Rocket vehicle loads and airframe design

Rick Newlands, Martin Heywood, and Andy Lee

Introduction

Rocketeers have generally followed a Darwinian evolutionary approach to designing their vehicles to withstand flight loads: if it breaks, try something else! This approach has worked remarkably well, but aircraft designers have better methods that allow prediction of the airframe loads *before* flight, especially for very high speed/high altitude flights that come to grief too often.

The major disturbances affecting the trajectory are thrust misallignments and winds. The wind, especially upper-atmosphere wind (such as the jet-stream) is by far the more significant of the two, as it affects the flight path, causing vehicle airframe bending and trajectory deviations.

A lot of the missile design literature from the 1960's has now been declassified, so we can see how they were designed against wind effects.

Recovery system (parachute) loads are not covered here; see our paper 'Recovery system design for large rockets' for these loads.

Part 1: Winds, atmospheric turbulence, and gusts

In addition to the compressive loads caused up the fuselage by the rocket motor's thrust, the vehicle is subjected to aerodynamic forces during flight which significantly affect the airframe design**.**

As well as the axial loads along the airframe caused by drag forces, surprisingly large forces normal to the centreline of the airframe (lift forces) can be created by gusts and other transient wind effects:

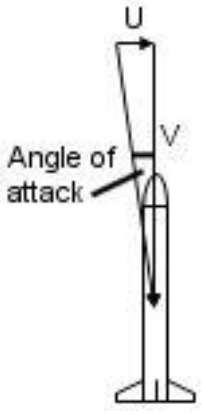

The sudden gust velocity vector 'U', when added to the vehicle's airspeed vector 'V', causes a small angle of attack α:
It's this angle of attack that creates normal forces that create the side-loads within the fuselage.
Rotational accelerations are also caused as the fins try to reduce the angle of attack α back to zero.

Ground winds:
Our rocket vehicles are too small to be broken just by the wind as the vehicle leaves the launcher, because we never launch in overly strong winds. But it's important to know how strong the wind will get as the rocket gains altitude, as the magnitude of any gusts it encounters depends on the strength of the wind.
The wind gets stronger with altitude as the rocket ascends the wind's boundary layer or 'wind gradient'.

An empirical engineer's formula for predicting the wind at altitude is: $V2 = V1\left(\frac{h2}{h1}\right)^P$ where

V1 is the wind measured at height *h1* near the ground, and *V2* is the wind speed at altitude *h2*. The exponent, *P* is unfortunately terrain dependant, but for the wide open-spaces around launch-pads the following rough values will do:
If *V1* is measured at *h1* = 3 metres above the ground and two-minute averaged readings are taken, then *P* = 0.2 if *V1* is less than 17 metres per second and *P* = 0.14 if *V1* is greater than 17 metres per second

This formula can be used up to a height of about 150 metres, which is typically the edge of the boundary-layer: above this height, the wind is constant with further altitude (though see below).

Wind shear and the synthetic wind profile:

Wind-shear is the rate-of-change of average wind speed with altitude: (it's a vector quantity, so includes changes of wind direction with altitude too.) Note that due to fine wind structures (nowadays suspected to be fractal) the wind-shear values increase the closer you examine them (vertical scale of distance measured e.g. by weather radar).

A sudden change in wind speed and/or direction causes angles of attack on the airframe which generate large side loads.

The first large wind shear occurs when the vehicle becomes free of the launcher, and encounters the wind, but the vehicle's airspeed is still very low so this doesn't load the airframe unduly.

Our 2nd flight of our FLARE hybrid was hit by a windshear after burnout but well before apogee; it caused the FLARE to tumble.

In order to judge the effect of windshear, we first have to know the strong winds at high altitudes, as these worst winds cause the worst windshears.

Various types of models representing the steady wind have been used to design launch vehicles, such as statistical, stochastic (power spectral), and synthetic models.

By far the easiest to develop and use for essentially vertical trajectories (therefore not for aircraft) is the synthetic wind profile.

Early American Launch vehicles and Britain's Black Arrow were all designed using synthetic wind profiles.

Construction of a synthetic wind profile starts with an altitude plot of worst-case winds. The plot for Cape Canaveral is shown here, and was compiled from quasi-steady (averaged over at least two minutes) wind values obtained from weather radar instruments at the site over several years.

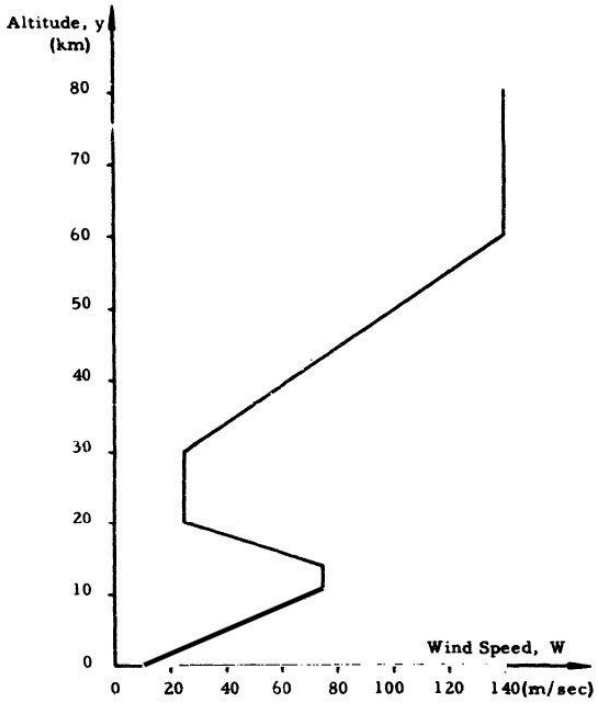

Ninety-five percent probability-of-occurrence wind speed profile envelope for Cape Canaveral, Florida

The 95 percentile wind plot shown here describes wind speeds that have only a 5% risk of being exceeded for any launch trajectory occurring at any time throughout the year. Note that the plots are for scalar wind speeds, i.e. the assumption is that all winds are blowing in the same arbitrary direction, which we can specify to be the worst-case vehicle loading direction. In reality of course, winds blow from all compass directions, though are still assumed to blow horizontally, as vertical winds have insignificant effects on Launch Vehicles.

In tabular form:

Altitude (km)	0	10	14	20	30	60	80
Speed (m/sec)	11	75	75	25.4	25.4	141.2	141.2

Note the first wind peak between 10 and 14 Km, this is common all around the world and describes the jet stream, a high-speed ribbon of air. The worst wind shears tend to occur here (Ref. 3) so this region tends to be the one that the vehicle structure has to be designed for.

To the wind profile, we then add a severe wind shear, usually at the same altitude as the strongest winds.

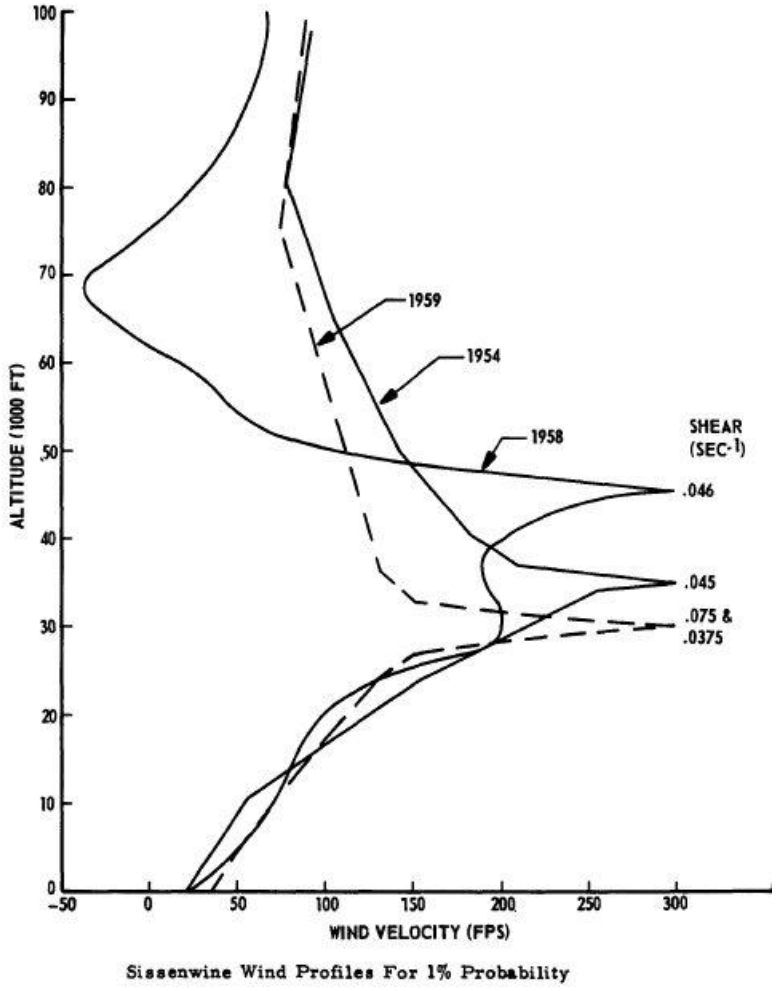

Sissenwine Wind Profiles For 1% Probability

Here are three synthetic wind profiles (Ref. 3) devised by the meterorologist Norman Sissenwine for windy regions of the USA. The latter profiles he released a few years after the first when better weather radar became available.
Note the windshears (where the curves are nearly horizontal: shears are the inverse of the curve gradient) at around 40,000 feet (12 Km) in all three profiles.
In the 1959 profile, two shears are noted: a lower one (0.0375 fps/per foot) for a 3000 foot investigation by radar, and a higher shear (0.075 fps/per foot) for a 1000 foot investigation. The lower shear is for the lower half of the point (wind increasing to 300), followed by the sudden drop in windspeed (the higher shear of 0.075).
Each profile is for a 99% probability: actual weather is expected to exceed these profiles (wind and windshear) only 1% of the year.

In tabular form, the 1959 Sissenwine profile is (converted to metric):

Altitude (km)	0	5.3	7.5	8.2	9.3	9.9	11.1	23.0	30.5
Speed (m/sec)	10.9	30.6	39.9	46.1	89.7	46.0	39.9	22.4	27.1

You could replace the zero-altitude data in the table above to suit your needs, for example input your launchsite's elevation above sealevel, and use the United Kingdom Rocketry Association's maximum allowed launcher windspeed of 8.94 metres/sec (20 miles per hour).

The Sissenwine profiles are a simple design tool that we can use for high altitude flights, and are easily fed into a trajectory simulator. Many vehicles including the British Black Arrow, used a Sissenwine profile for their design even though Black Arrow was flown over Australia rather than the USA.

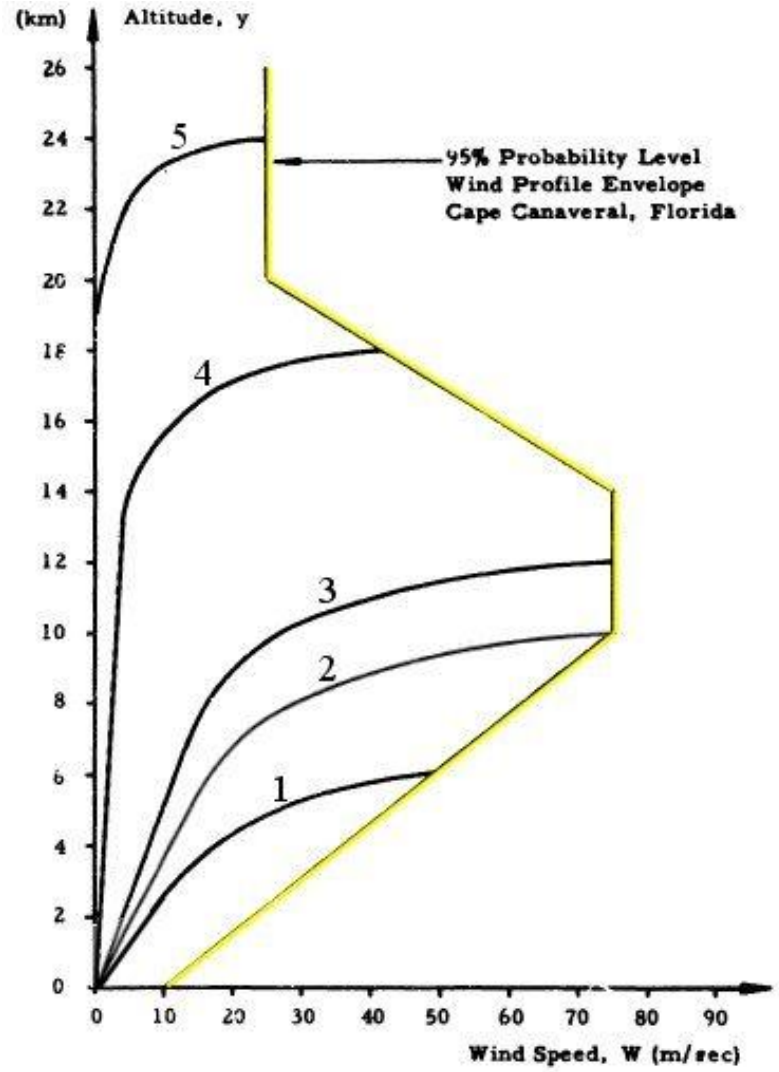

Example of synthetic wind profiles based on ninety-nine percent wind build-up rates to be associated with the ninety-five percent probability wind speed profile envelope at 6, 10, 12, 18 and 24 km altitudes.

Here is another, later, set of wind profiles (1 to 5) for Cape Canaveral, that were used to design early versions of the Saturn 5.
The yellow curve is the worst case (95%) winds that we saw earlier.
To this have been added severe wind shears (99%) at various altitudes (curves 1 to 5).
The wind speed percentile used is lower than the wind shear percentile used simply because the wind speed can be measured quickly by radar just before launch as a safety-check, and the launch halted if the wind is too high, whereas severe time-consuming number-crunching is needed to extract the wind-shear from radar data.
Since the chances of the worst-ever wind speed occurring at the same time and place as the worst-ever wind shear and worst-ever gust are fairly remote, the above Synthetic profiles are conservative.
So NASA reduces the shear and gust values by 15%, based on engineering judgment, before applying them to a sim. If curve 1 above was reduced by 15% then its gradient would be 15% steeper (lower shear) at all points.

Suppose, for example, that burnout (maximum airspeed of your vehicle) occurred at around 6 Km altitude (20,000 feet) so that you are interested in the effect of a windshear at 6 Km. Then you would use curve 1, which adds a windshear that meets the yellow wind curve at a height of 6 Km. Note that as you ascend curve 1, the gradient of the curve flattens, showing that the windshear is getting worse the closer you get to the yellow curve. This describes the aforementioned fact that the windshear gets more severe the smaller the length scale you examine it at.
Once curve 1 hits the yellow wind curve, you then just follow the yellow curve upwards from there.
Construction of the shear curves (1 to 5 above) is described in appendix 1.

Gust loads:
The wind shears are at their worst only at high altitudes, so probably can be ignored for low altitude flights providing your vehicle is sufficiently aerodynamically stable. Gusts, however, occur at all altitudes so have to be designed for.
The highest aerodynamic loads suffered by the vehicle's airframe occur at 'max q' where q is dynamic pressure (see glossary), which varies with the square of the vehicle's airspeed. Typically, the highest airspeed, and so the highest max q, occurs just at burnout.
If the vehicle is hit by a strong side-gust at max q, then the aerodynamic loads caused by the gust can cause structural breakup of the vehicle, if not designed for.
The gust is pictured as a horizontal layer in the sky that the vehicle ascends into.

Gust penetration effects:
The simplest gust representation is the immersion gust, where the vehicle is small enough compared to the gust wavelength (vertical length) that it can be considered as a point, and the appropriate aerodynamic loads are obtained instantaneously as the point traverses the gust. Our vehicles are small enough to use the immersion gust assumption.
However, a more accurate analysis includes penetration effects. This is known as quasi-steady gust penetration. 'Quasi-steady' means that aerodynamic inertia effects are ignored, and steady aerodynamic coefficients are instantaneously reached as each fuselage station passes the gust front. The vehicle is represented as several lifting sections (nosecone and fins) separated by essentially non-lifting sections (fuselage) in between.
There is therefore a time lag between the nose entering the gust and the fins entering the gust, i.e. the nosecone starts lifting before the fins do.

Modelling gusts
There are two approaches taken by aeronautical engineers for modelling gusts: a discrete gust, and continuous turbulence.
In the former, a single strong gust is added to the synthetic wind profile just as the vehicle reaches max q. This was the approach taken with the design of the UK Black Arrow launcher for example.

NASA has noticed particularly nasty sharp-edged gusts of 9 metre/sec amplitude above the steady wind, which represent convenient worst-case gusts.

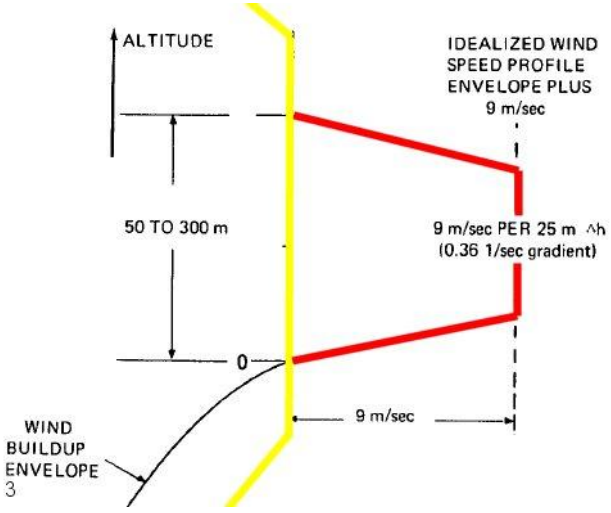

The Synthetic profile used by the Marshall Space Flight Centre for Saturn 5 design work took the 95 percentile wind and 99 percentile wind-shear (see previous pages), then slapped a 9 metre/sec gust onto it at the altitude of interest: here a gust (in red) has been added to windshear curve 3 from the previous diagram.
Note the buildup rate at the leading and trailing edges of the gust of 9 meters per second per 25 meters altitude.
Using engineering judgement, NASA then reduced the gust amplitude by 15% to make it a 7.65 meters per second gust.
Note that your sim must increment in height steps much less than 25 metres or it could miss the gust buildup curves and experience an 'infinite' windshear.

For Black Arrow's design, a 50 feet per second (15 metre per second) sharp-edged gust was added to the Sissenwine profile at max q (Ref. 7). This gust profile was rectangular, so was in effect simulating an infinite windshear (zero graph gradient) though their sim just regarded it as a perfectly valid step-impulse input. This was known to have been rather conservative, but on the other hand, the vehicle structure performed flawlessly.

I suggest that a less conservative gust might be 15 metres per second, but use the Sissenwine worst windshears of 0.0375 per sec and 0.075 per second for the leading and trailing edges of the gust.

The second approach adds a continuous stream of random gusts to the wind profile throughout the flight, and attention is paid to the flight around max q.
In the aerospace industry, several well-tried formulae have been developed to model gusts. A continuous series of gusts is known as turbulence.
The effect of turbulence can be modelled by summing the steady flow of air and a random, zero-mean turbulence velocity. Two central aspects of the turbulence velocity are the amplitude of the variation (the strength of the turbulence, its wind speed) and the frequencies at which they occur. Therefore a reasonable turbulence model is achieved by generating a sequence of random numbers that produces a sequence with a similar distribution and frequency spectrum as that of real wind.
Several models of the spectrum of wind turbulence at specific altitudes exist.
Two commonly used such spectra are the Kaimal and von Karman wind turbulence spectra:

$$Kaimal: \quad \frac{S_u(f)}{\sigma_u^2} = \frac{4L/U}{(1+6fL/U)^{5/3}} \qquad Von\ Karman: \quad \frac{S_u(f)}{\sigma_u^2} = \frac{4L/U}{(1+70.8(fL/U)^2)^{5/6}}$$

Here *Su(f)* is the spectral density function (whatever that is!) of the turbulence velocity and *f* is the turbulence frequency, σ is the standard deviation of the turbulence velocity (more on this below), *L* is a length parameter and *U* is the average wind speed.
Ref. 1 says that you can ditch the '1' on the denominator of either equation with very little loss of accuracy (4%). Cancelling out the resulting equations gives:

$$Kaimal: \quad \frac{S_u(f)}{\sigma_u^2} = \frac{1}{4.953f^{5/3}} \qquad Von\ Karman: \quad \frac{S_u(f)}{\sigma_u^2} = \frac{1}{2.665f^{5/3}}$$

Or in other words, both equations are proportional to $\frac{1}{f^{5/3}}$
so both equations can be reasonably modelled by utilizing white noise that has been altered by a suitable digital filter to give a spectrum of $\frac{1}{f^{5/3}}$ (this is now called pink noise).
Appendix 2 shows how to code this in C++ software.

The use of the standard deviation σ tells you something about the nature of turbulence. Being a natural phenomenon, it has the property that if you were to measure a sequence of gusts over a few hours and plot on a graph how often a gust of a particular strength occurs, you'd find that they have a Gaussian distribution as shown here:

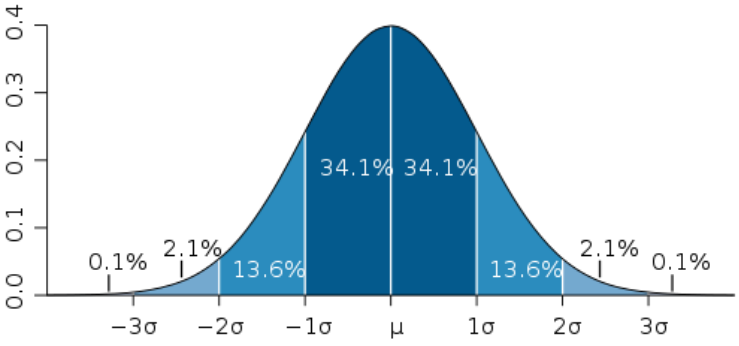

Or in plain English, the majority of gusts are weak, and cluster around the average wind speed μ. 68.2% of gusts occur within a width on the x-axis defined as between + and – one standard deviation σ, with only an extremely rare occurrence of a gust stronger than three σ.
In structural engineering, they define the 'turbulence intensity' as $I = \frac{\sigma}{U}$ which is the percentage that one standard deviation of the gusts is of the average wind speed.

Rearranging gives the gust strengths: $\sigma = IU$

Buildings will break above an ***I*** of 20%, so 15% would seem a suitable design maximum for our vehicles; this'll provide a safety-margin as we wouldn't fly in weather that gusty.

The whole point of feeding pink noise turbulence into a simulator is that the noise encompasses a whole range of frequencies, one of which should excite the natural pitching frequency of the finned vehicle and produce a large angle of attack response (this pitching frequency changes with airspeed, and loss of mass out the nozzle).

Which axes plane?

If your sim models a three-degree-of-freedom trajectory (X,Y, and pitch rotation) then the winds are input into the plane of the trajectory (X, Y).
However, if your sim is a six-degree-of-freedom trajectory then Ref.2 says that you would apply the varying winds (wind profiles, gusts, shears) at 90 degrees to the plane of the trajectory (=yaw plane) as this gives the worst loads, and just apply a constant wind in the plane of the trajectory (pitch plane) to bend the trajectory as the vehicle leaves its launchtower. However, for our nearly vertical trajectories either plane will do.

The response from the vehicle

Having fed gusts, winds, and windshears into the sim, the sim will produce a vehicle response. The data that we need output from the sim are: the dynamic pressure q, the normal and axial accelerations caused on the vehicle, the pitching angular acceleration *R* around the vehicle CG, and the maximum angles of attack on the fins and nosecone and the lifts that these cause. The sim described in our paper 'A dynamic rocket simulator' will give you this data.
You then apply these variables to the airframe to see what loads they create as we shall find out in the next section:

Part 2: Airframe loads

Both thrust and aerodynamic loads, which are treated mathematically as forces externally applied to the vehicle, vary in magnitude and direction, and are resisted only by the vehicle's inertia.
This 'inertial resistance' causes 'inertial loads' within the structure as it accelerates in response to the external forces.
The thrust is taken to act at the nozzle exit.

First approximation

As a first approximation, the French military suggest that the worst aerodynamic loads on the vehicle in windy conditions are of the same order of magnitude as the launch weight of the vehicle. This is borne out by sims of level 2 HPR (High Power Rocketry) class vehicles. However, a gust and windshear occurring at max q will give larger loads so don't fly in that kind of weather!

Aerodynamic loadings:

The distribution of the aerodynamic forces along the fuselage (in directions normal and axial to the long axis of the vehicle) determines how the vehicle suffers when gusts or windshears cause sudden angles of attack.

A typical such subsonic aerodynamic Lift force distribution looks something like this:

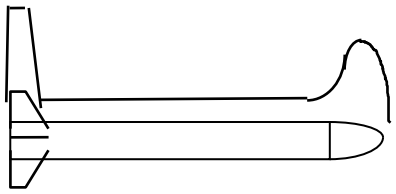

Note there is negligible 'carry-over' of pressure from the nosecone rearwards to the rest of the fuselage. In contrast, for supersonic airspeeds there is significant carry-over.

The fin Lift is concentrated over a very small axial distance, so items such as fins, shrouds, boat-tails, or large protruberances have local normal (lateral) forces that are most easily modelled as discrete loads acting at their respective centres of pressure.
Strictly, the nosecone is often too long to be dealt with as a concentrated load, but we shall do so, as nosecones are usually rigid structures: they will deform little.

Mass (inertial) loadings:

D'Alembert's principle (a fore-runner/pre-statement of Newton's 3rd law) describes mass as that which resists acceleration by causing an 'inertial load' (Newton called this 'Inertia': ma = F)
For a rocket vehicle, any unbalanced global external loads (thrust greater than drag, lift forces caused by a sudden gust-induced angle of attack) produce a global vehicle acceleration:

$$a_{global} = \frac{\sum F}{\sum m}$$ (Newton's 3rd law)

This acceleration causes local 'inertial loadings' caused by local masses (such as, say, equipment or payload masses) acting on their supports within the vehicle structure, which are straight-forwardly determined using D'Alembert's principle: $F = m_{local}\, a_{global}$

Don't forget that the global vehicle mass changes with time due to fuel mass being ejected out the nozzle during the burn.

To analyse the effect that the inertial loadings have on the vehicle structure, the total (global) mass of the vehicle is split into a distribution of discrete (local) component masses at various geometric stations along the length of the vehicle.
Because we are interested in airframe structural performance, the mass distributions we chose to consider for the analysis of normal components of inertial loadings may or may not coincide with distributions considered for axial components.

For example, the mass of a fuelled propellant tank is supported axially by a bulkhead at the bottom of the tank, therefore we can represent the axial inertial load of the tank mass (including propellant) as being concentrated at a single station at the base of the tank, whereas its normal component is supported by the tank wall, spread over the length of the tank.

Provided equipment masses are spread evenly around the vehicle central axis, (please make sure they are!), then structural near-symmetry results in inertial near-symmetry, with correspondingly negligible inertial coupling between the pitch, yaw, and roll planes, which means that motion in these individual planes can be modelled separately from one another.

Axial loads:
To evaluate the effects of combined loadings at some airframe cross-section which is distance 'x' from the base of the vehicle, all axial loads (thrust *T*, drag *D*) are simply added together:

$$A(x) = -T + D_{nosecone} + D_{fuselage} + D_{base} + D_{fin} + S\,\Delta P + a_x \sum_{X0}^{x} m_x$$

The sum is started from x_0 (the aft end). Note that it is important to assign the correct signum to the forces: the thrust is in the opposite direction to the drags so has the opposite sign. Also, forces causing compression in the airframe are signed negative by convention (tensile forces are positive).

The first term above, *T* is the thrust acting at x_0.

The $D_{fuselage}$ term is predominately the skin-friction drag of the fuselage between X_2 and X_1, its length. It can be expressed by integrating the drag per metre, or more usefully, drag per calibre 'C', where a calibre is x/d, where d is (the widest part of) the fuselage diameter:

$$D_{fuselage} = qS \int_{C1}^{C2} \frac{dD}{dC} dC$$

The D_{base} term is the combined base and boat-tail drag, acting pretty-much at the nozzle position.

The fins drag term is D_{fin}, acting at the fins centre of pressure.

$S\,\Delta P$ is the tensile axial force caused by any fuselage compartment internal pressure: ground-level air trapped within the fuselage as the vehicle rises into outside air that has less pressure at altitude. (*S* is fuselage cross-sectional area, and ΔP is the difference in pressure between inside and outside.)

The last term is the inertial load contribution of all masses tailward of the fuselage section '*x*', and is multiplied by the vehicle's axial acceleration 'a_x'.
Continuous mass distributions (such as fuselage walls) can be expressed as a mass per metre or per calibre then integrated over the length ($x-x_0$), as was done for the fuselage drag (above).

When applying the above equation, it's easier to start at the nosetip and work backwards, and then the drags and masses add up to a maximum at the engine thrust bulkhead.

For example, here's a vehicle just before engine burnout. The propellant mass is low, but the thrust is high. The other main inertial loads are the avionics and their batteries, and the main 'chute mass. Note the change in the graph sign caused by the thrust acting at the engine bulkhead; the structure tailwards of this bulkhead is in tension while the structure nosewards of this bulkhead is in compression. The graph slope is caused by the fuselage drag term.

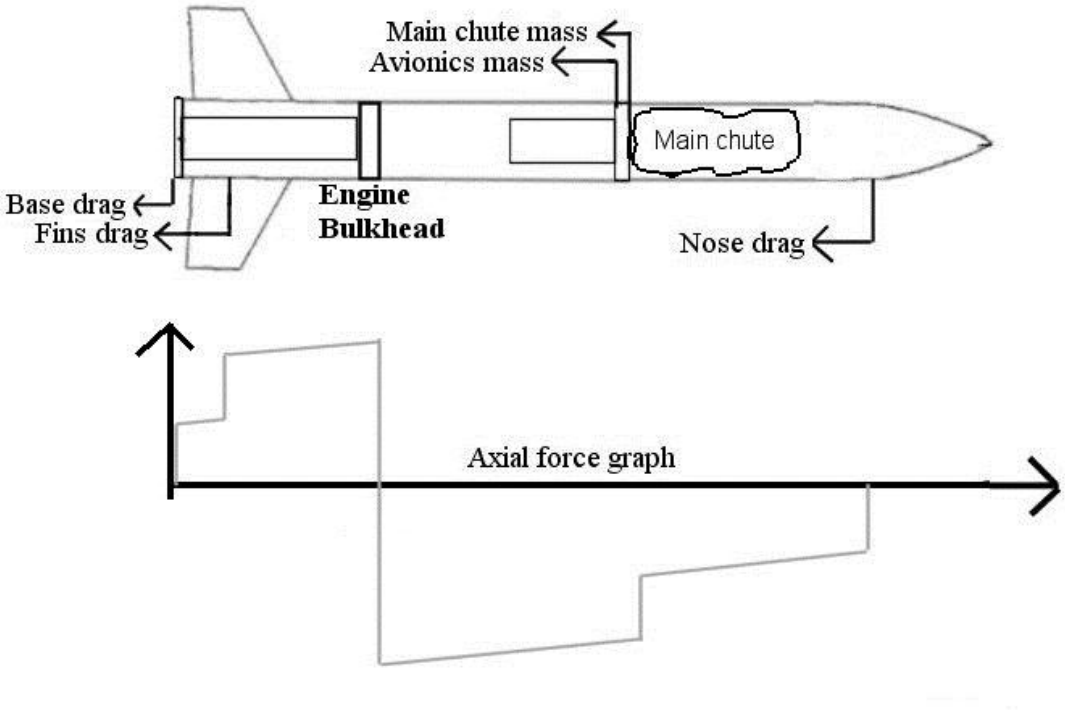

You can see that the largest loads suffered by the fuselage occur at the location of the engine bulkhead.
Note how individual loads are suddenly switched on only when the graph reaches their axial position; they 'don't exist' before their position is reached. This is the theory of the Method of Sections which is explained below.

The maximum aerodynamic axial load (max drag) will occur at max **q** (see glossary) whereas the maximum thrust will occur at launch. Maximum inertial loads will occur at maximum longitudinal acceleration: this is typically around engine burnout for a solid motor, but may be much earlier for the diminishing thrust from a hybrid engine.
Monitor your sim to find the point in the trajectory where the peak load on the engine bulkhead occurs.

Normal loads:
The loads normal to the long axis of the vehicle, in contrast, cannot simply be added together like axial load components.

The fuselage structure can be modelled as a cylindrical beam, which is unsupported at each end (known as a 'free-free beam').
Using standard engineering beam theory, the normal loads can be reduced to an equivalent system of Shear Forces and Bending Moments, (see below) and these equivalents *can* be simply added.

Beam theory:
Visualise a rigid 2-dimensional beam (i.e. of arbitrary cross-section) rigidly supported at the left-hand end by a wall, free at the other end, and loaded at some cross-section along its length by a suspended mass 'M':

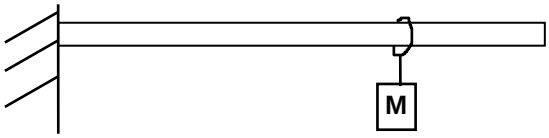

If the weight (Mg) of the mass 'M' is too great, it'll shear the beam apart at the point of its application. The beam is obviously designed to withstand this normal load, and note that the wall is obviously able to supply an equal load in the opposite direction (action and reaction), otherwise the beam would fall:

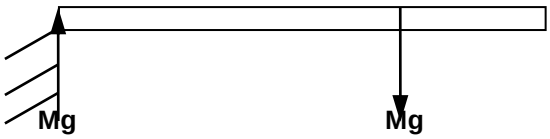

So the beam is in static equilibrium, i.e. there is no acceleration vertically.

The engineering method known as the Method of Sections can be used to analyse any longitudinal station along the beam.
The trick is to pretend that you've sliced the beam at the point of interest, and then you apply appropriate forces in appropriate directions to keep the beam stationary.
From the principle of action and reaction, the 'Shear' force you pretend to apply at the slicing point is equivalent to the shearing force actually occurring within the material of the beam at that station.

For example, if you were to make a pretend vertical cut in the beam at any point between the wall and the load 'M', you would have to apply an upward 'shear' force of magnitude Mg to support 'M', so a graph of actual Shear Force (V) within the beam, with beam distance, looks like this:

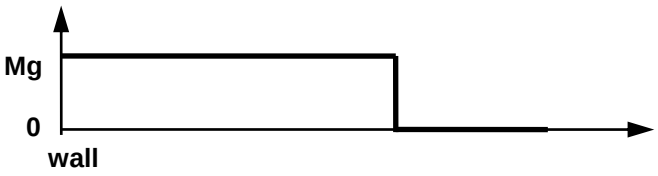

Obviously, if you cut it, the cut portion of the beam would rotate downwards about the cut, so you also have to apply a pretend Moment (B) for rotational equilibrium:

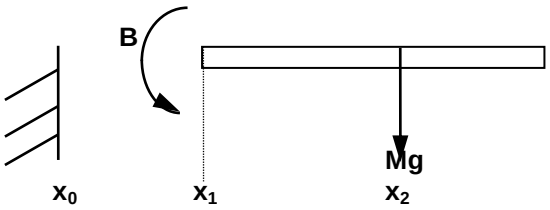

where $B = -Mg(x_2 - x_1)$

Again, from action and reaction this pretend Bending Moment at the pretend cut is equivalent to the actual moment occurring within the beam at the imaginary cut station, caused by the load M.

Note the distance dependence in Bending Moment calculations.
In actual fact, for simple beam theory, the Bending Moment graph is simply the integral of the Shear Force graph with respect to distance along the beam (x) therefore, for the complete beam shown above (beam itself assumed massless), the graph of Bending Moment with distance from the wall is:

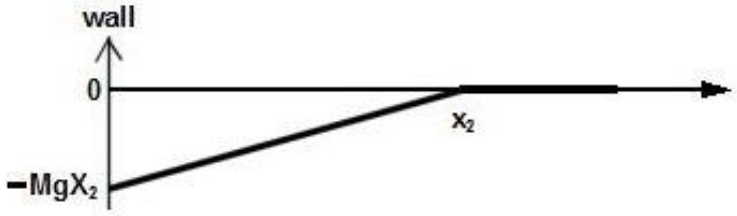

In this diagram I've used the usual sign conventions for Shear Force and Bending Moment (opposite to that for concrete structures) which are:

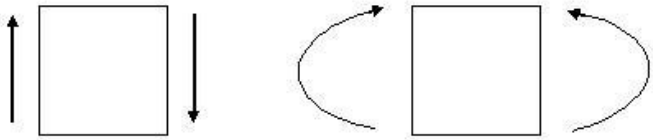

Positive Shear Forces will spin a segment of beam clockwise, and positive Bending Moments will cause a section of beam to sag downwards. There are lots of examples of Shear Force and Bending Moment diagrams on the web, usually for static beams.

In the real world, no beam is infinitely rigid, so the bending moment stretches the top surface of the beam, and compresses the bottom surface of the beam, so it's important to know the Bending Moment maximum values, and where along the beam they occur, because too great a stretch will tear/fracture the beam, and too great a compression will crumple the beam.
In the above case, the beam would snap where it meets the wall if too great a load was hung off it, because that's where the maximum Bending Moment occurs.

Rocket 'beam' theory:
In the above example, the beam was restrained from rigid-body movement and rotation by the wall, (though a little elastic deflection might occur) so that it was in static equilibrium.
In contrast, for a rocket vehicle fuselage, the only *partial* restraint against lateral aerodynamic loadings (and vectored thrust) is the rocket's inertia ('inertial loads', see above).
The Bending Moments created by the Shear Forces (caused by the aerodynamic loadings) are only *partially* restrained by the fuselage's rotational moment of inertia, and also by the aerodynamic moment caused by the fins, so the Bending Moments cause the vehicle to rotate around its centre of mass (C.M. or C of G) with some value of rotational acceleration. (Remember basic physics: unrestrained 'rigid' vehicles always rotate about their centre of mass.)
Also, the lateral loads cause the vehicle to linearly accelerate laterally.
So the vehicle is not in *static* equilibrium, but we can still use the above beam theory because it is in *dynamic* equilibrium: the vehicle's lateral acceleration a_y is the sum of the lateral loads on the vehicle divided by its total mass (Newton's second law). This acceleration causes 'inertial loads' within the fuselage which load it.

Lateral Shear Force:
This is caused by the externally-applied aerodynamic loads, which themselves are caused by a gust-induced angle of attack. As the vehicle accelerates laterally and also rotates, further Shear Forces are caused by 'inertial loads' from equipment masses.

The equation for the Shear Force occurring at (@) some longitudinal station 'X_k' is

$$V_s(@x_k) = L_{nose} + L_{boat} + L_{fin} - a_y \sum_0^{x_k} m - R \sum_0^{x_k} m\,(x_{CG} - x_m)$$

The sum is started at station 0 (the aft end).

Continuing the philosophy of the Method of Sections, the nosecone lift $L_{nosecone}$ term is only included if 'x' is nearer the nosetip than the nosecone's centre of pressure, and so on.
Note that I've assumed that the lift is concentrated at the nosecone's centre of pressure because the nosecone is rigid. In reality, it's spread over the length of the nosecone, giving:

$L_{nose} = qS\alpha \int_{C1}^{Ck} \frac{dL}{dC}\, dC$ where *C* is calibers (x/d) and α is angle of attack.

The boat-tail lift term $L_{boat\text{-}tail}$ is the boat-tail lift force. Note that boat-tail 'lift' occurs in the opposite direction to fin lift, so is negative.

The fin lift term is L_{fin} acting at the fins centre-of-pressure.

The a_y summation term is the inertial load contribution of all masses '*m*' tailward of the fuselage section 'X_K', and is multiplied by the vehicle's lateral acceleration a_y (metres/sec). Each component mass '*m*' gets 'switched on' as X_K passes it, and acts at its own position 'x_m'.

The last term in the Shear Force equation above represents rotational inertia loads caused by the rotational acceleration '*R*' (radians/sec) about the vehicle's centre of gravity 'x_{cg}'. For example, if the nose rotates such that the nose accelerates in the same direction as the nose lift, then the inertial Shear Force on the nose's mass is increased.
The 'x_m' in this term is the position of each individual mass.
Take care with your defined direction of positive *R*.

Continuous mass distributions (such as fuselage walls or main parachute) can be expressed as a mass per metre then integrated over their length, in which case the last two terms are:

$$-a_y \int_0^{x_k} \frac{dm}{dx}\, dx \; - R \int_0^{x_k} \frac{dm}{dx} (x_{CG} - x_m) dx$$

Bending moment:
The equation for the Bending Moment occurring at (@) some longitudinal station 'X_k' is obtained, as usual for a beam-like structure, by simply integrating the Shear Force equation above with respect to 'x':

$$M_B(@x_k) = L_{nose}(x_k - x_{nose}) + L_{boat}(x_k - x_{boat}) + L_{fin}\left(x_k - x_{fin}\right)$$

$$-a_y \textstyle\sum_0^{x_k} m(x_k - x_m) - R \sum_0^{x_k} m(x_k - x_m)\,(x_{CG} - x_m)$$

Note that the term for the fins starts at the fins position X_{fin} for example.

The translational and rotational accelerations are identical to those in the Shear Force equation above.

The 'x_m' in the last two terms is the position of each individual mass '*m*'.

Continuous mass distributions (such as fuselage walls or main parachute) can be expressed as a mass per metre then integrated over their length, in which case the last two terms are:

$$-a_y \int_0^{x_k} \frac{dm}{dx}(x_k - x_m)dx \; - R \int_0^{x_k} \frac{dm}{dx}(x_{CG} - x_m)(x_k - x_m)dx$$

Example:

As an illustration of how the above equations work, imagine a vehicle where the fins lift and nose lift are equal, and the CG is halfway between nose and fins centres of pressure. This will mean that the vehicle doesn't rotate as it accelerates laterally. Also assume, for the purposes of illustration, that the mass of the fuselage tube is zero.

We'll examine two loading cases: both are loaded with the same mass 'M', and in both cases, the inertial force 'F' resulting from the vehicle's upwards acceleration acting on 'M' is equal to the nose plus fins lift.

On the left, the mass is split into two halves which are positioned at the fins and nose centres of pressure. On the right, the mass is concentrated halfway between the nose and fins centres of pressure.

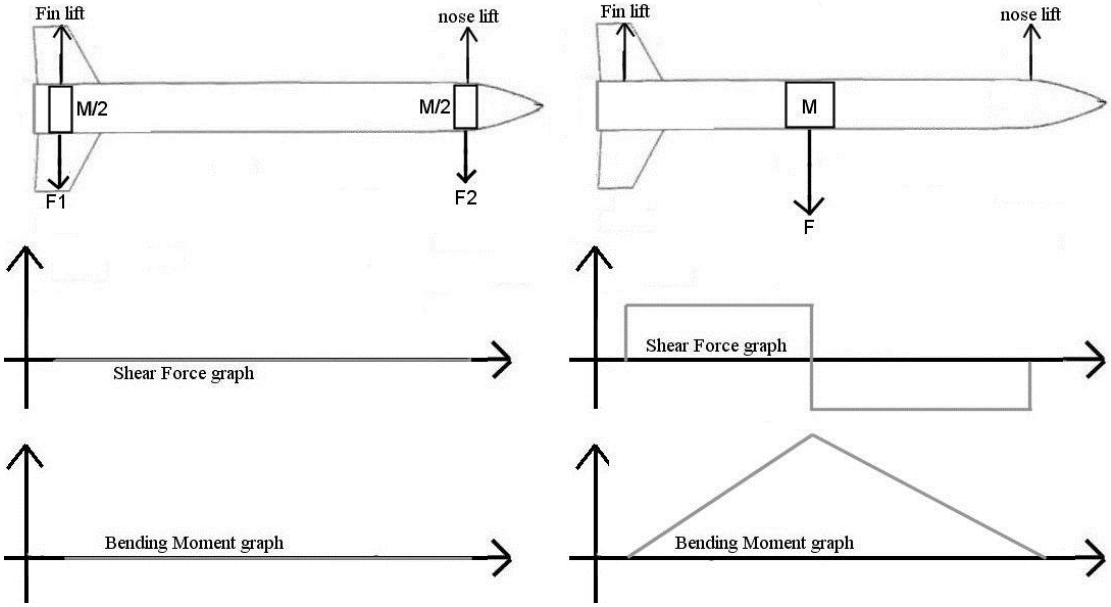

As you can see, on the left we have the minimum (zero) bending moment. F1 = the fins lift, and F2 = the nose lift (so F1 + F2 = F= nose lift + fins lift). The resulting net Shear Force is zero, so the Bending Moment remains at zero.

On the right, F= nose lift + fins lift again, but this time there are large (maximum) moment arms between mass 'M' and the nose and fin centres of pressure, so you get the maximum possible Bending Moment, occurring halfway along the vehicle.

The Bending Moment graph (derived by integrating the Shear Force graph) shows that it's therefore much better to position your equipment masses at the centres of pressure as on the left diagram: this bends the fuselage much less. The mass of our engines tends to be rearward anyway, but we can reduce bending by positioning our recovery systems and their batteries inside the nosecone rather than some distance behind it.

Note from the right-hand Shear Force graph that when analysing the vehicle using the Method of Sections, each load is only 'switched on' once the section of interest has been passed, and the Shear Force remains at that value until the next load is reached.

Similarly, from the Method of Sections, note how the gradient of the Bending Moment graph only changes once the Shear Force has been passed, i.e. in the above Bending Moment equation, you don't add the nose lift until 'X_k' is greater than 'Xnose'; if 'X_k' is less than 'Xnose' it's as if the nose lift doesn't yet exist.

Note how the Bending Moment graph starts at zero and ends at zero. That's because the ends of the vehicle are unrestrained.
Also note that a positive Shear Force causes a positive Bending Moment curve gradient, and a negative Shear Force causes a negative Bending Moment curve gradient.

For the above picture, the Bending Moment is positive, it would cause the vehicle to banana into a shape that sagged in the middle. (In structural engineering, negative Bending Moment causes the opposite of sagging which is known as 'hogging'; this may be a nautical term).

A more typical rocket vehicle loading looks like this:

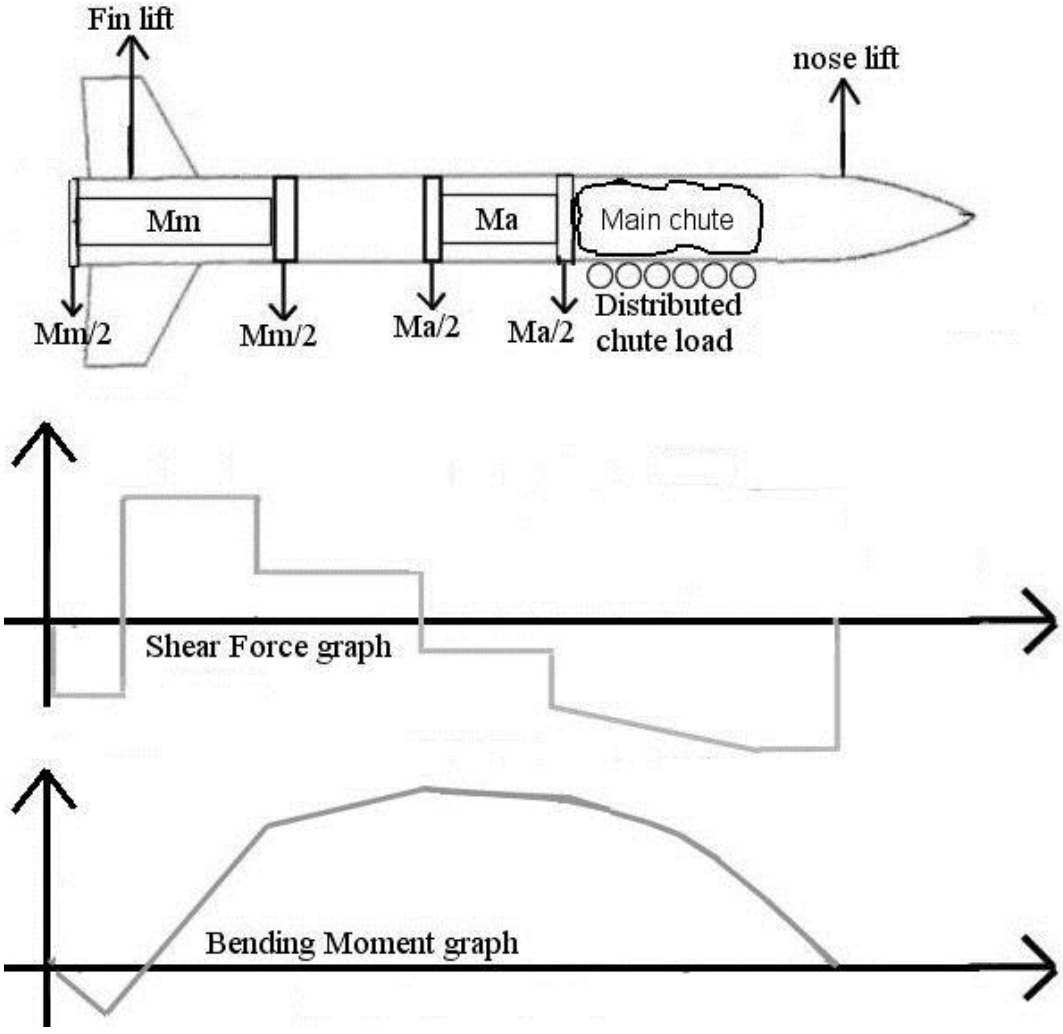

In this diagram, the vehicle is accelerating upwards due to the nose and fin lift, but it would also be rotating about its CG.

The resulting inertial loads from motor mass' Mm' and avionics mass 'Ma' are supported equally by fore and aft bulkheads, and the main chute gives a distributed load along the fuselage wall.
The point loads cause a sudden jump in Shear Force, whereas the distributed inertial loading of the main chute is a continuous downwards ramp, which causes the gradient of the Bending Moment graph to continuously decrease (a parabolic curve).
Again, the maximum Bending Moments occur about halfway along the vehicle here, where the Shear Force graph re-crosses the X-axis.

I haven't added a boat-tail to this vehicle, but it is easily added, acting at its own centre of pressure. Remember that boat-tail lift is in the opposite direction to nose and fin lift.

Worst loads:
Maximum Shear Forces (and therefore worst Bending Moments) will occur at the instant of maximum aerodynamic load, which is caused by a maximum of: dynamic pressure q times angle of attack α times lift coefficient C_L. Note that lift coefficient C_L varies with Mach number and has a peak value around Mach 1.

It so happens that with large launch vehicles, max q, Mach 1, highest sidewinds and worst windshears all tend to occur almost at the same time and altitude, which makes picking a worst-case point in the trajectory to design for easy: Black Arrow was designed for max q. With these vehicles, the engines were throttled-back around max q giving little longitudinal acceleration so their airspeed remained nearly constant over a moderately large vertical region. Consequently, their Mach number was more or less constant over this region too, so that lift coefficient C_L varied little over a fair vertical height.

For our smaller vehicles with their higher accelerations, these maximums probably won't line-up so easily, and we don't throttle-back at max q.
So what point in the trajectory would be the worst? It would be the product of dynamic pressure times *potential* angle of attack, i.e. if a gust or windshear caused an angle of attack at this point, it would cause the worst load.
So using your sim, monitor the product of dynamic pressure times **lift-curve slope** $\frac{dC_L}{d\alpha}$ at zero angle of attack and find the maximum:

Maximum potential worst load occurs at maximum of the product: $q \left|\frac{dC_L}{d\alpha}\right|_{@zero\ \alpha}$

where $\frac{dC_L}{d\alpha}$ is the sum of the lift curve slopes of the nose, fin, and boat-tail, which all vary similarly with Mach number. Although the boat-tail lift is negative and the lift of the fins and nose cause rotations in opposite directions, it makes sense to take the sum of the *magnitudes* of nose, fin, and boat-tail added together as an indicator of worst aerodynamic load. (For symmetric vehicles flying at zero mean angle of attack the nose, fins, and boat-tail all reach a peak value simultaneously: at around Mach 1.)

I've chosen the lift-curves in the limit of zero angle of attack, because you've got to pick some angle of attack to analyse. Finned vehicles fly mostly at zero angle of attack and pull at most only one or two degrees angle of attack in response to a gust at max q, so the value at zero angle of attack is practically the same.
The lift-curve slopes come out of the sim as they're needed for stability analysis (see our paper 'A dynamic stability rocket simulator').

Having determined the worst potential point in the trajectory you would then add a gust and/or windshear to your trajectory sim at that worst point to cause a peak angle of attack there.

What we are really interested in is that fact that the peak aerodynamic load will give the maximum lateral and rotational accelerations (maximum vehicle response), giving the highest inertial loads.
The response of the vehicle to a sudden gust or windshear is as follows, as shown here by subsequent snapshots of time in the trajectory:

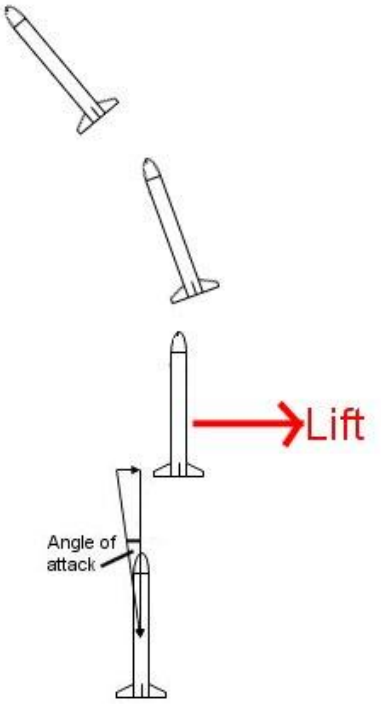

The lift of the vehicle caused by the sudden switch-on of angle of attack can be assumed to be instantaneous, certainly much faster than the period of the (rigid-body) response of the vehicle (note this assumption may fall over for tiny model rockets which react very quickly).

The vehicle accelerates in the direction of the total lift which builds up a lateral airspeed component that reduces the angle of attack. Simultaneously the fins rotate the vehicle to reduce the angle of attack further in a highly damped oscillatory motion.

So the moment of peak aerodynamic load occurs the instant the gust hits, before the vehicle response reduces the angle of attack and so reduces the aerodynamic load. Monitor this instant with the sim as this instant therefore gives the peak lateral acceleration and rotational acceleration:

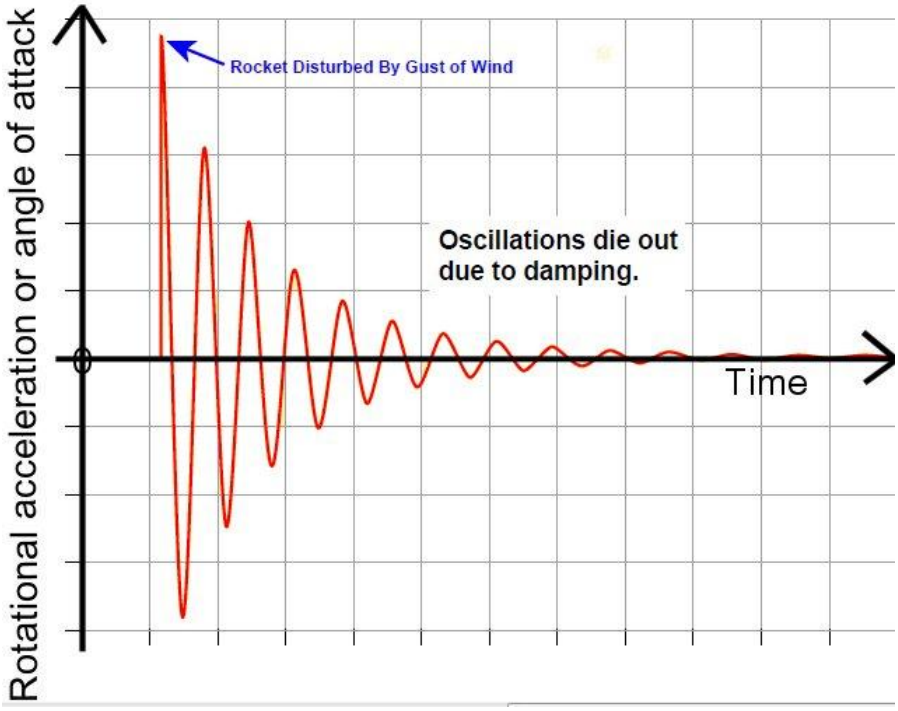

(Note that the angle of attack response is actually 180 degrees out of phase with the rotational acceleration so this graph shows the *negative* of the rotational acceleration.)

Dynamic loading effects:
In the above analyses, the assumption is that the vehicle behaves as a rigid body. The loadings thus calculated are known as quasi-static.

As well as the rigid-body motions, the vehicle fuselage also bends elastically out of shape slightly in response to the various loadings as if it were rubber.

What is the effect of this flexibility in terms of the loadings created? This analysis is called a dynamic loading analysis.

In the above Shear Force and Bending Moment analyses, the aerodynamic forces acted in one direction, which then instantaneously caused the inertial loads to occur in the opposite direction.

A dynamic analysis assumes the fuselage is flexible. As the fuselage also has mass, there is a brief time-lag while the fuselage mass responds to the aerodynamic load. The fuselage then 'twangs' like a ruler, with a particular 'twang' shape (known as a mode), and a vibration at a particular frequency.

Natural bending modes

the forces and masses along the fuselage cause the elastic bending to acquire a definite, predictable, shape or ' natural mode'. The shape is time-dependant: the fuselage will bend in and out of shape at a set frequency like a vibrating ruler.

And it's not just one mode, it's actually a combination of several independent natural modes: each mode vibrates at a certain (successively higher) frequency which may or may not be set in motion by gusts and turbulence of a similar frequency.

These bending modes have particular shapes which roughly look like this (note that the fuselage is free to move at both ends, i.e. the ends are unrestrained):
The solid line is the central axis of the fuselage, and the graduated x-axis is the unbent shape. Each mode is a half-wavelength higher than the previous one.

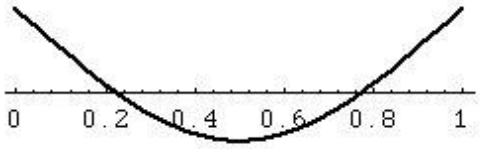

First bending mode of a beam with free boundary conditions

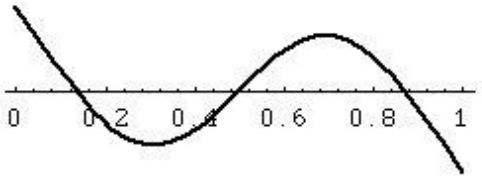

Second bending mode of a beam with free boundary conditions

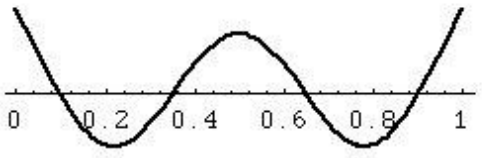

Third bending mode of a beam with free boundary conditions

And so on. It's worth noting that successively higher bending modes have less and less effect (smaller amplitude), so can be neglected: if analysing a full-size aircraft or large launch vehicle, attention would be paid to only the first three or four bending modes. For our small vehicles, we'd only be interested in the first mode (the fundamental mode).
What the above static pictures can't show is that these standing waves oscillate: half a cycle later, each picture gets turned upside-down.

The gust doublet
In aircraft design, it's often the practice to assume the worst-case of a second gust that occurs immediately after the first, of equal worst magnitude but in the opposite direction. And it's further assumed that the timing of the second gust after the first occurs at the same frequency as the first bending mode, so causes this mode to resonate.

In a nutshell, this has the effect on the vehicle that for a brief instant, the inertial lateral loads caused by the first gust (as calculated quasi-statically) are now acting in the *same direction* as the new aerodynamic load caused by the second gust, until the vehicle has had time to flexibly respond.
Having the aerodynamic loads and inertial loads occurring both in the same direction means they are simply added together to get a worst-case load.

Note that the probability of a maximum strength gust doublet occurring is much less than a single worst case gust occurring.

Bending effects
The bending is important to model, because it will cause the ends of the airframe (nose and tail) to pull different, possibly larger, angles of attack. And as the fuselage vibrates, this increases the inertial loads. Also, a bent tail will cause the line of the thrust force to miss the centre of gravity, so a rotation is caused by the thrust, as well as a sideways thrust component. Reference 2 says that the effects of bending on the pitch/yaw moments of inertia are minor and can be ignored, as can the bending effect on the CG position.

Needless to say, the bending modes cause merry hell with any inertial sensors (lateral accelerometers or rate gyros) within the fuselage. Fortunately, the bending mode frequencies are fixed and predictable (and measurable with a fuselage-shaker) so these frequencies can be filtered-out of the sensor data.

Mode analysis
Basically, the bending modes cause additional Bending Moments; terms which have to be added to the Bending Moment equation given earlier.
Structural analysis or testing will determine the amount of Bending Moment occurring for each individual mode deflection η_i, These moments can then simply be summed:

$$M_B(x_k,t) = \sum M'_{\eta i}\,\eta_i$$

The *M'* terms simply represent scalar coefficients which, when multiplied by the appropriate deflection, give the Bending Moment caused by that deflection.

Historically, matrix methods and an energy analysis would have been used to calculate the bending mode shapes (Ref. 2), but nowadays a Finite Element Analysis on computer would be used, where the mass of each element is included: the distributed mass and inertia along the fuselage must be lumped into discrete point mass elements.

It has been found from trial-and-error (Ref. 2) that for simple one-dimensional (length along the fuselage only) beam bending models the required number of fuselage mass stations should be approximately ten times the number corresponding to the highest elastic bending mode to be calculated.

For example, if three elastic bending modes are to be calculated, then approximately thirty mass stations are required to represent adequately the bending dynamics of the third mode. More mass stations than this won't increase the accuracy noticeably.

For our simple rocket vehicle shapes (no strap-on boosters) this could actually be quite a simple model which could be run on a home PC, although modelling our composite body-tubes presents additional difficulties because they're anisotropic (stronger in certain directions and weaker in others).

Applicability

But I'm not sure how all this would apply to the small, rigid, composite airframes of HPR class vehicles. They're deliberately exceedingly stiff, so may well not bend to any degree worth bothering about, even in the first bending mode, unless they are particularly long and thin.

Also, the frequency of even the first bending mode might be very high indeed, so only gust doublets occurring at high flight Mach number are likely to be closely spaced enough to set off the first frequency.

And the natural frequencies of solid rocket motors are much higher than the first few bending mode frequencies, so don't set them off.

And HPR airframes have coupler joints at the parachute compartments. These joints have a degree of rotational free play that is difficult to model mathematically. Then there's the mass of the parachute, which is free to move laterally; again difficult to model.

Frankly, if your HPR-sized vehicle is particularly bendy, you can attach a mechanical shaker to it to investigate its bending modes: suspend the vehicle in a cradle of elastic bands or bungees at its CG and alter the shaking frequency until it resonates with the first mode (you can see this, so photograph or video the bent shape against a graduated background).

Propellant slosh loads

We can ignore the inertial loads from the mass of waves sloshing across the nitrous oxide tank in our HPR hybrids: HPR vehicles are narrow enough that the slosh mass is a tiny percentage of the vehicle mass so can be ignored.

Part 3: Airframe structural design

With most of our HPR vehicles, the stiffened composite skin of the fuselage or 'body tube' serves as:

- The main aerodynamic surface,
- The main structural element,
- The walls of the various compartments (which if pressurised improve structural performance).

Because fuselage radii are usually much larger than skin thicknesses, radial stresses within the skin are unimportant.
Radial loadings associated with geometric and applied loading discontinuities will be ignored here, so at zero angle of attack, the only remaining forces acting on the fuselage skin are tangential, tensile, unit loadings caused by compartment internal pressures : $n_\theta = \Delta \mathrm{Pr}$, (where *r* is the fuselage radius) and also the axial unit loading resulting from the longitudinal (drag) loads:

$n_A = \dfrac{D}{2\pi r}$ where the axial drag load '*D*' is assumed negative in compression.

For thin skins and/or slender fuselages, the skin will buckle and the fuselage collapse under typical axial unit loadings n_A, therefore the skin can be stiffened internally by attaching longitudinal members (stringers) and discretely spaced transverse circular rings (frames). Or, one can use a sandwich construction as John Coker (see his websites) does.

The normal Bending Moment (see previously) resulting from aerodynamic/wind disturbances imposes an axial, varying unit loading:

$n_b = \dfrac{M_B}{\pi r^2}\cos\theta$ around the airframe in accordance with elementary beam theory.

The angular coordinate θ is measured around the longitudinal axis of symmetry of the vehicle, and is assumed zero (therefore giving maximum unit loading) when parallel to the loading plane (the angle of attack) producing the bending.
Because n_b and n_A are coincident and parallel, they add algebraically, and the airframe is in tension when $n_b + n_A > 0$, but, except for some joints, the controlling mode of structural failure under combined axial and bending loads is axial compressive unit loading at θ = 180 degrees or 0 degrees.

The other force associated with aerodynamic/wind disturbances is Normal Shear Force (see previously) which introduces a uniformly varying load around the airframe cross section of:

$n_v = \dfrac{V_R}{\pi r}\sin\theta$

The skin resists this shear unit loading.

Yet another loading is the wind pressure, or aerodynamic pressure acting normal to the skin. If this external pressure is greater than the compartment internal pressure, then all the strengthening benefits of a pressurised skin are nullified, and the stringers simply behave as beam columns, which must be stabilised against collapse by increasing their thickness.

Safety factors

Safety factors are there to deal with the inevitable uncertainties and simplifications in the loads analysis.

Rocketry safety factors are usually a lot lower than traditional civil engineering safety factors to reflect the fact that every extra gram of rocket structure counts. It's up to you to decide how much of a penalty extra structural mass is allowable on your particular vehicle versus the risk of structural failure.

With the quasi-static Shear Force and Bending Moment analysis, a suitable safety-factor would be to multiply the resultant loads by 1.5

For dynamic loading effects such as the gust doublet criteria, there is less probability of two worst-case gusts occurring one after the other at just the worst frequency to cause first-mode resonance so one would use a lower safety-factor of say 1.25

Then check the loads for both cases (multiplied by their respective safety-factors) and see which gives the worst case.

Testing

Exactly how much Shear Force or Bending Moment some location along the vehicle can withstand depends on the material used, and its shape/thickness etc.
Having calculated the largest Shear Force and Bending Moments, we need to replicate these forces and moments when the vehicle is stationary on the workbench, rather than waiting until it breaks-up in the air!
In particular, we need to load all the airframe coupler joints.

In theory, one could apply all of the aerodynamic loads at their correct point of application on the workbench, perhaps by suspending the vehicle horizontally by its nose and fins.
Then replicate the inertial loads using weights within the airframe.
However, this would be tedious; you only need to apply the worst Shear Forces and Bending Moments at the points that they occur, using weights and levers.

Fin stiffness

If the fins are not stiff enough, they can flutter and tear off, therefore fins should be designed for stiffness rather than strength. This can be achieved by giving the fin a little depth so that it is hollow with composite or sheet metal skins. Or use a sandwich construction.
Ref. 8 says that the chances of the fin fluttering decreases with the cube of the fin thickness-to-chord ratio, and decreases linearly with the fin **shear modulus** (effective stiffness). It also decreases with the square of decreasing fin **aspect ratio** which is why rocketry fins have such low aspect ratios. Leading-edge sweepback also helps.

Component requirements

The following were taken from our NRC competition build requirements (Ref.6), and serve as simple guidelines. Due to the crudity of the loading estimates, safety factors of two would be typically applied:

Motor mount strength requirements

The motor mount must be designed to take all of the thrust loads, both axial and lateral via the motor thrust ring. It must be designed to withstand and transmit into the body tube a force equivalent to twice the maximum motor thrust without permanent deformation. The mount must also withstand a lateral force in any direction equal to a thrust misalignment of 5° at the maximum thrust value cantilevered about the vehicle's centre of gravity.

Fin longitudinal loading

Each fin must be able to support a suspended load from its tip equal to twice the fin mass times the rocket's maximum axial acceleration occurring during any flight phase.

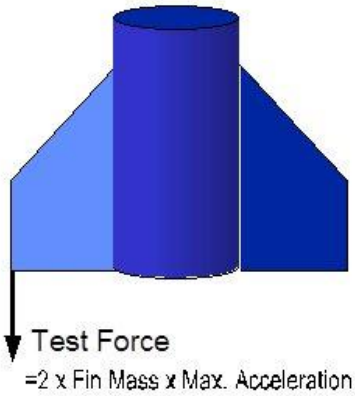

As a stiffness test, each fin must withstand a transverse load equal to the rocket's launch weight when suspended from the fin tip. When subjected to this load, the maximum lateral deflection measured at the tip must be less than 10° in either direction:

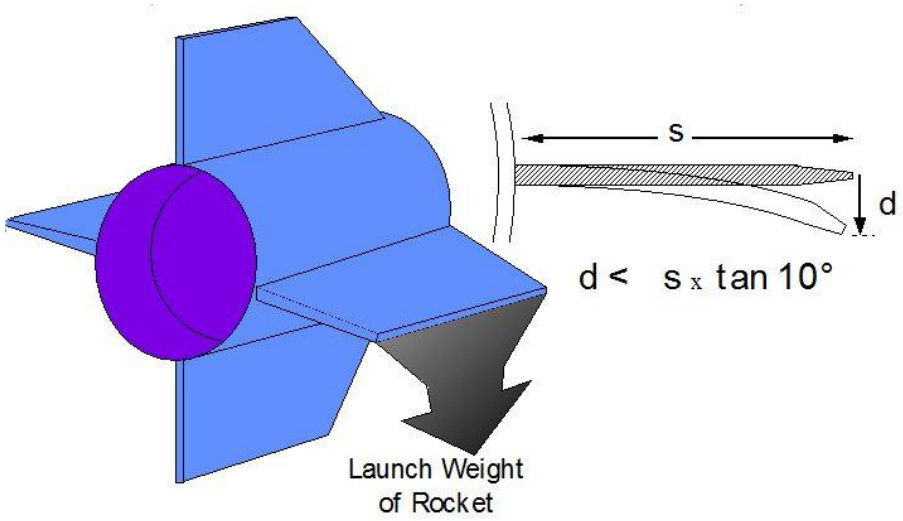

Fin allignment accuracy

According to Rocket Services UK, accuracy of construction, such as fin allignments, is absolutely vital for survival of the Transonic zone (around Mach 1). Misalligned fins (more than 2 degrees off the long axis of the vehicle), as well as creating more drag, can create vehicle angles of attack (and spinning) that could cause excessive Normal forces, breaking the vehicle.

It is good practice to investigate the effect of small misallignments in the trajectory sims to see if these use up any loading safety-margins. (Try misalligned nosecones and fins).

References

Ref. 1: Sampo Niskanen, “Development of an Open Source model rocket simulation software” Helsinki University of Technology Faculty of Information and Natural Sciences Master's thesis, Espoo, 20.5.2009

Ref. 2: Geissler, E. D, AGARDograph 115 “Wind effects on launch vehicles” advisory group for aerospace research and development (AGARD) Neuilly-Sur-Seine, France, Feb 1970 Downloadable at:
http://ftp.rta.nato.int/public//PubFullText/AGARD/AG/AGARD-AG-115///AGARD0G115.pdf

Ref. 3: AD 408259 “Wind shear response for missile systems, comparative study and design procedure” Lockheed missiles and space company, 1963 (downloadable)

Ref. 4: James R. Scoggins “Wind Design Criteria for the Saturn Vehicle” NASA Marshall space flight center, Aeroballistics division, 1963 (downloadable)

Ref.5: NASA SP-8035 “Wind loads during ascent” 1970 (downloadable)

Ref.6: Aspirespace NRC technical requirements 2000 (based on Les Jeunes en Espace strength criteria from their Mourmelon launching campaigns.)

Ref.7: Private communication with our good friend Cliff Maidment, aerodynamicist to the Black Arrow rocket programme.

Ref.8: Dennis J. Martin, “Summary of flutter experiences as a guide to the preliminary design of lifting surfaces on missiles”, NACA TN 4197, Feb 1958

Glossary:

Geometric definitions:

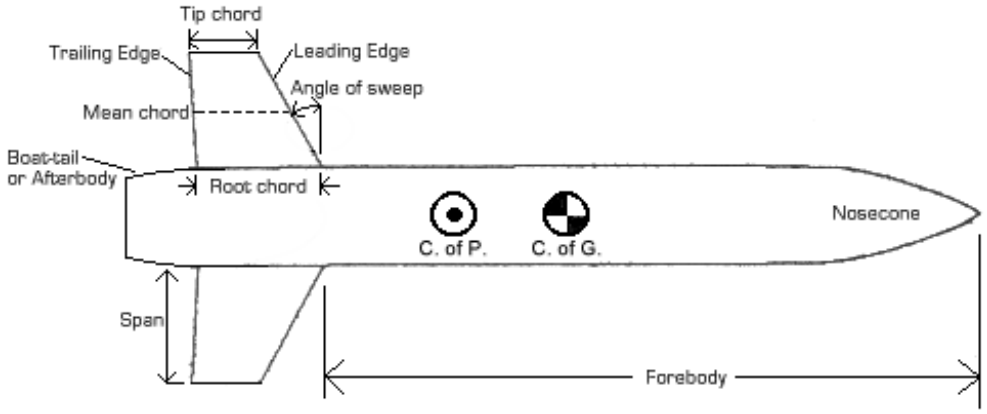

Angle of attack: α (or Angle of Incidence)
This is usually referred to as 'alpha', and corresponds to the angle between the airflow direction (usually the Freestream direction) and some vehicle or fin datum.

Aspect ratio AR:
A wing or fin's span divided by its mean chord (see above diagram).

Calibers, Calibres:
In rocketry, vehicle dimensions are usually divided by (compared to) the diameter of the thickest part of the fuselage so that rockets of different size can be compared: this diameter is therefore one Caliber.

Centre of Gravity, centre of mass (CG):
The point within the vehicle that is the centroid of mass, the balance point.

Centre of Pressure (CP):
The point on the vehicle's surface where the average of all the aerodynamic pressure forces from some component (e.g. nosecone) act.

Dynamic pressure: (q)
All aerodynamic forces scale directly with the kinetic energy term $\frac{1}{2}\rho V^2$
ρ being volume-specific mass = air density, and V = flow velocity (airspeed).
This kinetic energy term is called Dynamic Pressure (q), to distinguish it from its Potential energy counterpart of static pressure (P).

Lift-curve slope: $\left(\frac{dC_L}{d\alpha}\right)$
The gradient of the lift versus angle of attack graph.

Reference area: (S)
See Drag (equation)

Shear modulus: (G)
In stress analyses, the shear modulus (or modulus of rigidity), is defined as the ratio of shear stress (applied load) to shear strain (resulting twist), or in other words the higher the shear modulus, the less the fin material will twist under an applied torsional load.

Appendix 1: constructing a synthetic wind profile

Start with data of windshear versus altitude: here's the graph for Cape Canaveral (metres per second per Km):

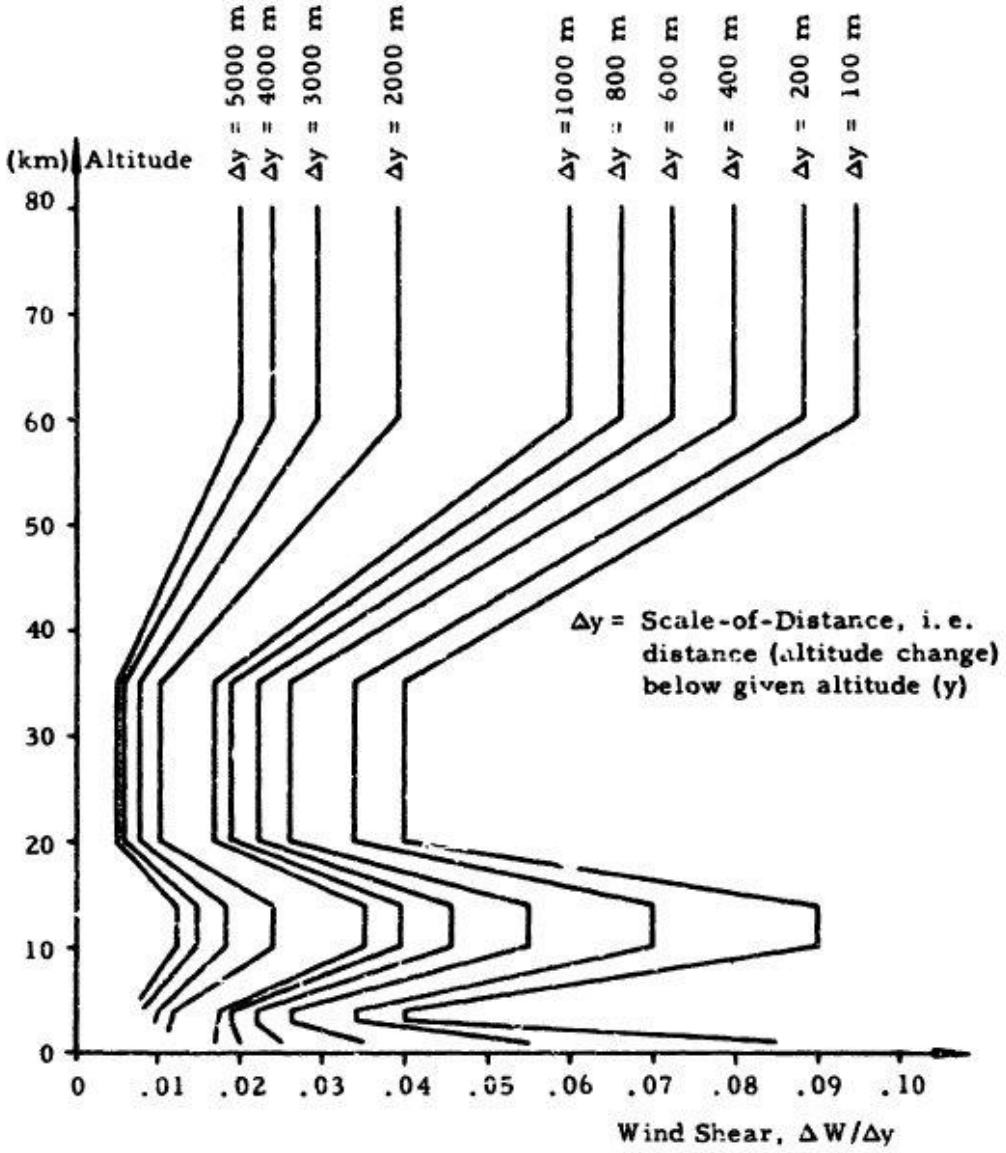

Note that there are several curves for different scales of distance (Δy). For example, the 100m curve was obtained from an instrument (radar) that could discern wind differences at intervals of 100m altitude. Note that due to fine wind structures the wind-shear values increase the closer you examine them (smaller scales of distance) which is why the curves aren't equal. If you were using your own windshear data you might only be able to provide one or two different scales of distance.

Note that the curves are composed of straight line segments, this is deliberate to allow one to linearly interpolate between altitudes.

Next, convert this windshear data into a wind speed change spectrum, corresponding to a change in altitude (scale of distance), and plotted as a function of altitude. Wind speed change is found by multiplying the above windshear data by the above scales of distance:

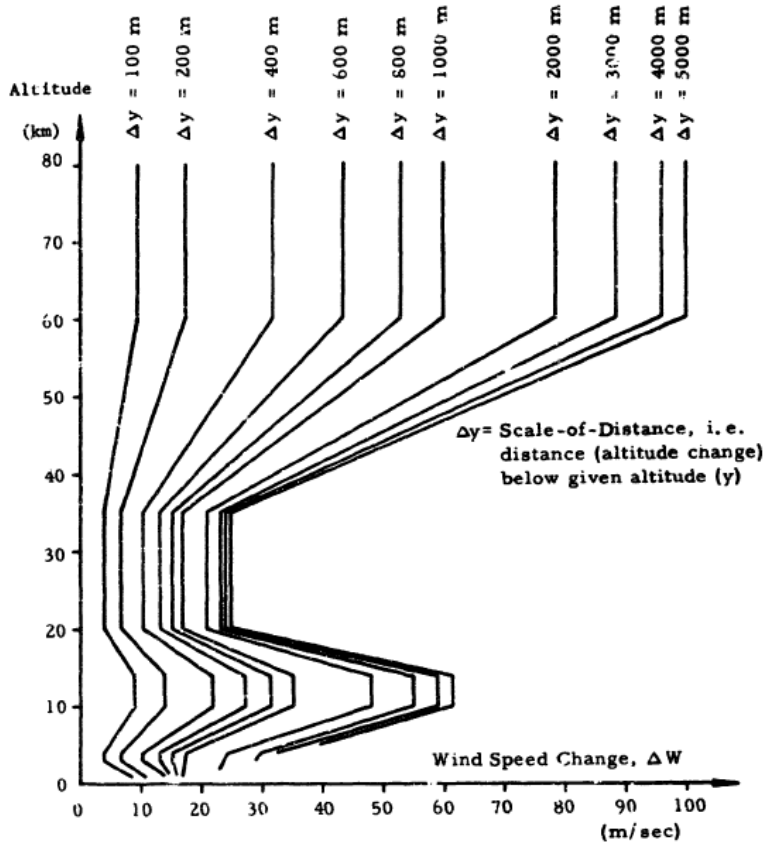

In tabular form, the above graph is:

Altitude Km	Windspeed (metres/sec) per scale of distance (metres)									
	100m	200m	400m	600m	800m	1000m	2000m	3000m	4000m	5000m
0.1	10.7									
0.2		12.4								
0.4			15.5							
0.6				15.3						
0.8					16.1					
1						16.9				
2							22.9			
3	3.9	6.8	10.3	13.2	15.1	17.2	23.5	28.8		
4	3.9	6.8	10.3	13.2	15.1	17.4	24.1	29.7	32.5	
5										39.5
10	8.9	13.9	21.7	27.3	31.2	35.2	47.8	54.9	58.7	61.4
14	8.9	13.9	21.7	27.3	31.2	35.2	47.8	54.9	58.7	61.4
20	4	7.0	10.3	13.2	15.2	16.8	20.8	22.9	24.0	24.7
35	4	7.0	10.3	13.2	15.2	16.8	20.8	22.9	24.0	24.7
60	9.5	17.6	31.9	43.4	53.1	60.0	78.3	88.0	95.5	99.6
80	9.5	17.6	31.9	43.4	53.1	60.0	78.3	88.0	95.5	99.6

Next, we take a scalar wind speed profile such as the 95 percentile for Cape Canaveral (see main text).

Then we select a target altitude ‘hp’ on this graph upon which to superimpose a windshear. This might be the altitude that our sim says that ‘max q’ occurs, which is the maximum dynamic pressure (see glossary) suffered by the vehicle and would typically occur just at engine burnout.

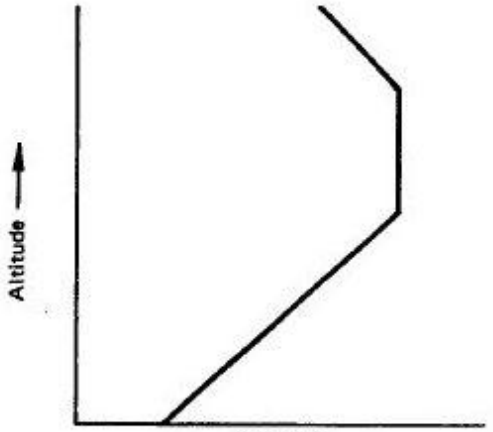

Step 1: Select wind-speed envelope.

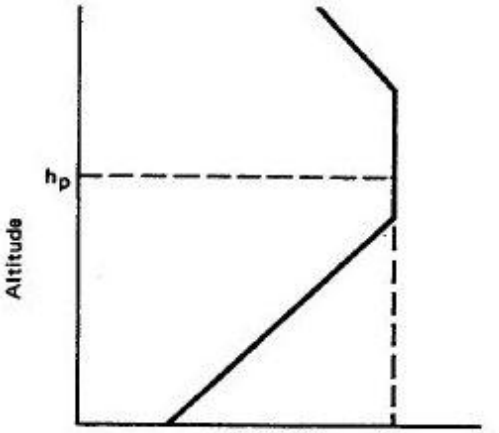

Step 2: Select desired altitude of peak wind.

Then starting from altitude hp, we subract the above wind speed change spectrum tabular data. Start at the smallest scale of distance which is 100 metres, for example if hp is 12Km, then the above data for 12Km and 100m scale of distance is 8.9 metres per second wind speed change. So plot a point whose ‘x’ is Vp minus 8.9 and whose ‘y’ is 100m lower than hp.

Then plot the other scales of distance values in the same way, joining the resulting wind-buildup curve to the origin of the graph by a straight line that is tangent to the curve:

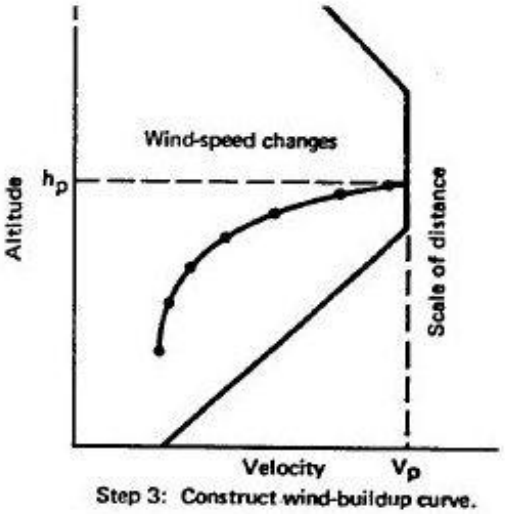

Step 3: Construct wind-buildup curve.

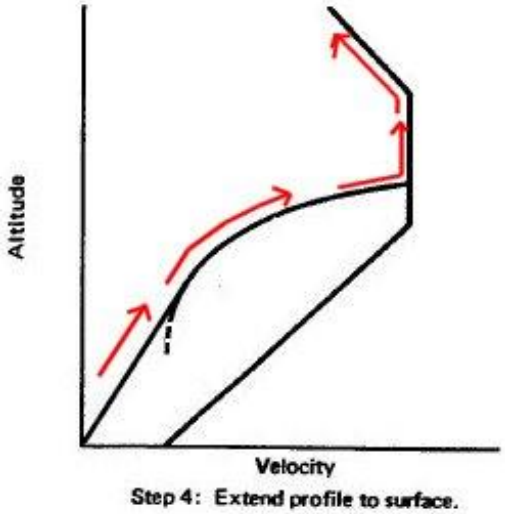

Step 4: Extend profile to surface.

As the vehicle ascends, let the wind speed follow the red arrows shown here.

To reduce the windshears by 15% simply reduce the wind values in the above table by 15%, keeping the altitude values the same. This increases the gradient of the curve (reducing the shear) by 15%

Appendix 2: numerically modelling turbulence

As real turbulence almost has a Gaussian distribution, the first step is to generate white noise with a Gaussian distribution. Some programming languages (such as java) already provide Gaussian white noise, but C++ doesn't, so first we'll generate this:

```
  // generate Gaussian white noise (i.e. white noise with a Gaussian amplitude distribution)
  // see http://www.dspguru.com/dsp/howtos/how-to-generate-white-gaussian-noise
  // The Central Limit Theorem states that the sum of N randoms will approach normal
  // distribution as N approaches infinity.
  // The drawback to this method is that X will be in the range [-N, N],
  // instead of (-Infinity, Infinity) and if the calls to rand() are not truly independent,
  // then the noise will no longer be white. Jeruchim, et. al., recommend N >=20 for good
 // results.

  white_noise = 0.0;

  for (N = 0; N < 100; N++) // use 100 randoms (N = 100) for more accuracy
  {
    // generate random number between 0.0 and 1.0
    bob = rand() / 32767.0; // rand() gives random integer between zero and 32767
    white_noise += bob;
  }

  /* for uniform randoms between 0 and 1, mean = 0.5 and variance = 1/12 */
  /* so adjust white_noise so that mean = 0 and variance = 1 (this is called standard normal) */
  white_noise -= N / 2.0;        // set mean to 0
  white_noise *= sqrt(12.0 / N); // adjust variance to 1
```

I checked that this code does indeed produce Guassian white noise: it does.
So we now have the Gaussian white noise (a sequence of Gaussian random numbers). We now pass it through a suitable digital filter, called the Infinite Impulse Response filter (with two memory terms) to generate a pink spectrum of turbulence:

```
  // pass white gaussian noise through an infinite impulse response filter (see Ref.1 page 57)
  // with two memory terms (called poles) to generate required pink noise

// initialisation, put before main loop
static double memory1 = 0.0, memory2 = 0.0;
static int count = 0;

// main loop
  alpha = 5.0 / 3.0;
  first_pole  = -alpha / 2.0;
  second_pole = (1.0 - alpha / 2.0) * first_pole / 2.0;

  if (count == 0)
    pink_noise = white_noise;
  else if (count == 1)
    pink_noise = white_noise - first_pole * memory1;
  else
    pink_noise = white_noise - first_pole * memory1 - second_pole * memory2;

  if (count < 5)
    count++;
```

```
 // save the last two pink_noise values
 memory2 = memory1;
 memory1 = pink_noise;

 // one standard deviation (SD) of 1000 pink noise values output from this software
// (on my computer, check what values yours gives)
// The most recent versions of Microsoft Excel have an SD function.
 const double standard_deviation = 2.179639788;

 normalised_pink_noise = pink_noise / standard_deviation;

// input turbulence intensity in percent (see main text)
 wind_standard_deviation = (turbulence_intensity / 100.0) * average_wind_speed;

// turbulence speed in metres/second
 turbulence = normalised_pink_noise * wind_standard_deviation;
```

Note that Ref.1 says that this software should be looped continuously at 20 Hertz, which is slower than most rocket trajectory sims loop. In this case, linearly interpolate the turbulence values with time before feeding them into the sim.