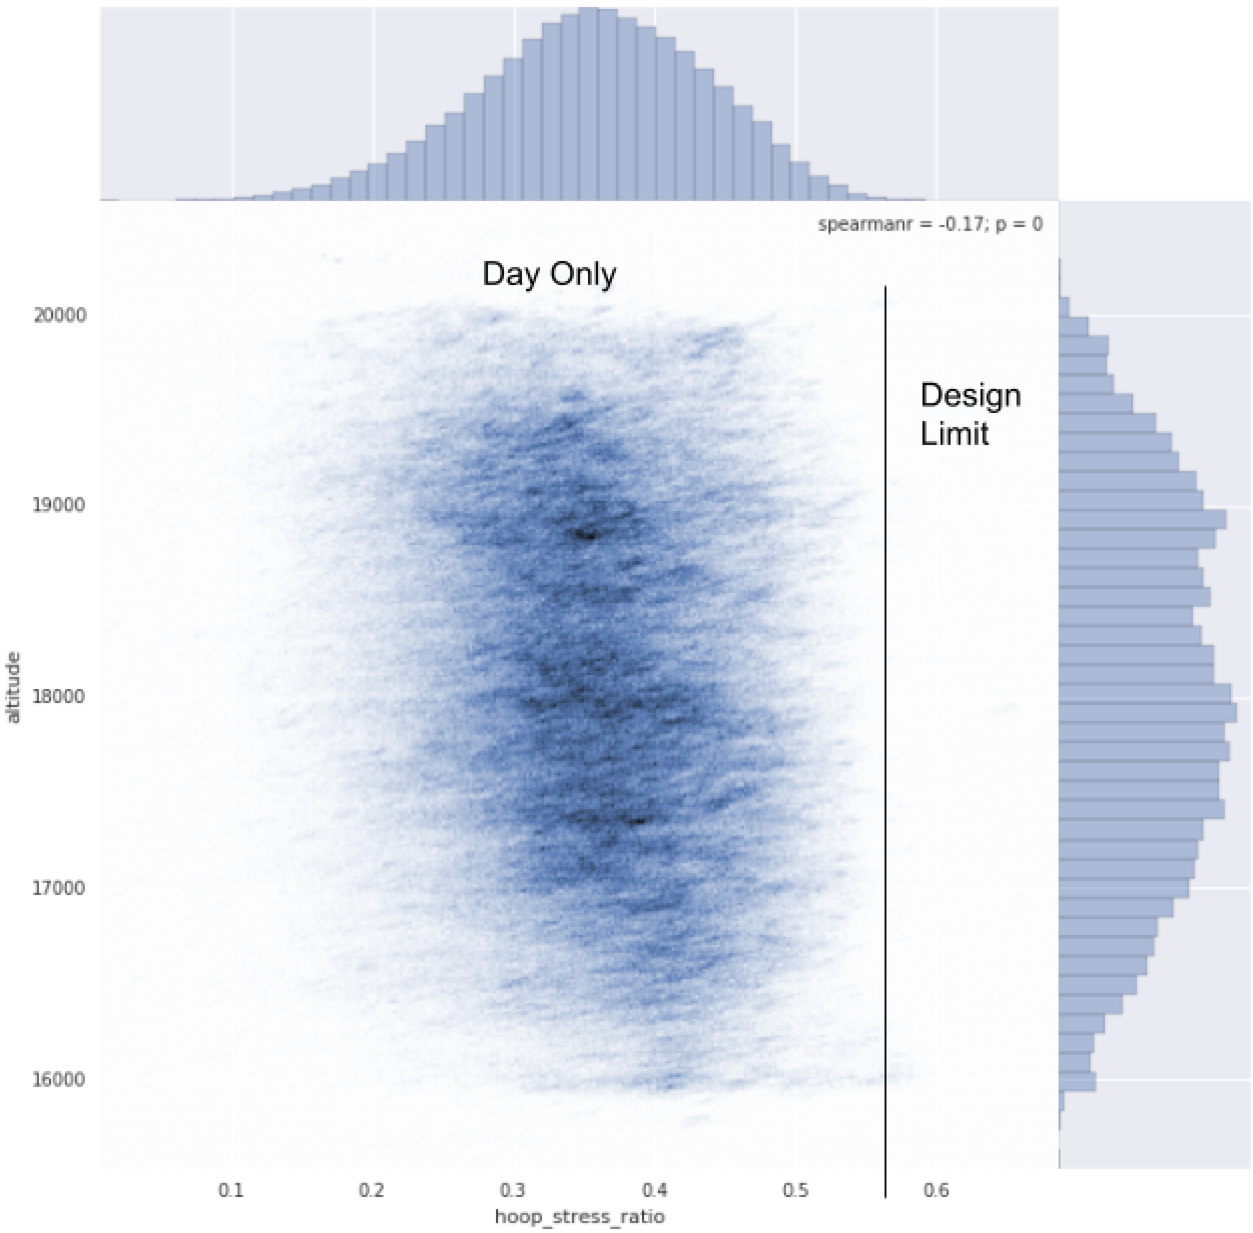

This is a combination of the following 2 effects of temp vs altitude and pressure vs altitude:

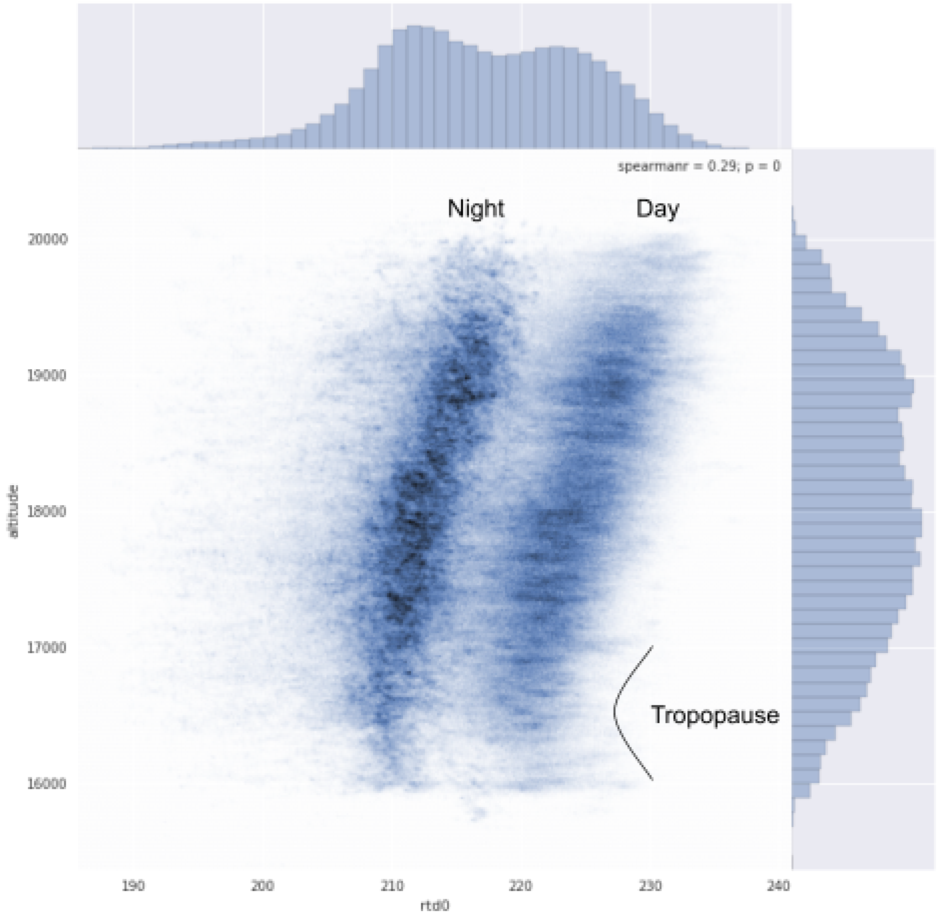
spearmanr = 0.29; p = 0
Night
Day
Tropopause
20000
19000
18000
17000
16000
altitude
190
200
210
220
230
240
rtd0

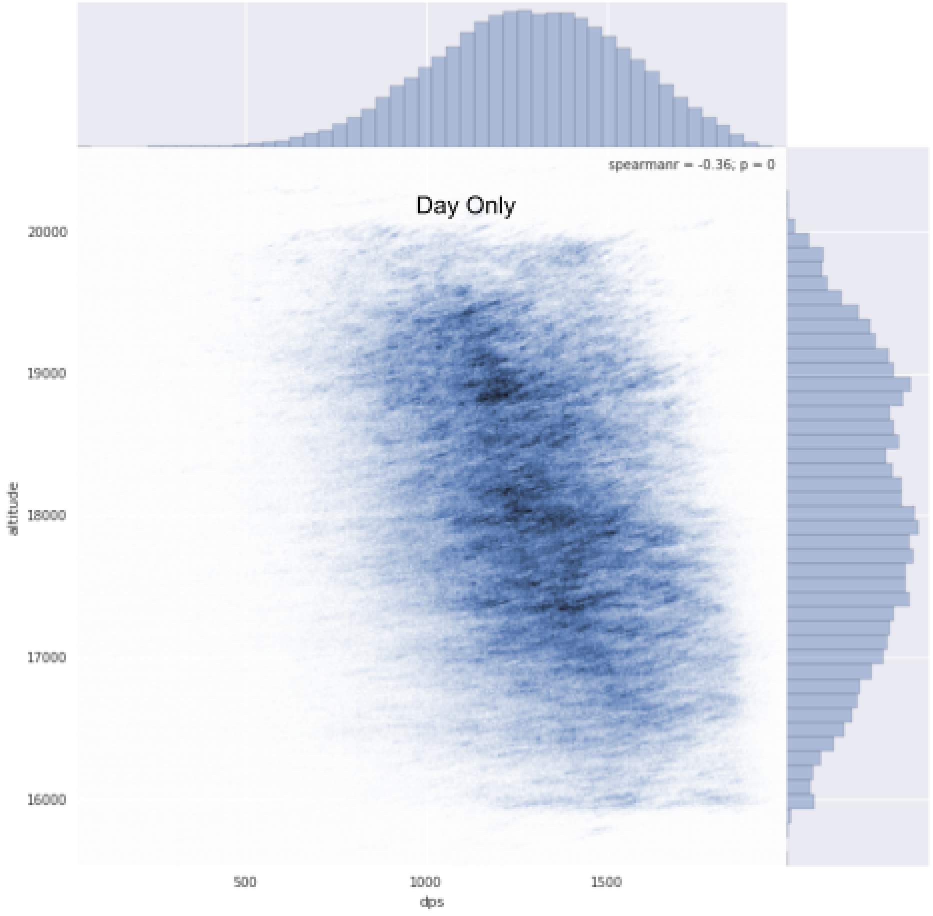
spearmanr = -0.36; p = 0
Day Only
20000
19000
18000
17000
16000
altitude
500
1000
1500
dps

99th Percentile of Temperatures (Lower Design Limit)

Thermal runaway (below) is hard to design for, so, for system design purposes, we use the 99th percentile of temperatures and zero pressure as the end of life for a balloon system. Previously, we set this at 1 km down from float altitude but may in the future set the requirements differently such that we can ride out storms at lower altitudes. This is around 195K or -78C.

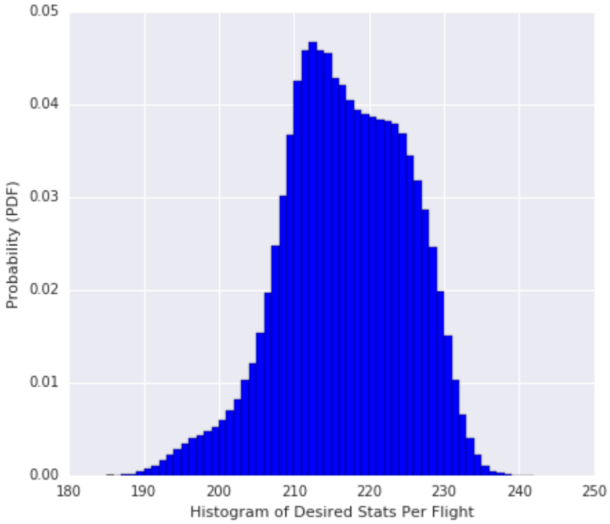

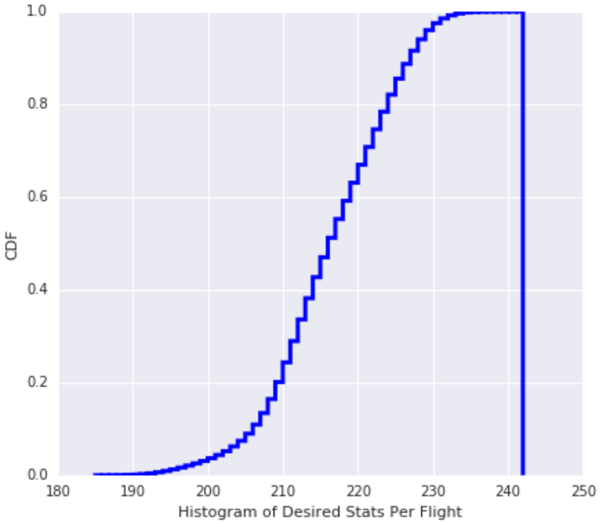

Proximity to the Cloud Deck (Edge Case)

Being above storms is bad. They are cold (see above) and often full of lightning and electrical gradients. They also create turbulence in the vicinity that may damage the balloon directly or cause the payload to impact the envelope. We should try to figure out where these are, are going to be, and then steer clear of these if we can.

Zero Pressure Leading to Thermal Runaway (Lower Operating Limit)

The initial statement with ballast was that zero pressure events with racooning (altitude oscillations due to self-arresting thermal runaway) are ok. This is still likely true for the envelope shell but is not proven to be 100% harmless, especially in the case of extremely long lived balloons. I think for most cases, this assumption is valid. Thermal runaway is the lower operating limit for the system.

A zero pressure event almost always occurs at night at low envelope temperatures:

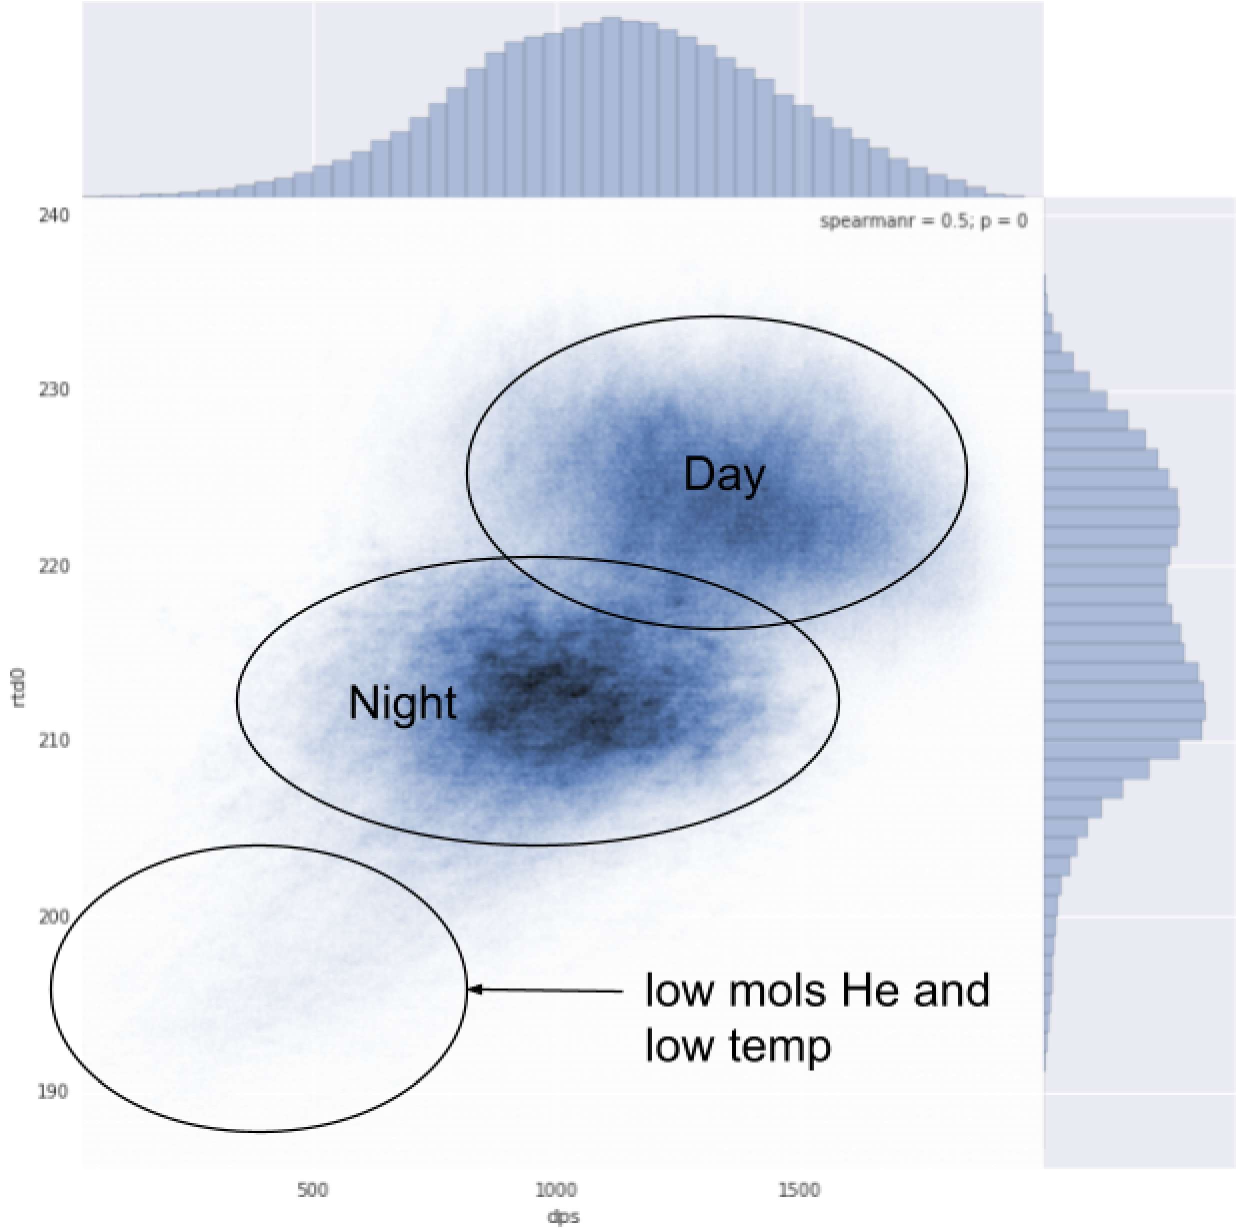

Tilt and Payload Contact with Envelope (Edge Case)

When a balloon gets close to zero pressure while maintaining 40 degrees of tilt, a case that happens at low altitudes, our payload can puncture the balloon. This is due to the payload weight rotating the base plate into the balloon when there is very little load in the tendons. We need to avoid this case with yin-yang balloons. It should not be an issue with the reverse ballonet design.

Tilt data (in radians 0.7~40 degrees) vs altitude for OL yy designs:

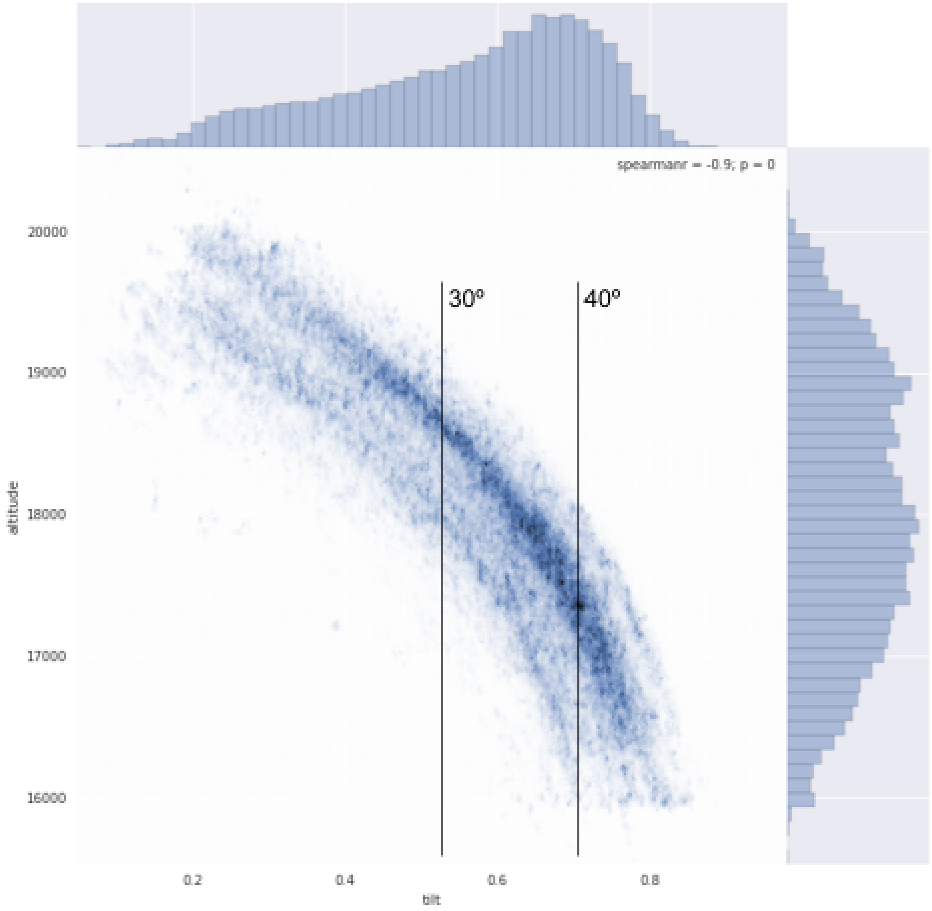

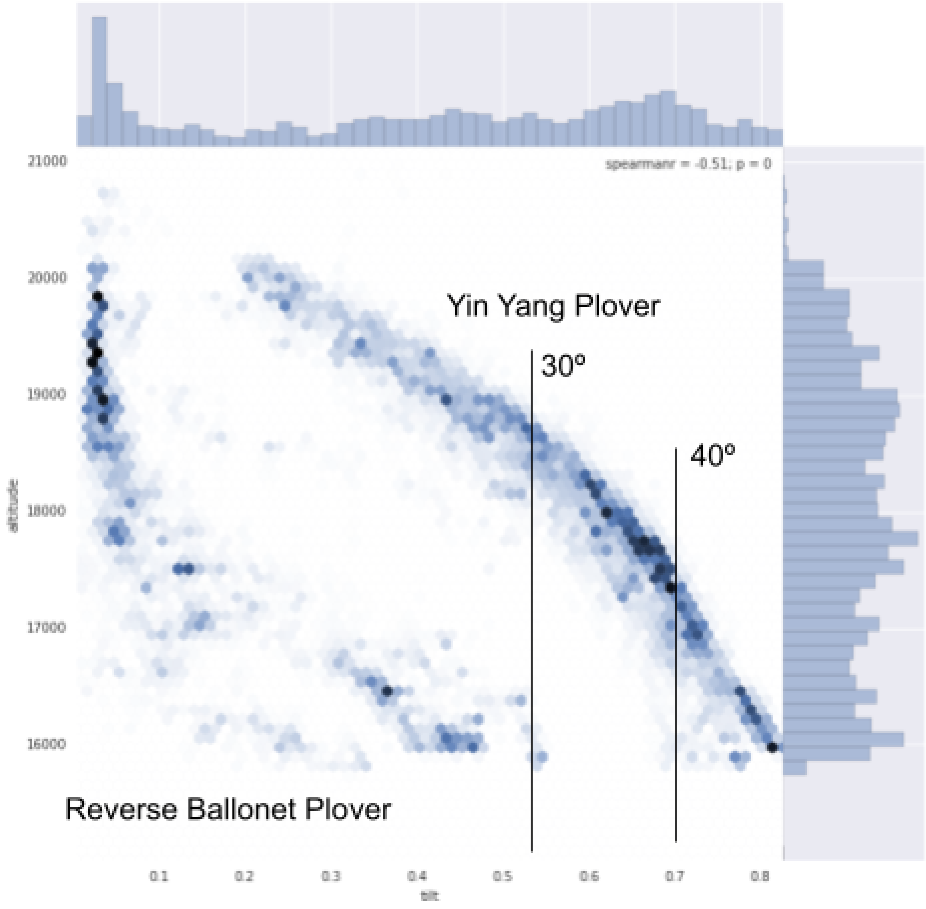

On an OL ying-yang balloon system this translates into roughly the following requirement:

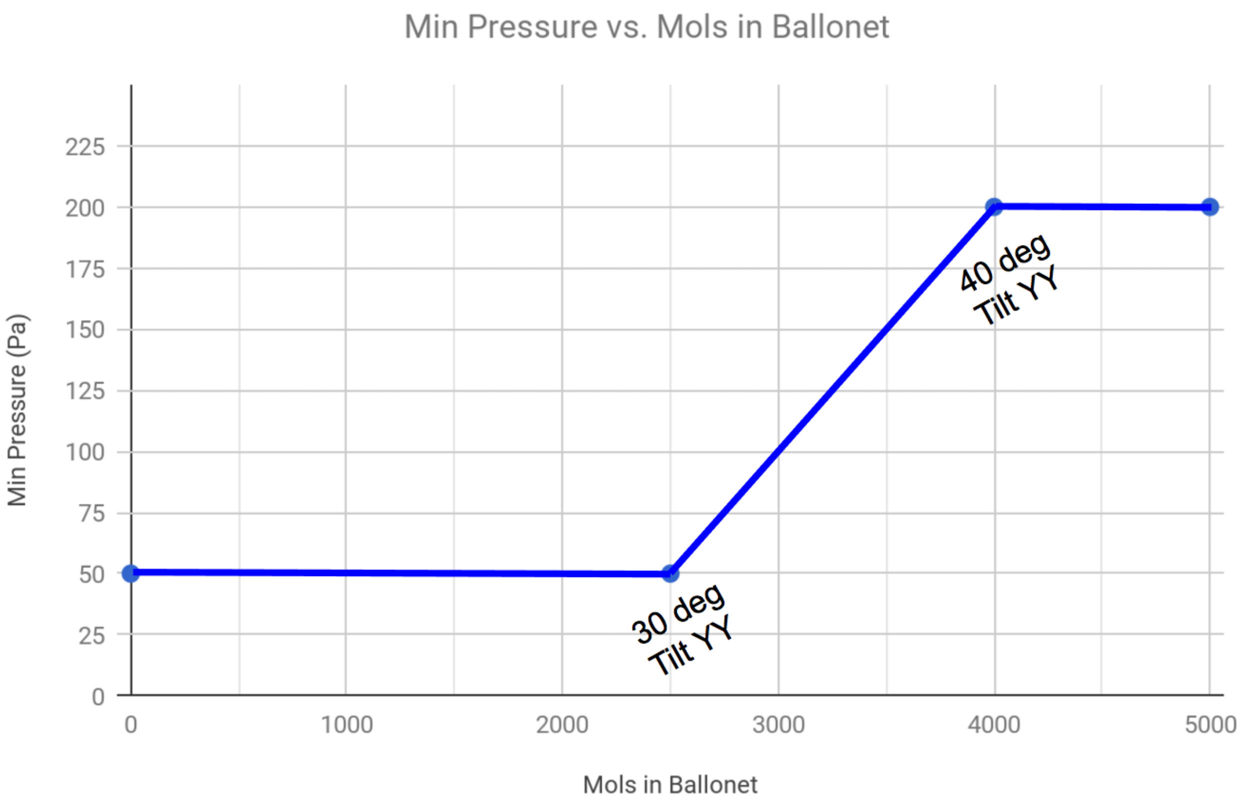

Stay Between the Lines

What is the optimal way to keep the balloon pressure as low as possible while minimizing racooning events for the first portion of the balloon's life? How much ballast do we need for our get out of jail free card to work and when do we say the balloon's life is effectively over?

The Lower Design Limit (approximated with est.zero_pressure.static -17)

The Lower Operating Limit (approximated with est.zero_pressure.dynamic)

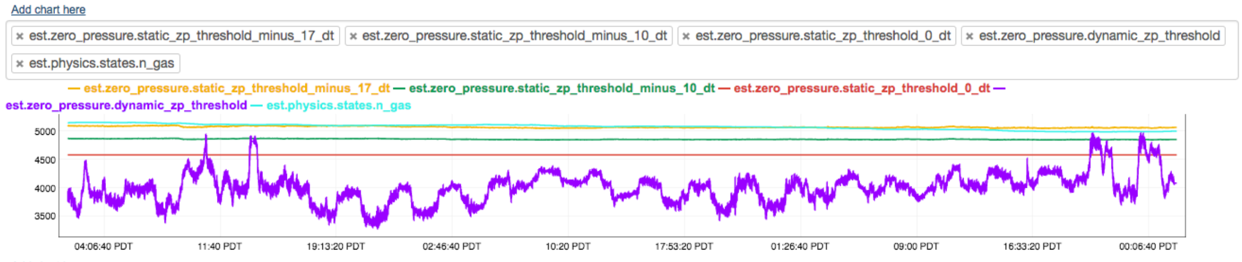

The blue line in Figure 1 is the plot of mols of lift gas in the envelope over time. Our balloons lose lift gas so this number slowly decreases. The yellow, green, and red lines are molar thresholds at which the balloon will lose superpressure if the supertemperature reaches -17, -10, and 0, respectively. The way to think about this is that if the number of moles of lift gas (yellow line) falls below the -10 threshold (green line), the flight no longer has the ability to survive a night where the temperature falls below approximately 203 K or the 95th percentile. The -17 threshold (yellow line) marks the 99th percentile or 195K and the lower design limit.

The purple line is the dynamic zero pressure threshold. It plots the number of moles of lift gas we must have had to maintain superpressure given the lift gas temperature the balloon was experiencing at that point in time.

When the blue and yellow lines intersect we have hit our lower design limit. When the blue and purple lines intersect we experience a zero pressure event.

Creating lines for static or dynamic_zp states at different altitudes may also be useful.

The Upper Design Limit (hoop_stress_ratio < 0.5773)

A ballast drop can be executed safely when doing so will not cause the balloon to exceed its upper design limit. If we create a state predictor like the zp_static_threshold but for how close the balloon is to its design limit in the 95th and 99th percentile cases of temperature and pressure, staying between the lines will be possible.

Hoop_stress_ratio is a function of temperature and pressure and the limit is currently written based on a number of factors to include film thickness variation, estimated stress concentrations, creep over time, UV degradation, and factor of safety.[2]

The current upper design limit will need to be added to the balloon fragments so it can be adjusted based on our best engineering analysis of the tolerance stack for balloon film.

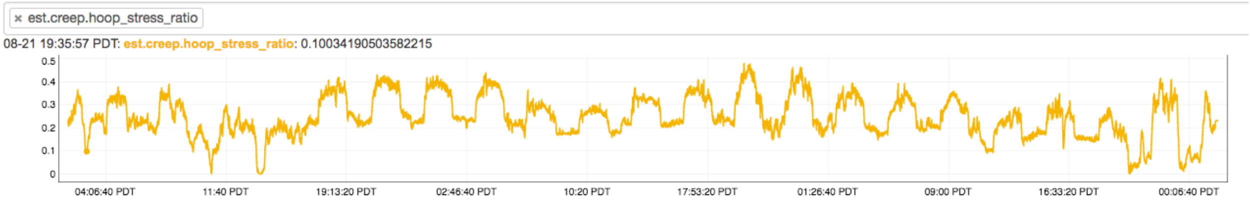

Automation

Automation will be added to balance avoiding potential thermal runaway against long term stress on the balloon envelope. In this case, thermal runaway is considered an uncontrolled loss of altitude that has substantial impact to operations, e.g., falling below the air traffic control flight floor for our flights, and not any zero pressure event.

We will add a precomputed lookup for the expected probability of stress beyond a threshold -- either accumulated or peak -- on the envelope over time as a function f of the current envelope state of the form

$$P[\text{bad stress threshold}] \sim f(\text{mols, expected time left, system mass})$$

---

[2] Still requires full engineering buy-in on the exact value and another large doc

The function f will be computed by Monte Carlo for each class of flight system and account for expected upwelling IR, temperatures, etc., the flight would be historically expected to encounter.

We will then evaluate a decision rule of

```
If P[thermal runaway] > k * P[bad stress threshold]:
    Drop some ballast
```

This works ballast release mechanisms that can handle many small quantity releases, as well as if only a single large release is supported. This decision rule is also self-correcting in the sense that the closer we get to beginning to thermal runaway the better we are able to estimate this probability.

Here we are effectively saying that we should only drop ballast when we must, and then adjusting our risk tolerance (how fast we are to proactively drop ballast) based on how bad we expect the stress to be on the envelope later in life. If the flight is later in life and lower on mols then we expect we will be fine without ballast. In contrast, on the first day of the flight we should be quite cautious about dropping ballast as it will almost certainly hurt us soon after despite surviving the night. We always prioritize surviving the night, i.e., we will never drop a balloon out of the sky to preserve ballast.

The Future is Now

The Zero Pressure Conundrum was a grand scheme to get more with less. Where are we now?

- **Ballast** is now a standard component of our fleet
- The flight engineers have developed a **standard operating procedure** around ballast.
- We still need to **gather data** on zero pressure and thermal runaways.
- With that data we need to **improve our model** and prediction of zero pressure and thermal runaway. Much of this is around better understanding of the errors in our weather forecast models precisely.
- We are hard at work making a **ballast mechanism with multiple controlled drops**
- **We have better sensors**, ambient air temperature in particular, that feed our leak estimates and predictive models that in turn removes more margin from our operations and get even more days of life per balloon.