Supersonics and Waveriding

Intro

The important thing to note about a Waverider is that the name is a misnomer: a Waverider is *not* a particular shape of vehicle; it's any vehicle that to a greater or lesser degree Waverides, where ‘Waveriding’ is the key concept we’ll discuss here.
First we have to discuss supersonic aerodynamics in order to understand Waveriding, as Waveriding can only occur at supersonic speeds.

Items in **bold** appear in the glossary at the end of this paper.

Supersonic flow

Picture a rowing boat moving slowly across a calm lake. Water piles up in a hump against the bow as the boat moves along, or rather, the Kinetic energy of the boat is partly converted into Potential energy, here expressed physically as the mean height of the hump of water at the bow above the flat surface of the lake.
In aerodynamics rather than water dynamics, air pressure is the Potential energy.

Just as the hump of water at the bow of the boat extends a little upstream of the boat, so the pressure-field extends a little upstream of the nose (and wing-edges etc) of an aircraft travelling below Mach 1. In the boat’s case, the approach of the hull is transmitted upstream by a series of tiny potential-energy waves, in this case conventional water-waves, radiating out from the boat.
Aerodynamically, millions of incredibly weak *pressure* waves radiate concentrically upstream, ‘informing’(in an energy-transfer sense) the air upstream of the presence of the approaching aircraft.

Now greatly increase the boat’s speed: The waves no longer radiate out ahead of the hull, as they physically don’t travel as fast as the boat, so there’s no way the water just upstream can ‘know’ (in an energy sense) of the approaching hull, and so the water piles up in a thin, abrupt, bow-wave *at* the hull.

Pressure-waves also have a fixed speed of travel, known as Mach 1 or the Speed of Sound, since our ears react to pressure waves which we perceive as sound.
Aircraft flying at Mach 1, just as with the high-speed boat, experience an abrupt, piled-up pressure-wave known as a Shockwave. At higher speeds the shockwave bends backwards from the nose just as the bow-wave around a boat does. Now it is more correctly known as an Oblique Shockwave.

What *is* a shockwave? It is not a physical entity in itself, it just demarks where some of the kinetic energy of the incoming airflow gets converted suddenly, inefficiently, and noisily; (the air is 'shocked') into large changes in aerodynamic properties, such as abrupt changes in temperature, pressure, and flow speed. This occurs across a boundary only a few air molecules wide. At sea level this corresponds to fractions of a millimetre thick, but at the top of the atmosphere, the shockwave might be a metre or more across, at which point the aerodynamicist throws up his hands and has to give up.
However, at low level, the shockwave can be considered to have minuscule thickness and so changes in aerodynamic properties of the flow occur instantaneously upon passing through this infinitely thin shockwave. That makes the maths much easier: shockwaves are described mathematically as discontinuities.
Always remember that a shockwave is a dissipative phenomena, energy is being wasted: a supersonic flow is decelerating to a lower speed. Conversely, when the flow is accelerating to a higher speed you do *not* get a shockwave.

Wedge flow

Before describing the complex flow around real supersonic aircraft, we'll simplify the picture by starting with the flow around a simple two-dimensional shape. When Aerodynamicists talk about a two-dimensional shape, they're talking about a wing or aerofoil that is infinitely wide in span: i.e. infinitely long in the third spanwise dimension.

This means that if it is a wing sat in a windtunnel, then it literally goes from wall-to-wall, and is joined at the wall; there are no effects caused by the end of a wing, as happens in the real world.

Dealing with only two dimensions makes the maths a lot easier: classically, the way one proceeds in subsonic as well as supersonic aerodynamics is that you start with a known set of aerodynamics for a two-dimensional shape, and then you add to this three-dimensional effects, and this is where it gets complicated, and the maths goes up by an order of magnitude. It’s worth noting that just as with subsonic aircraft, these 3D effects are all adverse; lift is lost, and so 3D effects need to be minimized.

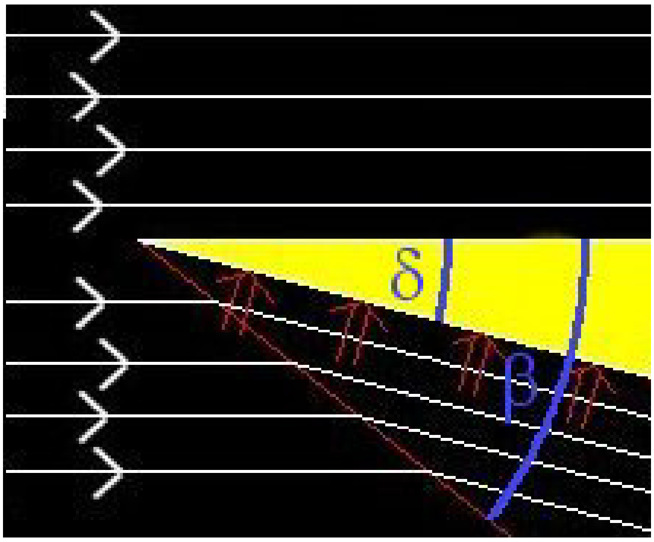

We'll start with two dimensions for simplicity: the shape shown here is known as a wedge; and does indeed look like a door wedge.

The top of this wedge happens to be parallel to the flow; it is neither windward nor leeward of the incoming air and so the top surface flow is completely unaffected by the wedge.

We’ll ignore the upper surface from now on as the lift is generated from the lower surface.

The underside of the wedge is seen by the supersonic airflow; it's to windward if you like, which inclines the flow at an angle called the angle of attack (denoted, oddly, as delta δ in supersonics, rather than the more conventional alpha α), and this causes a corresponding shockwave. This is inclined at an angle (beta β), and shown here beta isn’t 90 degrees, so the shock is known as an Oblique shockwave.

The compressed airflow on the wedge underside flows parallel to the underside, and the oblique shockwave moves out into the flow and settles at the required angle necessary to facilitate the flow-turning.

Oblique shockwaves turn flows in an interesting way: air flowing through a shockwave experiences sudden changes, such as a rise in pressure and drop in velocity *but only in a direction at ninety degrees to the shockwave.* The component of the flow parallel to the shock is unchanged, and a simple vector diagram of these two components (normal and parallel to the shock) will show that this bends the airflow toward the shock as it passes through it, similar to the way glass refracts light-rays.

The angles delta and beta are measured relative to the direction of the original incoming airflow before it encountered the shock, and are related to each other, and also to the Mach number *M* of the incoming flow, according to the following rather unhelpful mathematical formula:

$$\tan\delta = \frac{1}{\tan\beta}\left(\frac{M^2\sin^2\beta - 1}{\frac{\gamma+1}{2}M^2 - (M^2\sin^2\beta - 1)}\right)$$ where γ = 1.4 for air.

Interestingly, this formula gives two solutions, so there are two distinct and separate angles for beta that can occur for some angle delta, although in practice you only get one or other of the angles occurring. Solutions to this formula are given in the NACA tables of compressible flow, and graphically as:

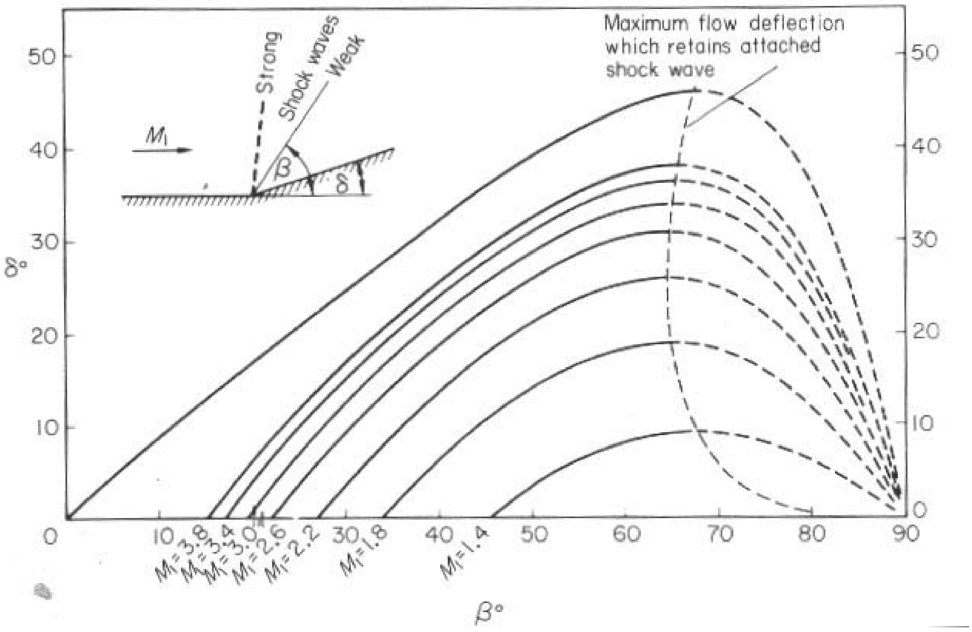

The two solutions are called the Weak solution (solid line above) and the Strong solution (dotted line above) respectively. The reason for these arbitrary names is that the Strong solution gives a bigger change in flow pressure across the shock wave, and hence as felt by the underside of the vehicle, than the Weak solution.

Which of the two solutions you will get is generally dictated by the shape of the object that the airflow encounters, although it is not always so clear in advance; sometimes the flow can flip between the two solutions, usually with disastrous consequences to the aerodynamics that you are investigating. The engines of the SR 71 Blackbird used to suffer from a condition that the pilot called 'unstart' where the shockwave across the lip of the engine intake could flip from one flow state to the other, causing a dramatic drop in thrust of the engine. So the SR 71 pilot always had to be on his toes was because if you are doing Mach 3 with two side-by-side engines and one of them suddenly shuts down, the aircraft will yaw like anything!

As examples of the two flow solutions we'll take two extremes:

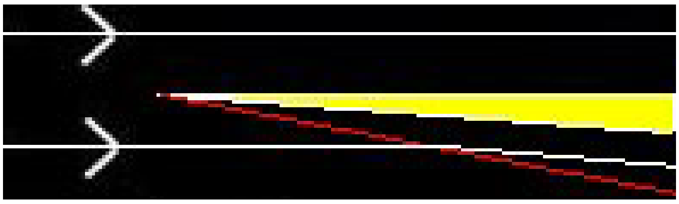

The first example shown in this diagram has a very small acute angle of attack which causes the weak shock solution. The shock angle beta is only slightly larger than the angle of attack delta and so the shockwave is physically lying very close to the underside of the vehicle.

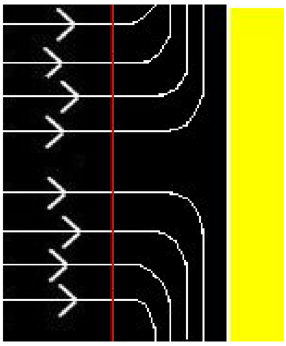

In this next diagram we see the incoming airflow hitting the vehicle at an angle of attack of 90 degrees; it's effectively hitting a brick wall, and so the flow is forced to turn a right-angle. There's no physical way a supersonic airflow can turn a right angle without slowing down to practically zero airspeed first; it literally rams into the wall and temporarily stops dead, giving what is known as Stagnation conditions because the flow is stagnant. The flow to the right of the shock is wholly subsonic as supersonic flow can't make the turn.

This is the most extreme condition that the incoming flow can suffer, and so this results in the highest pressures and temperatures on the vehicle, and so the Stagnation condition is a condition that aerodynamicists worry about; it's a very good worst-case indication of whether things are going to melt.

In the compressible flow chart above, you'll notice that there's a peak angle of attack for each Mach number plot. At a higher angle of attack than this peak, the flow must drop subsonic to make the turn.
Referring to the diagram above, you'll notice that the shockwave stands off the vehicle by a certain distance known as the (detached) shock stand-off distance. It has to, because at some point the flow drops through Mach 1 as it slows down to near zero speed at the wall, and the shock will occur where the flow is *at* Mach one.

In practice, the strong shock solution for attached shocks almost never happens; it never happens for the flow situations we'll discuss in this paper. Conversely, detached shocks are nearly always strong shocks.

So generally, as the angle of attack increases, the weak shockwave angle beta increases until the shock detaches as a strong shock.

Expansion fans

So that's a shockwave, which you may well have heard of, but for completeness, there is another equally important supersonic phenomena that you may not have heard about, called a Prantl-Meyer expansion fan, named after the two people who originally researched it.
In the following diagram the top surface of the wedge is not parallel to the incoming flow it's leeward of it, and so the flow will expand over the leeward upper surface through an expansion fan. Prantl-Meyer expansion fans are drawn *diagramatically* as a series of little waves known as Mach waves. It's assumed that in practice there are an infinite number of infinitely weak Mach waves, each of which gently turns the flow by an infinitesimal amount, and so an expansion fan is a gentle process of continuous expansion, unlike the sudden abrupt change occurring through a shockwave.
I don't know if anyone has ever actually seen an individual Mach wave, but the Prantl-Meyer expansion does exist.

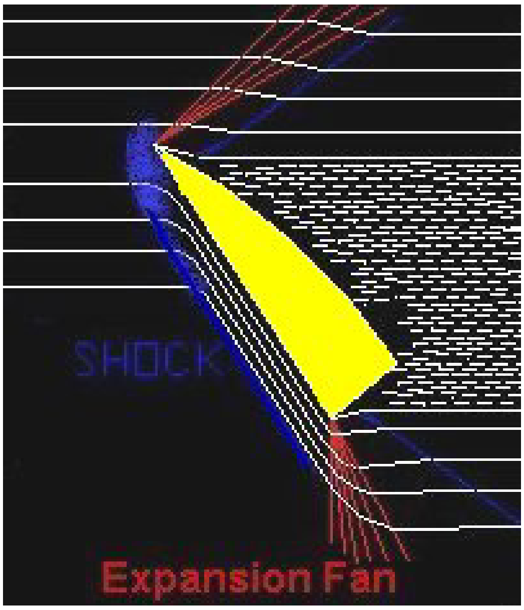

The Prantl-Meyer expansion fan is just as important as a shockwave; the two are essentially opposites. The shock causes a compression, an increasing pressure, whereas the expansion fan causes a decrease in pressure.

The picture above shows a wedge during re-entry: the angle of attack is at around 60 degrees, which gives maximum lift. At such a high angle, the underside shock is of course detached from the nose: the blue shading between shock and nose indicates a region of subsonic flow.
The dashed area is effectively dead, slowly recirculating air, a very low-pressure wake.

The most important thing to remember about shocks and fans is that they are effects, not causes. The cause is the geometry and inclination of the vehicle surfaces.

To recap, interferences in the main flow such as jumps in pressure, are limited to a speed of Mach 1, and so can't flow upstream against a Supersonic flow.
So Supersonic airflow, on encountering a gently inclined sharp edge, such as the nose or a wing leading edge of an aircraft, (even if it incorporates sweepback), is cleft in two by the edge into two distinct flows: over the upper and lower surfaces of the wing. Because pressure waves from either flow can't travel upstream back to the leading edge, then the two flows remain totally distinct, and they won't meet up again until exiting the vehicle at the wing trailing edge.

Blunt edges

At a blunt nose or leading edge however, or a sharp edge that is at too high an angle of attack, the flow nearly stops dead momentarily so it can negotiate the corner and so a local region of subsonic flow has to occur around the edge: the shock is detached. Pressure waves *can* travel upstream within the region of subsonic flow.
Assume that the wing is at a positive angle of attack so that the flow over the underside of the wing is at a higher pressure than the flow over the upper surface of the wing.
The higher pressure of the lower surface flow can now affect the upper surface lower-pressure flow within this subsonic region: some of the underside flow spills up and over the nose and leading edges, especially if they're swept back, and so some of the compressive lift under the wing is lost.

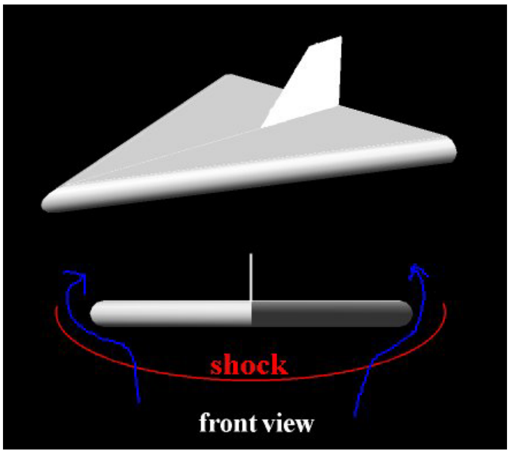

On the Space Shuttle which has very blunt nose and leading edges because of its particular thermal protection ethos as we'll discuss later, about 25% of its possible wing lift leaks away up and over the nose and leading edges.

Waveriding vehicles on the other hand, deliberately incorporate aerodynamically sharp edges to keep the shocks attached (below the shock-detaching angle of attack), and prevent this leakage.
What do we mean by sharp? the front edge of a real wing or wedge can never be made infinitely sharp. If it was infinitely sharp then it would have infinitely small cross-sectional area at the tip: then there would be very few atoms available to conduct even the smallest amount of heating away from the tip, and so it would overheat and melt and quickly become a rounded tip. And even if it could be made infinitely sharp, there is a real world effect called the boundary layer which gives even the sharpest edge a slight roundness.
One of the American X-series of vehicles had razor-sharp wings; so sharp that it required protective covers over the wing leading edges to protect the people working around it.

To all practical purposes, if you have a vehicle large enough to be man-carrying then any tip sharper than a millimetre or so radius can be regarded as aerodynamically sharp.

The slight roundness of a real-world tip causes a strong detached shock at the tip. Note that the subsonic region of flow at the tip is just a very local phenomenon around the tip: the sharper the tip, the smaller this region, and the less leakage can occur around the tip.

Because the edges are sharp on Waveriding vehicles, the shocks are attached to the edges; the vehicle then looks as if it is riding atop an attached shock, hence the name ‘Waverider’. Don’t be fooled by this ‘magic carpet shock’ description, the shock is just an after effect of sharp edges, it isn’t the cause of the Waveriding.

3D flow and real vehicles

As we’ll see, the whole point of Waveriding is to minimise unwanted and detrimental 3D effects so as to preserve the ideal 2D wedge-flow pictures above as much as possible.

In 3D a shockwave is known as a shock-surface.

One real world effect is the fact that unlike the 2D wedge, no wing has infinite span. At subsonic speeds, it's the fact that the wing does come to an end that causes vortices which cause a fair amount of the drag, known as induced drag.
Spillage of air from the high pressure lower surface, round the wing edges, to the low pressure upper surface occurs, which causes the vortices. Whenever the wing creates lift, induced drag is created.
At supersonic speeds, this spillage has the potential to occur, and designers should try hard to minimize it as it loses underside lift.
In addition, shock waves are created at the wing tips and a whole gamut of three-dimensional effects, all of which are unwanted, makes computation of the flow around the wings considerably more difficult.

Historically there were two different approaches taken to try to translate simple 2D wedge flows to 3D real vehicles.

First approach:

The first approach was experimental, and was used by the legendary spacecraft and aircraft designer Harrison 'stormy' Storms. He was an experimental aerodynamicist working on the NACA Langley research center's supersonic windtunnel.
He noticed that by bending the tips of a supersonic delta wing down at a slight angle (negative dihedral, known as anhedral) then the underside flow would not leak out from the tips, and over the leading edges, nearly as much as it would normally do. The reduction in lost lift over the edges greatly improved the aerodynamic efficiency (Lift-to-drag ratio) of the XB70 Valkyrie Mach 3 supersonic bomber.

By bending the wing tips down he was creating a physical barrier to restrict the flow spilling sideways over the sides.

With a flat underside, part of the flow merely has to turn 90 degrees to leak outward:

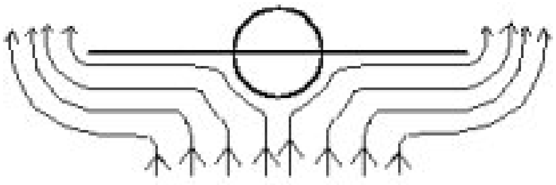

But by adding ‘sidewalls’ to the underside, the underside flow now will have to try to flow into the incoming flow, which is much more difficult. Trapping the incoming air this way in an underside cavity just as in the cup of a parachute or the underside cavity of a hovercraft, is the secret of Waveriding: the air can only escape out the rear of the wing.

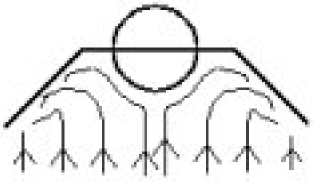

It was a simple solution, but very effective.
You can see the same method used on the British Blue Steel supersonic missile, and the cancelled TSR2.
The underside flow was less 3D because there was less spanwise flow but it was still complex.

Imagine a typical lifting-body toy model (e.g. a Space Shuttle) held upside-down and at an angle under a running tap. The water will hit the underside and spread out over the surface, but as well as running out over the back, the water will also spill out over the edges because the underside is either flat or convex: a spanwise flow is established: This water flow is more than just an analogy to supersonic airflow: the water reacts to changes in height on the model undersides, exactly as airflow under the real vehicle does.

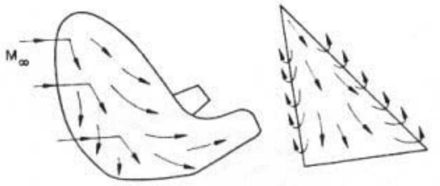

Flow spillage from convex and flat-bottomed vehicles.

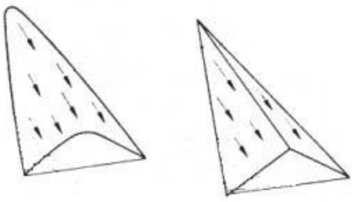

Flow containment by concavity

The underside of a Waveriding model is concave, therefore the water (airflow) will tend to stay in the middle of the underside before running out over the back, inhibiting spanwise flow and spillage.

Waveriding is therefore a two strategy ethos: sharp edges, and concavity, to keep the underside airflow from spilling over the edges.

Second approach:
The second approach was a numerical one, taken by a designer called Terry Nonweiler. Way back then in the days before computers, he wanted to make the computation of the flow on the underside of a supersonic aircraft as simple as possible, so he devised a three-dimensional aircraft underside which behaved physically and mathematically *exactly* like the flow under a simple two-dimensional wedge: no 3D spanwise flows or spillage whatsoever. This allowed him to use the two-dimensional wedge equations straight off.

At first, he strove to prevent this spillage purely to keep his equations solvable, but soon he realised that the spillage caused loss of lift, so solving this problem would kill two birds with one stone.
Terry also wanted to use a delta wing, as delta wings were all the fashion at the time for re-entry vehicle studies, the late 1950's.

There was good reason for the popularity of deltas, because the more you sweep the leading edges of an aircraft back, the less they heat up at supersonic/hypersonic speeds. This is because the air is striking swept leading edges at an angle rather than straight on.

He came up with what was called the 'Caret wing', so called because as seen from the front it resembles the ascii caret symbol ^ on your keyboard.
Like the XB70, the ‘Caret’ wing had pronounced anhedral, it resembles a conventional delta wing that has been folded down the middle, and indeed can be made this way as a paper aeroplane.

Imagine Nonweiler’s thought-processes running thusly:

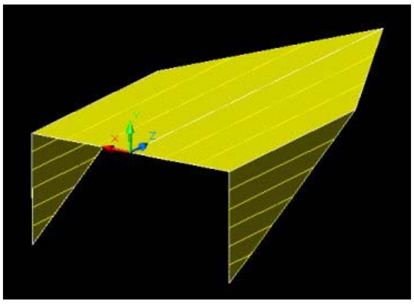

The first way to cure some of the spillage on a deltawing is simply to sidewall the wingtips:

Hmm, not overly aesthetic; flyable, but not terribly landable. Shockwaves will shoot off the side-walls to interact with the main underside shock, complicating the mathematical picture, and causing localized heating.

Here, the wing has been split into a series of small wedges and sidewalls. For the wedge parts of the wing, most of the flowfield is handleable mathematically.

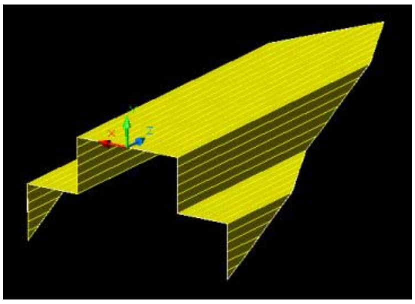

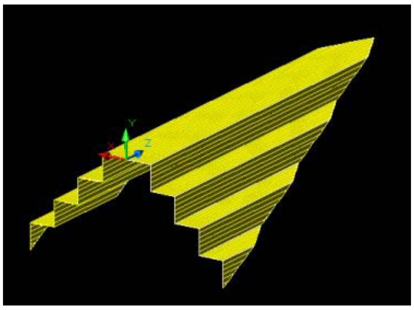

More of the same, a larger number of smaller sub-wedges.

Let’s increase the number of wedgelets to infinity to blur out the sharp steps across the span.

In the limit of infinitely small wedgelets you get a smooth, simple delta-wing, folded down the middle to give negative dihedral.

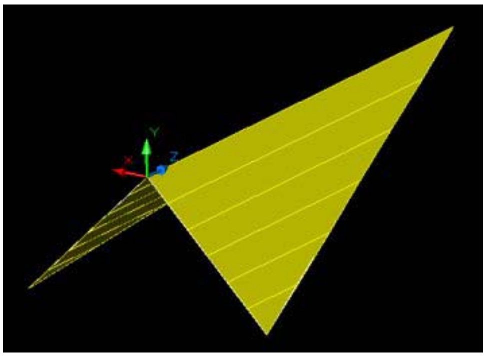

At a particular combination of angle of attack and Mach number, the shockwave on this Caret wing is a flat plane joining the wing edges, and the flowfield is essentially two-dimensional due to zero spillage:

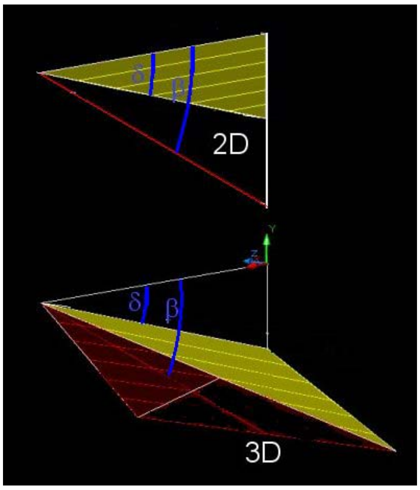

A good deal of work was carried out on the Caret wing in the Sixties and Seventies, when it was only possible to analyse the airflow for such a simple shape, to the extent that many people consider that all Waveriding craft are Caret wings.

It's also wrongly assumed that the Caret wing will only work for specific combinations of Mach number and angle of attack, known as 'on design' conditions; infact the Caret wing will happily perform at other speeds, it's just that the shocksurface will no longer be a flat plane, and so what? The shocksurface will be convex at speeds and angles of attack below 'on-design', and concave above 'on-design'. The mathematics of these conic shock surfaces are sadly much more complex of course.

The Caret is just one of an infinite variety of shapes that will support a planar '2D' shock and flow system:

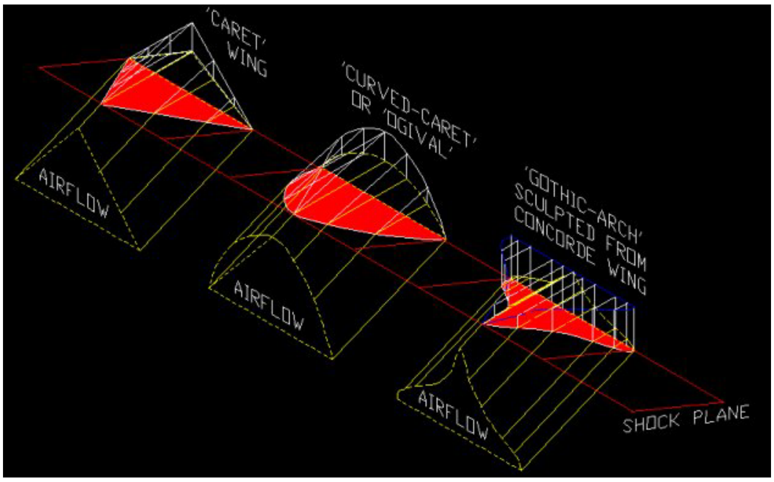

Actually, only the underside areas near to the wing leading edges need be concave to 'wall-in' the airflow; What happens to the flow downstream of the leading edges is somewhat immaterial if the wings are delta-like. The underside of the Caret wing for example can be filled-in to a large degree, which gives more fuselage internal space, and a lower vehicle centre of gravity (CG) to aid stability.

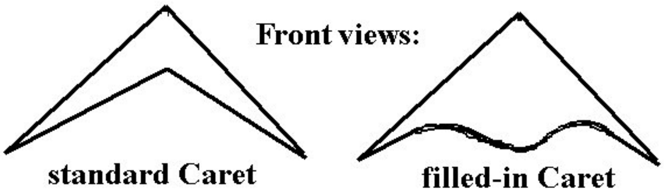

The old worries about Caret Waveriders not being structurally efficient and also being top-heavy because the centre of lift of the wings are below the centre of gravity; these worries disappear when you fill in the central cavity.

Some degree of concavity can be added to just about any supersonic delta-winged aircraft; every little helps to improve lift.

Types of Waverider

When you say 'Waverider', people form one of two distinct mental images in their minds.

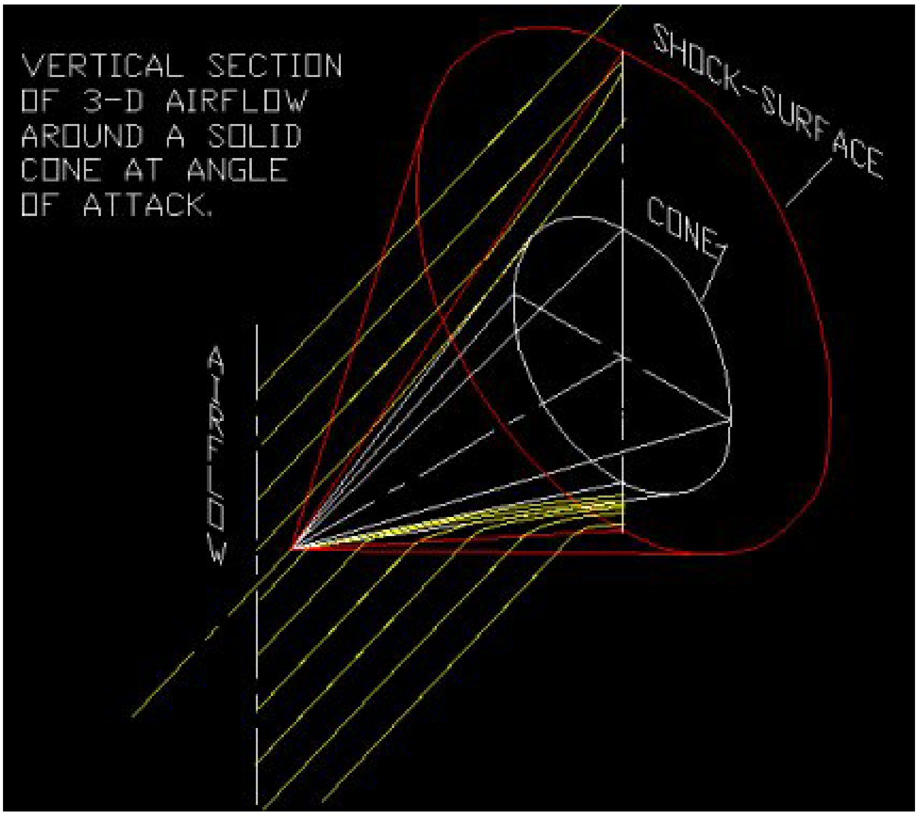

They either picture the original Caret wing or they picture the more modern cone-derived Waveriders, so called because instead of basing the flow on the mathematics of a two-dimensional wedge, these Waveriding craft are constructed from the three-dimensional flow around a low-drag nosecone.

You obviously get a curved shockwave, plus three-dimensional effects. But again, the weak shock is 'attached' to the leading edges. You can consider the Caret wing with its planar shock to be a cone-derived Waverider where the radius of curvature of the shocksurface is infinitely large.

The trouble is that the modern cone-derived Waveriders require a great deal of mathematics to describe the flow, and so they tend to be calculated on computer using Computational Fluid Dynamics (CFD).

The problem there is that the people doing the programme and analysis are Computational Fluid Dynamicists. These are computer kids who have absolutely no idea what a real aeroplane looks like, they've never even built a model one.

They computer-generate some wonderful shapes; some of which will fly, but a lot of which could *never* fly. Some of the shapes actually defy basic shockwave physics; however, if the computer says it's okay, then they publish them in learned journals. Oh, they say that the design is designed to maximise this, and minimise that, and is tweaked such and such a way, and what they end up with resembles a flying shoehorn, it's as useless as a real-world aircraft concept as the basic Caret.

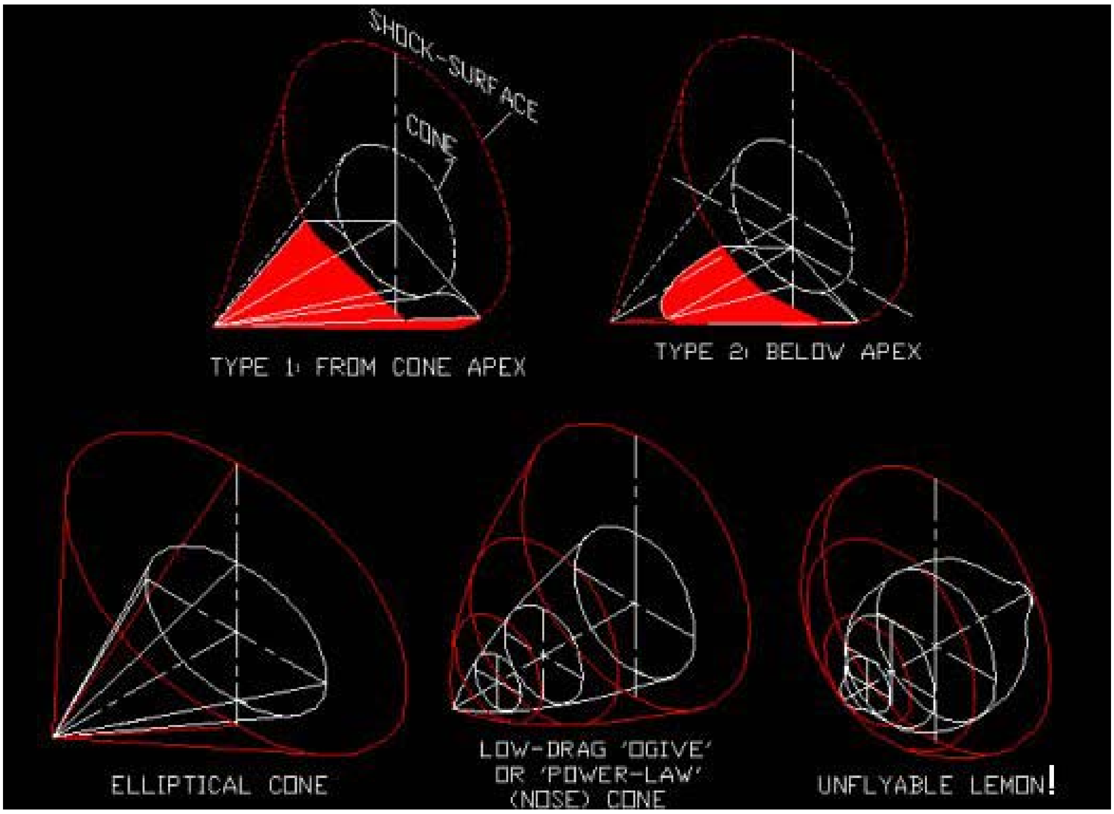

The better approach is to start with something flyable, then add concavity!

Vehicle design and increasing sweepback

One problem the CFD kids have with their cone-derived waveriders is that they do like their classic science fiction spacecraft: When viewed from above, (the planform), the leading-edges of Sci-fi spacecraft tend to look like a parabola. So as you go from the nose to the tail, the angle of sweepback is increasing all the time.

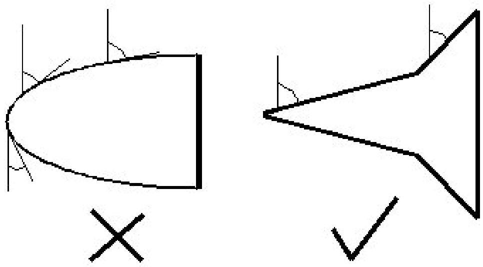

And this is precisely what you do *not* want to do.

If you look at the Concorde, Space Shuttle, the cranked delta of the F16Xl and suchlike, you will notice that except for the rounded tips where they are trying to minimise heating, the sweepback starts at a large angle at the nose, and then is *reduced* towards the tail; and that's actually much more stabilising, especially at high angles of attack.

Doing it the other way, is destabilising.

You can make a Caret wing fly stably at low speeds quite stably, but you cannot with a lot of these cone-derived beasties; I know, because I have made and flown models of both.

The design ethos should be for high angles of leading-edge sweepback at the nose, which spin-up a pair of vortices.

Picture horizontal tornadoes emanating from the leading edges and spilling over the upper surface. At large angles of sweepback, for example 70°, these vortices are extremely powerful, both mathematically and physically in terms of the velocity of spin. This spinning air induces the air flowing past it which then speeds up, and so from Bernoulii's principle, this faster-flowing air has lower pressure. This lower-pressure air is above the wing, so causes lift.

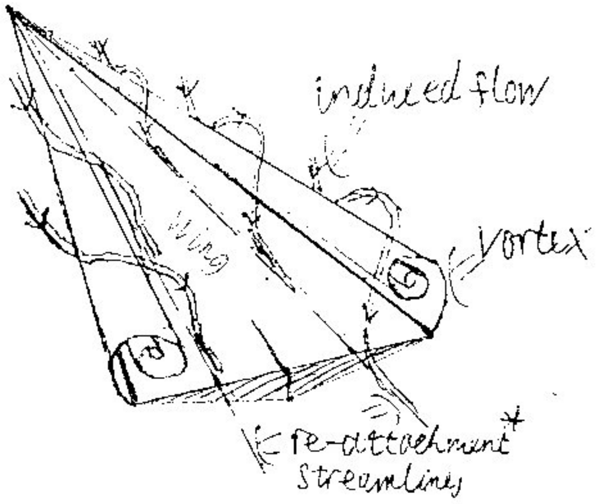

As you reduce the sweepback angle, the power (spin speed) of these vortices decreases. So what is generally done is that on the front of the vehicle, for example the Space Shuttle, you have about 70° of sweepback, to spin up a very fast pair of vortices, and then further aft the sweepback is reduced, to around 45°. This gives the famous double Delta platform as shown with a tick next to it on the diagram on the previous page. (the reduced sweepback aft gives a more aerodynamically efficient wing).

The powerful vortices forward energise the less powerful vortices rearward to keep them spun-up even at high angles of attack, and prevent the delta equivalent of a stall (where the vortices' spin drops to zero).

Therefore, double-delta and cranked delta wings can fly at much higher angles of attack than a straight delta, which is very useful for landing at low speed.

Anhedral and flyability

Going back to the humble Caret wing, every aeromodeller will look at a Caret and pronounce it unflyable because of its negative dihedral (known as anhedral). They know that positive dihedral (having the wingtip higher than the wing root) aids stability whereas negative dihedral causes an unstable aircraft.

However, there is a little-known law contrived by one of the fathers of the delta wing: Dietreich Kuchemann, that every five degrees of leading-edge sweepback has the same stabilising effect as one degree of dihedral.

And so you can balance the anhedral required for the Caret wing with a sufficient amount of sweepback to get a net stabilising effect. Caret wings are remarkably stable at low speeds, even at the high angles of attack required for landing, because if you think about the geometry, the leading edges of the Caret wing are somewhat like the leading edge slats on conventional aircraft; the airflow sees them as drooped.

Caret wings can also exploit ground effect: the incoming airflow gets trapped between the runway and the underside cavity, the same trick used on the fast passenger carrying catamarans.

Aspire and our friends have had tremendous fun over the years flying Caret-wing models off of kites, cliffs, and piggy-backed off of rockets at apogee.

Reentry

There are two design aspects to reentry vehicles: mitigating the g-acceleration and dealing with the aerodynamic heating.

The current generation of winged space tourism designs are suborbital; From an apogee of just over 100 Km, peak Machnumber during reentry is just over three, so heating isn't an issue.

The g's build up during the pull-up manoeuvre from vertical plummet to near-horizontal glide:

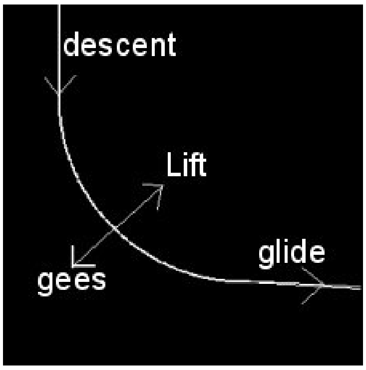

It's important to perform this manoeuvre as high up in the atmosphere as possible, to mitigate both the g's suffered, and also the peak **dynamic pressure** (max aerodynamic loads).
This is because the atmosphere doesn't increase linearly with decreasing altitude, it increases exponentially with decreasing altitude.
So the *rate* at which the atmosphere is thickening around the craft as it descends at some vertical velocity is much gentler higher up, so if re-entry is performed higher up, then the deceleration to low speed is spread-out over a much larger vertical 'braking' distance, which lowers the gees.

On top of this there's the simple issue that there's less height to fall between 100 Km and the top of the atmosphere compared with 100 Km to the lower atmosphere: you simply haven't built up so much speed.

To perform this pull-up manoeuvre high up requires a lot of Lift; from the **Lift equation**, as the air density is low, and the airspeed moderate, then the craft has to be designed to generate a high lift coefficient and have a large wing area for its mass.
Supersonically, the lift coefficient increases with angle of attack up to a maximum angle, then decreases.
At such a high angle of attack, the shocksurface will detach from the leading edges, and yet it has been found in the windtunnel that aircraft with a region of underside concavity can

maintain their underside flow attached to their leading edges to higher angles of attack than those with a flat or convex underside; this means that Waveriding craft can attain higher lift coefficients than comparable flat or convex-underside deltas, their lift peaks at about 60 degrees angle of attack.
For these reasons, simulations show that a Waveriding craft performing Spaceship One’s pull-up manoeuvre fromvertical dive to horizontal glide peaks at 3 gees rather than the current 5 gees.

Hypersonic flow and heating
Re-entry from orbit occurs at hypersonic Mach numbers.
There’s no fundamental Mach number above which the flow changes immediately from Supersonic to Hypersonic, unlike the change from subsonic flow to supersonic flow, which definitely occurs at Mach 1. This is because the term ‘Hypersonic flow’ covers several categories of flow properties; thermal, chemical, etc., not all of which change at the same Mach number.
When supersonic flow is going fast enough to cause chemical effects it can then be called Hypersonic flow. Molecules of the gases that make up the air are generally dumbell shaped, or 'diatomic'. But at hypersonic speeds the air molecules are literally shaken apart upon passing through the shockwave; this is called disassociation: the diatomic molecules are split into a hail of little spherical bullets. So hypersonic is a rather arbitrary term because the different gases disassociate at different Mach numbers.

Let’s pick Mach 5 as the onset of Hypersonic flow anyhow.

As an interesting sidenote, Sir Isaac Newton, Britain's most prolific mad scientist, had a go at trying to describe aerodynamics in mathematical form. Unfortunately he did not know that the air molecules are diatomic; and dumbells can sop up heat by spinning and vibrating in many more ways than monatomic single bullet shaped molecules.
Diatomic molecules can vibrate in about five different ways at moderate temperatures, and this fundamentally affects how such gasses behave. Newton did not know that, and he also did not know about the viscosity (syrupiness) of the air, and so his aerodynamic formulae do not work at low speeds.
Newton simply modelled molecules as spherical bullets that do not interact with each other and simply smack into the surface of the vehicle, colliding elastically, and so he simply worked out what the change in momentum of each sphere would be and summed the effects of all the spheres hitting the surface to calculate pressure.
Interestingly though, diatomic molecules ‘break’ (disassociate) into monatomic molecules as they pass through the hot shock and boundary layer around a hypersonic vehicle. They then are just a hail of spherical bullets and so Newton's simple formulae of long ago work surprisingly well, and give good first estimates of the lift and drag of hypersonic vehicles!
The flow is then called Newtonian flow.

My fellow Waverider researcher Gordon Ross has quipped the term 'megasonic' which corresponds to Mach 25. This is a useful yardstick because spacecraft entering the Earth's atmosphere from low Earth orbit (such as the Space Shuttle) encounter the top of the atmosphere at Mach 25 or 'megaMach 1'. A large yardstick is needed because space probes plunging into the atmospheres of other planets tend to do so at extremely high Mach numbers.
The Apollo 13 command module holds the manned speed record at Mach 36 (megaMach 1.4) as they were in a hurry to get home.

Heading and sharpness of edge

It was discovered in the 'fifties that the sharp noses and wing leading edges of rockets soon melted off at hypersonic speeds. It was found both mathematically and experimentally that the heating *q* experienced by a nose or leading edge depended upon the nose radius *R* of the edge:

$q \propto \sqrt{\frac{\rho}{R}}$ where ρ is air density.

Consequently, most vehicles flying above Mach 5, such as the Space Shuttle, have profoundly rounded noses and wing-edges; seriously blunt.

On the face of it, this is bad news for Waveriding craft which require sharp noses and wing leading-edges (tiny *R*).
Although sharp edges do get a lot hotter than rounded ones at the same air density, their improved lift means that Waveriding craft can fly at much higher altitudes where the air-density ρ in the above equation is lower: you can see that this decreases the heating.
Note also that reduced density means that the chances of laminar flow occurring over most of the underside are good.
Laminar flow is flow where the layers of air next to the vehicle skin slide smoothly over one another, so that the hot, outer layers of a hypersonic flow are kept away from the skin.

So infact while Waveriding craft invariably do suffer higher re-entry heating, the heating is just within the limits of current materials technology.
A hot, sharp, insulating tip would quickly shatter due to thermal stresses, and would doubtless soon melt as well, so as Nonweiler has shown, the nose and leading edges must be heat-conductive rather than the insulative tiles of the Space Shuttle.
Townend and Nonweiler performed a study of a winged 'ambulance' re-entry vehicle to be used to ferry injured astronauts back to Earth from the International Space Station.
The nose and leading edges of their 'SLEEC' craft were to be constructed of a 16mm long solid wedge of niobium followed by 70mm of solid graphite coated with niobium.

Both niobium and graphite are heat-conductive. Unfortunately they're also very heavy, and niobium is ungodly expensive. The other problem is that hot metal expands, and if severe thermal stresses are to be avoided, movable metal scales or 'shingles' are often proposed. Unfortunately, experience with the X15 aircraft showed that the joints between metal scales interfered with the air next to the vehicle's skin, causing it to become turbulent. This mixing lets hot hypersonic air 'touch' the skin and locally roasts it.

Sadly, high-hypersonic Waveriding craft are an idea ahead of their time, waiting for new materials capable of taking very high heating without shattering, melting, or expanding much. One day such materials will be developed (see http://www.space-rockets.com/sharp.html for some new ceramics) and then the benefits of Waveriding can be realised for re-entry from Earth orbit.
Townend and Nonweiler's study showed that the use of their Waveriding craft allowed the re-entry gees to be held at just over one to minimise further injury to an injured astronaut during re-entry, and their vehicle had such good aerodynamic performance that it could glide to land at almost any airport in the world.

Nonweiler had a further novel approach to re-entry heating:
Presuming thermally conductive edges, then this heat could be channelled to cooler sections of the vehicle.
He noted that the flowfield over the upper surface of a Waveriding craft is separate from the hot underside (unlike the Shuttle, where hot spillage from underneath roasts the upper fuselage), therefore Waveriding craft were ideally suited for re-radiative cooling:
The heat received from the edges, and the metal underside of the vehicle, is conducted up through the hot vehicle skin and outer fuselage structure, and is radiatively dumped off the cooler metal topside into the upper surface flow, and ends up in the wake behind the vehicle.

The Space Shuttle partly re-radiates underside heating back off the glassy underside heat shielding, but for this to work effectively, the heat-tiles have to be at a much greater temperature than the underside flow, so some what unbelievably, the Shuttle is deliberately flown lower in the atmosphere during re-entry to deliberately heat the heat-tiles much more than is required!
The remaining heat pours into the tiles and so the shuttle's re-entry is therefore time-limited. Even insulating heatshield tiles let a trickle of heat through, and it is hoped that the Shuttle lands before the tiles leak too much heat into the vehicle's aluminium structure (which is nicely flammable).
In contrast, a Waveriding craft can re-radiate heat off the top surface at the same rate it receives it underneath, i.e. the vehicle is in Thermal equilibrium, so it can take as long as it likes for re-entry as there isn't a net build-up of heat to reach dangerous levels. This also makes Waveriding craft the better candidate for long hypersonic flights, such as a hypersonic airliner.

In summary, Waveriding supersonic aircraft out-perform other supersonic aircraft, and hypersonic waveriding craft will do so one day, just as soon as new materials have evolved to handle the heating issues.

Glossary:

Dynamic pressure: (q)

All aerodynamic forces scale directly with the kinetic energy term $\frac{1}{2}\rho V^2$

ρ being volume-specific mass, or air density, and V = flow velocity.
This kinetic energy term is called Dynamic Pressure (q), to distinguish it from it's Potential energy counterpart of static pressure (P).

Lift (equation):
Lift is a force generated by aircraft at right-angles to their flightpath.
The equation used to calculate lift is simply the *lift coefficient, Cl*, times **dynamic pressure**, times some reference area '*S*', i.e: $L = \frac{1}{2}\rho V^2 S\ Cl$ (ρ = atmospheric density.)
For aircraft, this reference area 'S' is the total wing area.