How the Wing Got its Lift—Expanded Version

P. Marchal

The popular explanations of why a wing produces lift typically involve some (often rather obscure) combination of “Bernoulli’s Theorem” and “Newton’s Third Law”. These explanations do not come close to correctly describing the physics of the phenomena. An explanation rooted in the physics of aerodynamics can unfortunately not be summarized in a few soundbites—the generation of lift is indeed a complicated story… That does not mean, however, that a qualitative explanation based on sound principles cannot at least be attempted—that’s my goal in this note[1].

- Let’s first go through some definitions and assumptions. In the following, we’ll concern ourselves exclusively with two-dimensional (2-D) flows around airfoils (the shape of the cross-section of a wing). All real flows are three-dimensional, of course—by two-dimensional, we mean a flow that does not *change* along the third dimension. If we think of an infinitely long wing, for example, the flow in any given cross section will be identical to the flow in any other cross-section--none of the flow characteristics (e.g., velocity, pressures) will depend on the position of the section along the span of the wing. In real-life, wings are of finite span, but the flow close to mid-span for a high aspect ratio wing will be a good approximation of a 2-D flow, and so will be the flow around a wing of finite span bounded by the vertical walls of a wind tunnel, which is how experimental airfoil characteristics are determined. These two-dimensional characteristics are the reference from which the properties of the flow around a wing of finite span are calculated.
- We will consider the flow incompressible. This might not look like an obvious assumption in the case of air, but it can be shown that, as long as the Mach number (the ratio of air velocity to the speed of sound) is small--less than 0.4 is a generally accepted limit--air is for all practical purposes an incompressible fluid, and we can consider the density constant everywhere.
- We’ll also assume that air is inviscid, that is, that it has no viscosity at all. This is an obvious simplification, but it has been confirmed experimentally that, as long as some conditions are met, it is a good first approximation in many cases. We’ll identify these conditions as we progress in our analysis, and will revisit the inviscid assumption as needed.
- And finally, we’ll mostly consider steady-state flow—at any point in space, the fluid characteristics do not vary with time.

Let’s start by investigating the actual flow of a fluid around a symmetrical airfoil immersed in a uniform flow, at zero angle of attack. By design, an airfoil is *streamlined*, that is, it is of such a shape that the fluid will closely follow its contour, without any flow separation. This is one of the conditions required for the validity of the inviscid assumption.

---

[1] Note to the reader familiar with aerodynamics—in an attempt to make this write-up as accessible as possible, I purposely stay away from the traditional concepts of conservation of vorticity, circulation, Kutta-Joukowski theorem, etc. I limit myself to the more familiar and intuitive concepts of pressure and velocity, sacrificing some rigor along the way.

The picture below displays a flow visualization corresponding to that situation, obtained in a wind tunnel. The air flows from left to right, and smoke injected at regularly spaced intervals upstream of the airfoil allows us to observe the trajectories of fluid "particles" as they pass near the airfoil—the *streamlines*.

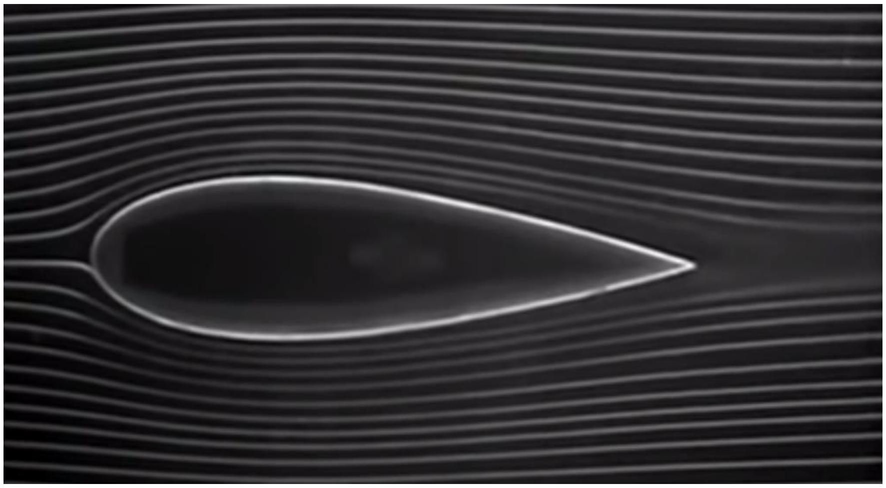

In the absence of any obstacle, the streamlines would be equidistant horizontal lines. The presence of the airfoil forces the streamlines to part away shortly upstream of it, and they rejoin downstream. Far from the airfoil, the streamlines are mostly undisturbed from uniform flow, but close to it, they curve to follow the contour, and are noticeably squeezed together nearing the point of maximum thickness.

From this simple picture, we can derive some useful physical insights about the flow. Let's look at it in more detail.

In the following 3-D view of a slice of with *w* of our wing of infinite span, we identify two neighboring streamlines. Each streamline in the 2-D section corresponds in three dimensions to a stream surface.

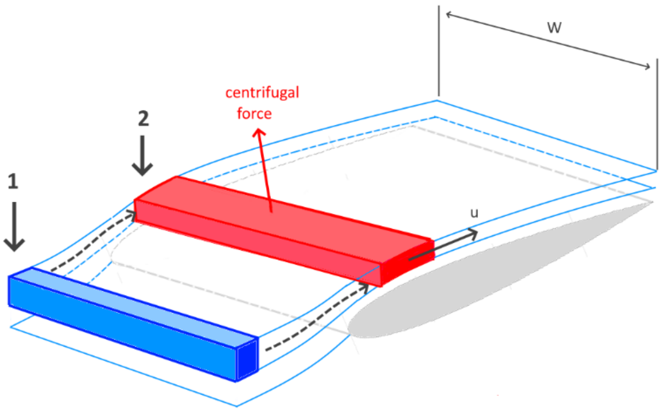

Sandwiched between these two stream surfaces is what we'll call a stream tube. Air that enters the stream tube at the upstream position 1 will forever be confined to it--for example, the fluid element colored in blue in the figure will travel down the stream tube, and after a certain time, will be found in location 2 (colored in red).

So, we can form a mental picture of the flow around the airfoil as a “layer cake” of stream tubes stacked upon each other. We can dig deeper and look at the internals of a stream tube, i.e., we can apply Newton’s laws of motion to the red fluid element at position 2. We do just that (with some math—but we keep it to a minimum!) in Appendix A--the interested reader is encouraged to refer to it. The main conclusions are the following:

- Rule 1. A stream tube will be curved only if the pressure on one side of the tube is greater than on the other side--the higher pressure “pushes down” on the stream tube and bends it towards the lower pressure. The greater the pressure difference, the greater the curvature, and the higher the fluid velocity, the higher the pressure difference for the same radius of curvature. At position 2, for example, the tubes are curved towards the airfoil. For each layer, the pressure “below” will be lower than the pressure “above”, and as a result we’ll see a decreasing pressure across layers if we approach the airfoil from above. To visualize this, we can imagine that, since the particle path is curved, centrifugal force tries to “pull” it away from the surface, creating a suction force acting on the surface.
  As a rule, pressure will always decrease when traversing layers towards the center of curvature and increase in the opposite direction.
- Rule 2. The quantity $p + \frac{1}{2}\rho V^2$ is constant—Bernoulli’s Law! (p is the pressure, V the fluid velocity, ρ the density.) Given our assumptions, Bernoulli is just another name for Newton’s second law of motion. If the pressure at position 2 is less than at position 1, the velocity will be higher.

In addition, we’ll note that the flow rate through a given stream tube is constant, independently of the streamwise location. That implies that, if the velocity at position 2 is higher than at position 1, the stream tube must be “thinner” at that position. This is Rule 3.

Going back to our “layer cake” model, we can now build a qualitative picture of the flow around the airfoil. Since the airfoil is symmetrical and placed at zero angle of attack, the flow will also be symmetrical about the airfoil chord line. To simplify the reasoning, we’ll discuss the upper half only.

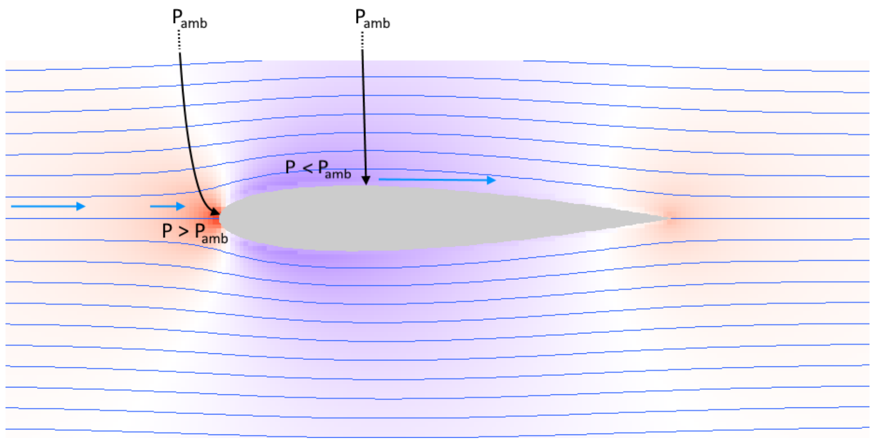

Far upstream, all layers (a.k.a stream tubes) are parallel and of equal thickness (uniform flow), and the pressure is uniform and equal to the ambient pressure. Approaching the airfoil, the lower layers face an obstruction (the leading edge) and have to curve upwards, in the process forcing the upper layer upwards also. According to rule 1 above, this flow pattern corresponds to an area of pressure higher than ambient around the leading edge, where the layers are concave upward (the centrifugal force "pushes" fluid particles towards the leading edge, "pressing" on the surface). The highest pressure is on the contour of the airfoil, and decreases towards ambient moving away from the airfoil across the layers (along the arrow on the left for example). As the fluid moves downstream, following the contour of the airfoil, the layers turn concave downward and as a consequence, the pressure will decrease from ambient when moving towards the airfoil from far away, following the arrow on the right for instance (the centrifugal force "pulls" the particle away from the surface). Thus we'll observe an area of pressure lower than ambient on top of the airfoil. To this lower pressure corresponds a higher velocity (Bernoulli), and thus a thinning of the layers. Moving further downstream, the curvature of the layers decreases, and even reverses in the vicinity of the trailing edge, where the streamlines return back to horizontal. The pressure on the airfoil will accordingly gradually increase as the flow progresses downstream, returning back to close to ambient (actually slightly positive) near the trailing edge. The increase in pressure results in a decrease in velocity and a thickening of the layers.

Using only a few rules deriving directly from Newton's law of motion, we were able to get a pretty good feel for the flow around the airfoil! In summary:

- The presence of the airfoil forces the streamlines to curve.
- The curvature of the streamlines induces variation in the pressure in the vicinity of the airfoil—higher than ambient around the leading edge, and lower on the top (and bottom) sections.
- Variations in velocity and steam tubes thickness follow from the changes in pressure.

If we want to get quantitative results (e.g., the actual pressure on the airfoil surface, so that we can calculate lift), things are getting more complicated, since we don't know a priori the exact shape of the streamlines. And the shape of the streamlines determines the pressure, which determines the velocity, which changes both the pressure (the centrifugal force driving the pressure change varies with the square of the velocity) and the shape of the streamlines (by Rule 3)… Fortunately, there are computational methods (the mathematical underpinnings of which are beyond the scope of this note) that allow calculating the flow velocities that match exactly the rules previously discussed. Such a computer code was used to generate the pressure field displayed on the figure above (red areas pressure greater than ambient, blue for lesser than ambient—the deeper the color, the larger the difference). Experimental data acquired in wind tunnels show that the predicted pressure distribution on the surface of such an airfoil is pretty accurate—we have seemingly a good tool at hand.

What happens to the flow as we increase the angle of attack to a small positive value—say 10 degrees?

Once again, turning to flow visualization provides a qualitative understanding. The pattern on the top side of the airfoil is similar to the one for zero angle of attack, with a stronger curvature of the streamlines around the leading edge—the pressure on that side will be below ambient, more so than for a zero angle of attack. The flow pattern around the bottom part is different though—the curvature of the streamlines away from the airfoil, combined with "thicker" stream tube layers (which implies lower velocity than on the top side), point to a higher pressure than ambient over a large portion of the bottom of the airfoil. Lower pressure on top, and higher pressure on the bottom, combine to produce a net upward force—we have lift!

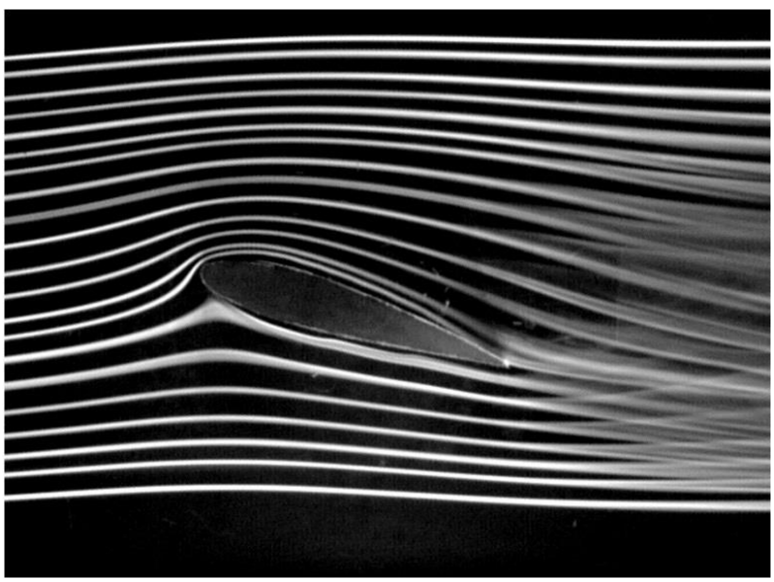

However, if we run the computations with the new angle of attack, we obtain results that do not match the empirical data at all…

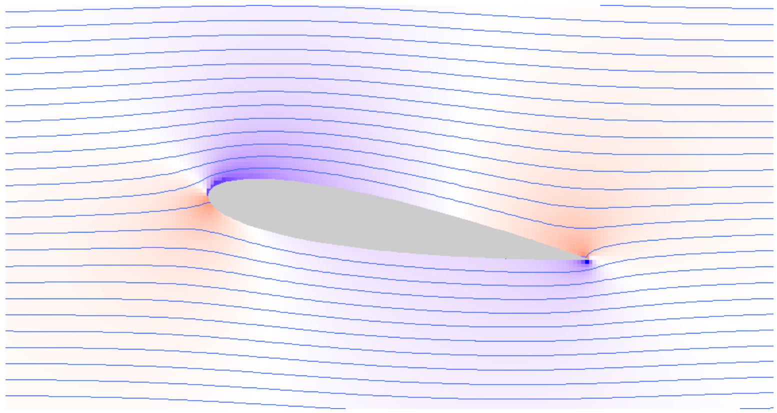

The streamline pattern is markedly different, especially around the trailing edge, and more concerning even, the calculated lift is exactly zero! What went wrong?

Brief historical note before answering this burning question. The problem we just encountered has been known since the middle of 18th century, when the French mathematician d'Alembert noted that, under the assumptions we made here, the total force applied by a fluid in motion on an immersed object will always be null—no drag, no lift… the so-called d'Alembert paradox[2]. It is only in the early years of the twentieth century, with the development of boundary layer theory by Prandtl and the theory of lift by Kutta and Joukowsky, that theoretical aerodynamics reached a point where lift and drag were understood and useful predictions could be made. By that time, the Wright brothers had been flying for years (first powered flight in 1903). Good thing the practitioners didn't wait for the theoreticians…

Back to our calculated pattern of streamlines. The odd thing about it is that the streamlines curve sharply around the trailing edge, as shown if the zoomed-out figure below.

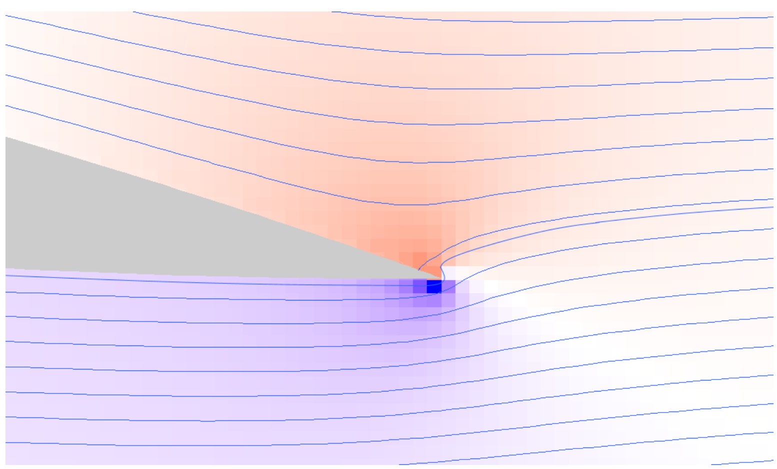

We know intuitively that this is not reasonable—in real life, the flow will separate from the bottom of the airfoil at the trailing edge and will tend to proceed straight downstream. It can be shown that any fluid (like air) that has *any* viscosity, even a very small one, will separate at the trailing edge—ignoring

[2] This is not entirely correct--d'Alembert paradox strictly addresses drag only. It can be shown however that, in a lot of cases, there is no applied force at all--no lift, no drag.

the effect of viscosity at that location is just not an acceptable assumption. This trailing edge separation impacts the flow around the airfoil—it acts like a "barrier" to the flow around the trailing edge, forcing the streamlines to rearrange themselves. The dashed blue line represents a reasonable approximation of the boundary between the flow passing on top and bottom of the airfoil:

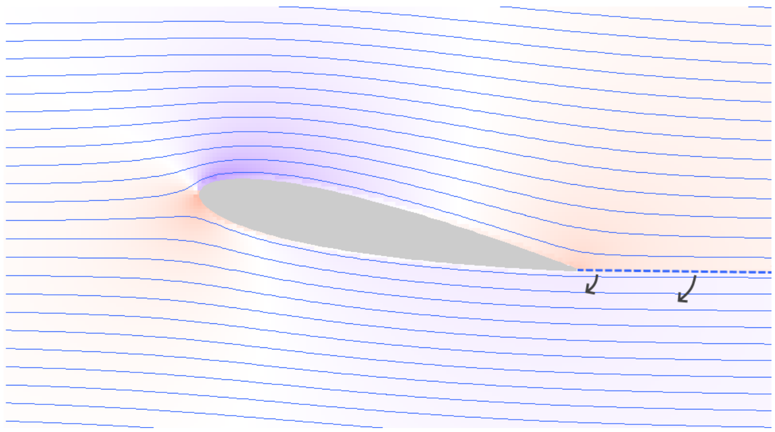

In accordance with the "curvature rule," the pressure above the separation boundary is higher than below (and thus, by Bernoulli, the velocity above is lower than below). This is not sustainable however—the separation boundary is not a solid surface, it is just another streamline. Under the effect of the pressure differential, it will deform in the direction indicated by the arrows, until the pressure differential vanishes, ending in the configuration illustrated below. [3]

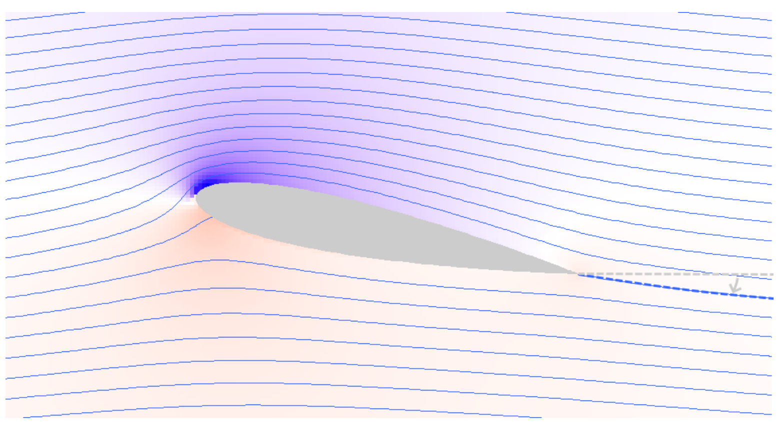

This pattern is now comparable with the one from flow visualization, with the streamlines from the upper and lower sides of the airfoil joining smoothly at the trailing edge. There's no pressure differential across the trailing edge anymore, and so (by Bernoulli), the velocities of the flow from above and from below will be equal where they join. And there is no flow separation there anymore either. This pattern along the trailing edge is known in aerodynamics circles as the "Kutta condition." Visually, the Kutta

[3] This is an over-simplified picture. The actual behavior is more complex, involving rotational motion, but results in the same final flow configuration.

condition corresponds to a streamline pattern that follows smoothly the surface of the airfoil, from leading edge to trailing edge, and leaves the trailing edge tangentially to the camber line.

Let’s now compare the flow and pressure patterns between the original calculation and the one meeting the Kutta condition, including this time a representation of the flow velocities as arrows, the length of which are proportional to the speed at the arrow origin. Two streamlines closest to the surface of the airfoil are highlighted for emphasis.

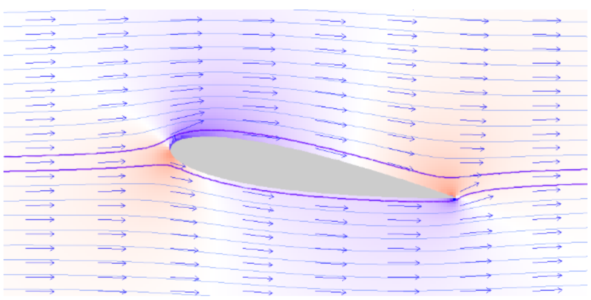

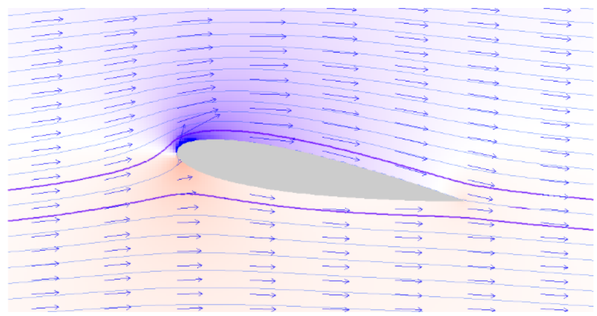

Under the pure inviscid condition (a.k.a, “ideal fluid”, on the left), the streamline pattern is almost symmetrical with respect to the center of the airfoil—the curvature of the streamlines is such that areas of high pressure exist at the bottom of the leading edge, and top of the trailing edge, and areas of low pressure exist at the top of the leading edge and bottom of the trailing edge, balancing each other. Under the Kutta condition, on the other hand (on the right), an upwash upstream of the leading edge, and a downwash downstream of the trailing edge, distort the symmetry of the flow and results in streamlines curvature such that the top side experiences lower pressure only (except for a small recovery zone near the trailing edge), and the bottom side experiences higher pressure over its entire length—producing net lift.

We can acquire a deeper insight by looking at the difference between the flow velocities under the two regimes. The next figure displays the difference between the velocities under the Kutta condition and the ideal fluid regime (the velocity scale is increased compared to the previous figures).

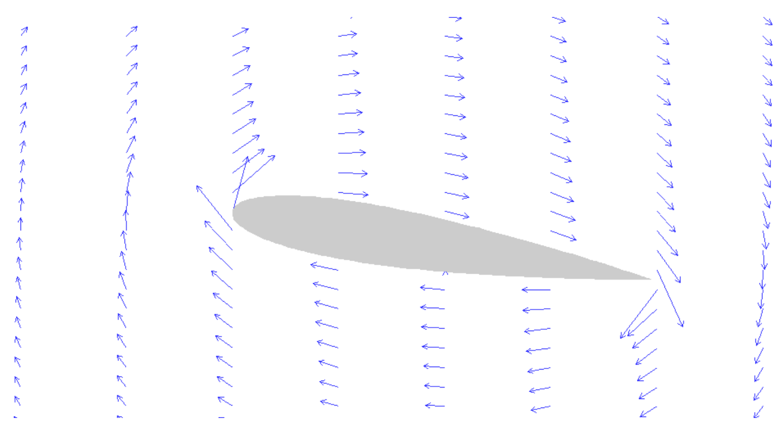

A striking pattern—meeting the Kutta condition is equivalent to superimposing a vortex-like flow to the original ideal fluid flow, a flow that reduces the velocity on the bottom of the airfoil (and thus increases the pressure, by Bernoulli), and increases the velocity at the top, reducing the pressure. It also shows the genesis of the upwash (the vortex flow displays upward velocity near the leading edge), and downwash (downward velocity near the trailing edge). Historically, recognizing the existence of this rotational flow has been central to the development of the theory of lift by Kutta and Joukowski.

We were able to paint a qualitative picture of the flow around an airfoil at positive angle of attack, and explain the generation of lift. Let's summarize. We'll start from a symmetrical airfoil at rest, and start a forward motion

- The air attempts to flow according to the rules of motion for inviscid fluid we discussed earlier.
- This resulting pattern involves the air making a sharp turn around the trailing edge, which is not physically possible for a real, viscous fluid.

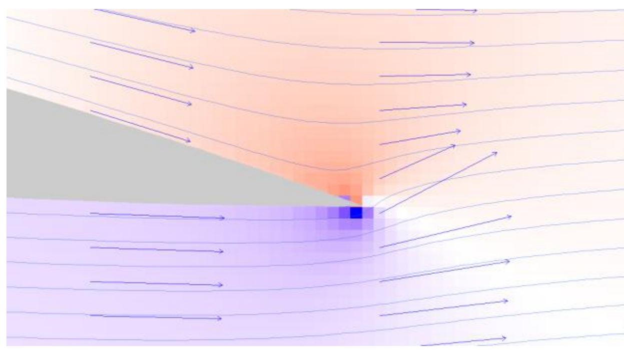

- As a result, the flow along the bottom at the airfoil separates at the trailing edge, proceeding in a generally downstream direction.
- The flow adjusts to this new configuration, resulting in a higher pressure just above the separation line than just below.
- Under the effect of the pressure difference, the separation line curves downward until the Kutta condition is met: the flow from the top and bottom of the airfoil merge smoothly at the trailing edge, with equal velocities—no more flow separation.

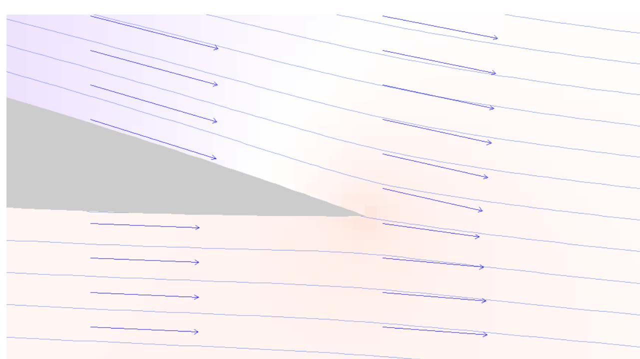

- The flow pattern that results from meeting the Kutta condition results in an increased pressure (and decreased velocity) along the bottom of the airfoil and a decreased pressure (and increased velocity) on top. The pressure difference results in lift.

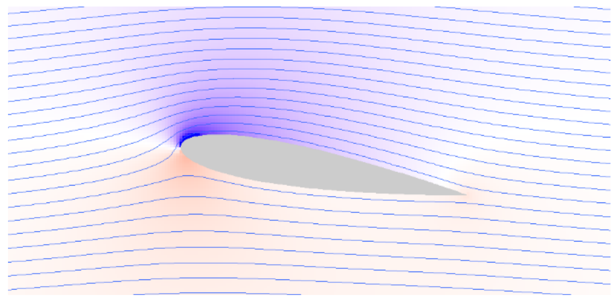

Additional Notes

1. A popular theory of lift states that air particles must travel faster over the top of the airfoil to meet at the trailing edge the particles traveling at the bottom. There is no physical reason for adjacent particles at the leading edge to meet at the trailing edge—it can actually be shown that a particles traveling at the top will arrive at the trailing edge *ahead* of the particle traveling at the bottom if the airfoil generates lift, as illustrated below.

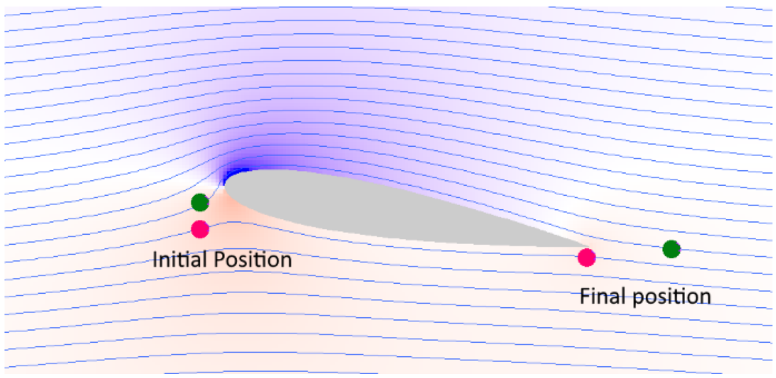

However, as we have seen, the vortex-like flow induced by the Kutta condition results in a decreased velocity along the bottom surface of the airfoil, and an increased velocity along the top (compared to the purely inviscid flow pattern, which generates no lift). By Bernoulli's law, the pressure will thus be lower at the top, and higher at the bottom, producing lift. The Kutta condition is the reason for the velocity difference between the top and bottom of the airfoil, and thus the cause of lift. And the Kutta condition itself is a consequence of air having some viscosity. No viscosity means no drag, but it means no lift either!

2. What about the case of an asymmetric airfoil? The following figure depicts the flow around an asymmetrical airfoil at zero angle of attack (a NACA 4415 airfoil, 4% camber with maximum camber at 40% chordwise).

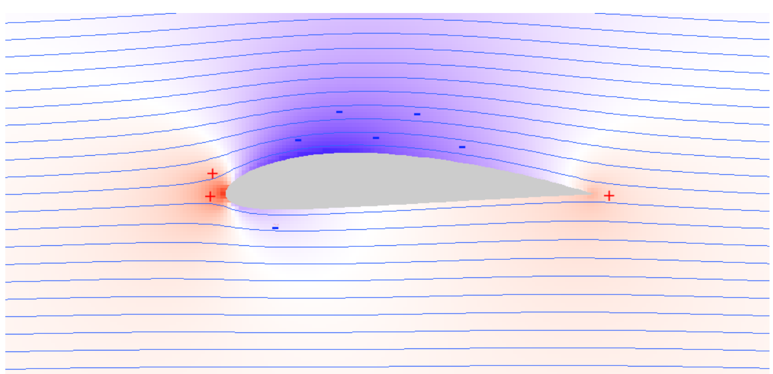

The geometry of the airfoil is such that the bottom is almost flat, while the top is more curved than the corresponding symmetric airfoil. Consequently, the streamlines (at zero angle of attack) will exhibit more curvature at the top than at the bottom. The curvature is such that both top and bottom surfaces see a lower pressure than ambient, but more so at the top (due to the higher curvature), and thus there will be a net lift generated.

3. What's the relationship between lift coefficient and angle of attack? The figure below displays the relationship between the predicted lift coefficient and angle of attack for the symmetrical airfoil (NACA 0015) and cambered one (NACA 4415).

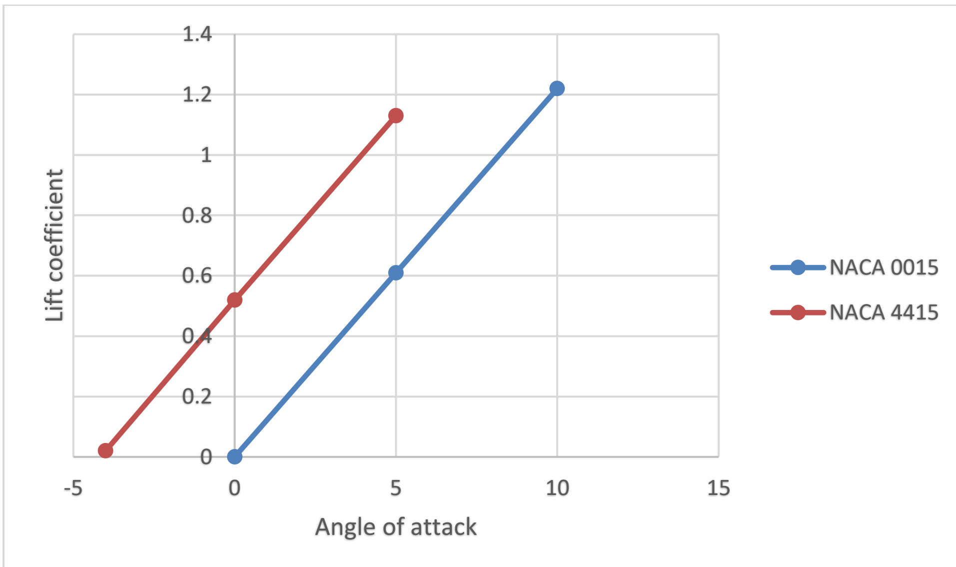

In the absence of flow separation (i.e., stall), the lift coefficient is a linear function of the angle of attack, with a slope slightly greater than 0.1 per degree. The slope is near identical for both airfoils, the main difference between them being that, as discussed above, the cambered airfoil exhibits a positive lift coefficient (0.5) at zero of attack.

4. The Kutta condition arises from fluid viscosity, and determines a flow pattern that generates lift, but our analysis did not otherwise explicitly include viscosity. As a result, the calculated drag is still zero! Viscous effects are significant only in a thin layer around the airfoil surface (typically a fraction of an inch for a typical general aviation aircraft), the so-called "boundary layer." An analysis of boundary layer flow is beyond the scope of this note, but it is the key to estimating airfoil drag, as well as the maximum lift coefficient that can be obtain before the occurrence of flow separation on the top of the airfoil—the stall. And while the exact shape of the airfoil is not a major determinant of lift (besides the effect of camber on the zero angle of attack lift coefficient), it plays a decisive role in the development of the boundary layer, and thus on the drag and stall characteristic of an airfoil. Modern airfoil design techniques typically attempt to maximize the lift-over-drag ratio, subject to a number of application-specific constraints (e.g., stall characteristics, aircraft mission profile, structural and manufacturing considerations).

5. Our analysis is based on the principle that the air follows closely the airfoil boundary. An obvious question is then—why is this true in the area of reduce pressure on the top surface of the airfoil?

Wouldn't the air elements tend to "fly off" the surface under the effect of the centrifugal force? The important thing to remember is that the pressure distributions we discussed are relative to the ambient (atmospheric) pressure. The numerical value of these pressure differences are small compared to the ambient pressure, and every part of the wing surface is subject to a large pressure directed into the surface, "pressing" the air onto it.

To visualize this we'll first display the pressure differentials around the airfoil as arrows pressing or pulling on the airfoil surface, using an arbitrary scale.

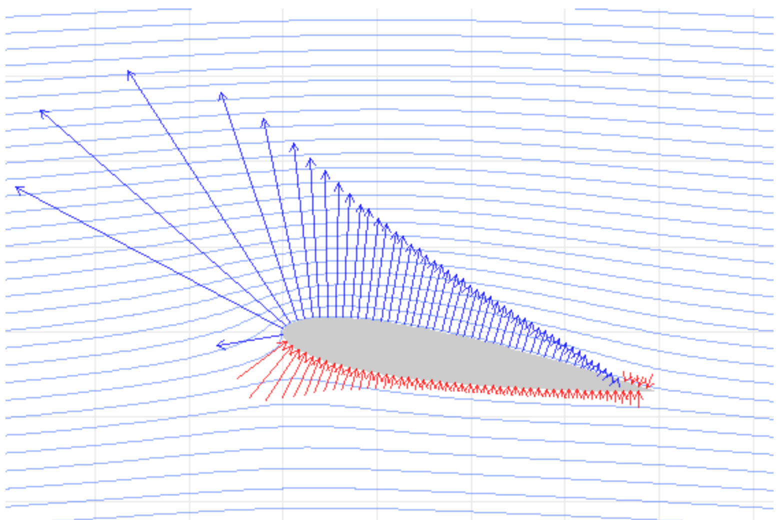

And then, here is the same diagram displaying the absolute pressure.

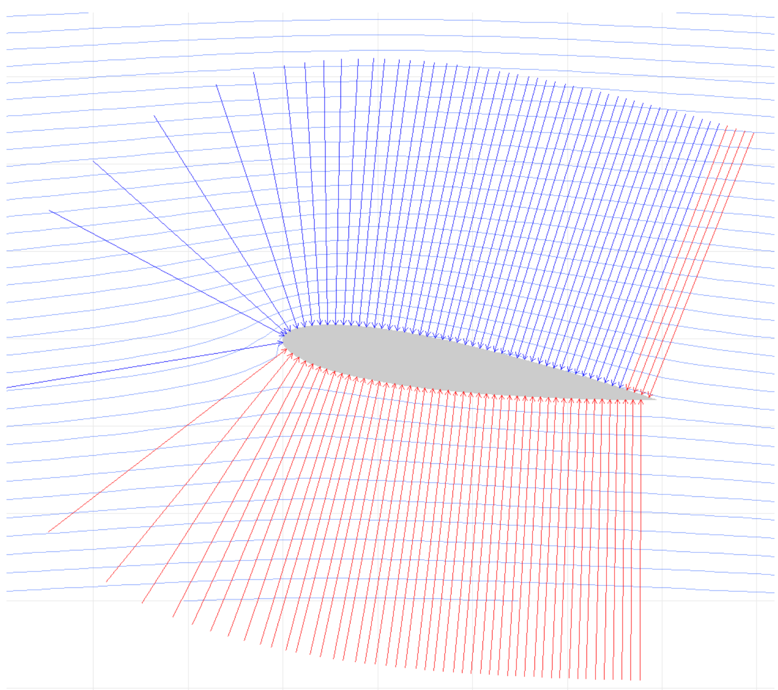

The pressure differences are barely noticeable on this graph--the maximum negative pressure difference, near the leading edge, amounts just to 10 % of atmospheric pressure, and most of the wing surface is within less than 5% of it. But these small differences are enough to keep our airplanes in the air!

Appendix A—A Deeper Dive

Below we show a 3-D view of a slice of with w of a wing of infinite span, where we identify two neighboring streamlines. Each streamline in the 2-D view corresponds in three dimensions to a stream surface.

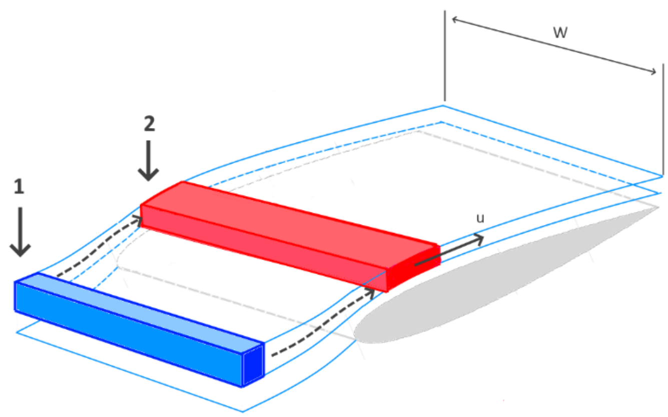

Sandwiched between these two stream surfaces is a stream tube. Air that enters the stream tube at the upstream position 1 will forever be confined to it--for example, the fluid element (a.k.a "particle") colored in blue in the figure will travel down the stream tube, and after a certain time, will be found in location 2 (colored in red).

Let's consider the red element at a specific time t, and apply to it Newton's law of motion. To simplify the figures, we'll look at the 2D side-view of the steam tube, remembering that that the width is w.

Let's first consider the laws of motion applied in a direction perpendicular to the stream tube.

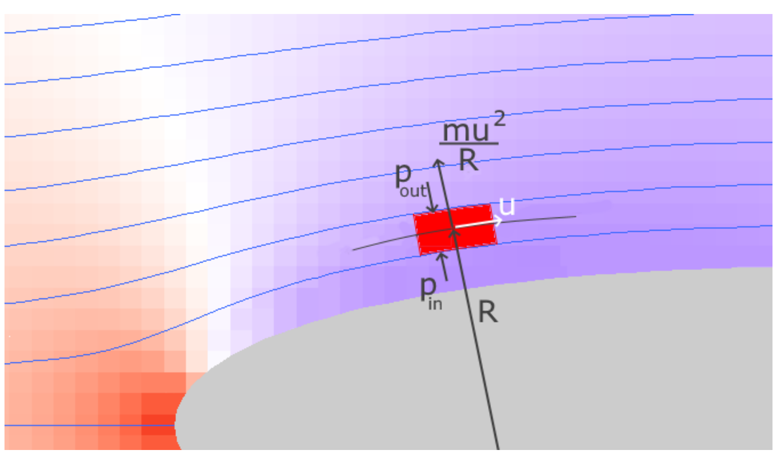

The steam tube is curved at the red particle location, with a radius of curvature R as indicated on the figure. The air particle flows with a velocity u from left (upstream) to right (downstream), and as a mass m. We'll call the side of the stream tube towards the center of curvature the "inside" side, and the opposite side the "outside" side, and will denote pressure on the inside p_{in} and the pressure on the outside p_{out}. Just like an airplane in a turn, by Newton's second law if motion, the fluid element is subjected to a "centrifugal force" trying to pull it away from the center of curvature[4]. The value of that force is $\frac{mu^2}{R}$, it is proportional to the mass of the particle, the square of the velocity, and inversely proportional to radius of curvature. Just like for an airplane in a turn, this centrifugal force must be opposed by a reaction in the opposite direction. In the case of the airplane, this is the horizontal component of lift resulting for the bank angle. In the case of our fluid element, it will be force exerted on its inner and outer sides by the fluid pressure. To balance the centrifugal force, the outside pressure must be greater than the inside pressure: $p_{out} > p_{in}$, or equivalently, $p_{in} < p_{out}$. At the location of the element under consideration, the pressure near the surface of the airfoil will be less than far away. We can visualize this by imagining that the particles in the stream tube, under the influence of the centrifugal source, are trying to "pull away" from the airfoil surface, creating a suction force—hence the deeper blue coloring near the airfoil surface (blue indicates a negative pressure relative to ambient pressure—the lesser the pressure, the more saturated the color).

The opposite is true near the leading edge. In that area, the streamline curvature is in the opposite direction—the centrifugal force will tend to "push" fluid elements towards the leading edge, resulting in a higher pressure (hence the deeper red color near the leading edge, indicating a higher relative pressure).

The general rule resulting from this discussion is that, whenever stream tubes, and thus whenever streamlines, are curved, there exists a pressure difference (technical term is a "pressure gradient") in the direction perpendicular to the streamlines. The pressure always *decreases* when traveling *towards* the center of curvature, and *increases* when going *away* from the center of curvature. Where the streamlines are parallel, the pressure is constant across them. So, from the mere inspection of the streamline pattern, we can easily deduce qualitatively where the areas of higher and lower pressure relative to ambient will be (remembering that, far from the airfoil, the pressure is constant and equal to ambient pressure).

Let's now consider the laws of motion applied along the length of the tube.

[4] Note to the purists--I apologize for introducing the "centrifugal force" here… The analysis should really be conducted in terms of centripetal acceleration and centripetal force. But I think the more intuitive concept of centrifugal force will make this section more accessible to most.

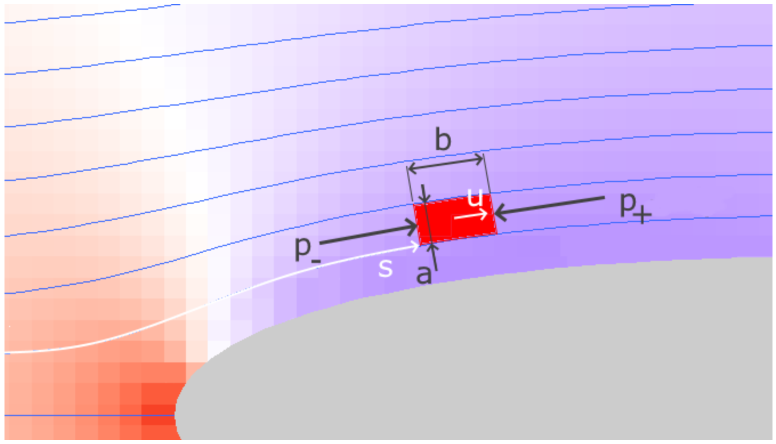

We'll call the height of the element a (the thickness of the tube at position 2), and its length b, where b is a taken as a very small length. The element is located at some distance s measured along the lower streamline. Pressure p_- acting on the upstream face exerts a force $F = awp_-$ on the element, where $a\,w$ is the surface area of the face. This force pushes the element in the downstream direction. Similarly, the pressure p_+ acting on the downstream face exerts a force awp_+ that pushes the element towards the upstream direction. If $p_+ > p_-$, the net force $aw(p_- - p_+)$ will push the element towards the upstream direction, opposing the direction of its motion, and thus slowing it down. So, qualitatively speaking, if the pressure increases along a stream tube, the fluid velocity u will decrease, and if the pressure decreases, the velocity will increase. Let's express that quantitatively (you can skip the derivation if you don't like math!). The next figure shows the same red fluid element after a short amount of time dt (the $\underline{d}$ in dt means "a small difference" in t). The dashed gray outline represents its previous position at time t. During the time interval $\mathrm{d}t$, the element will have traveled a distance $ds = u\,dt$, and the new velocity will be $u + du$.

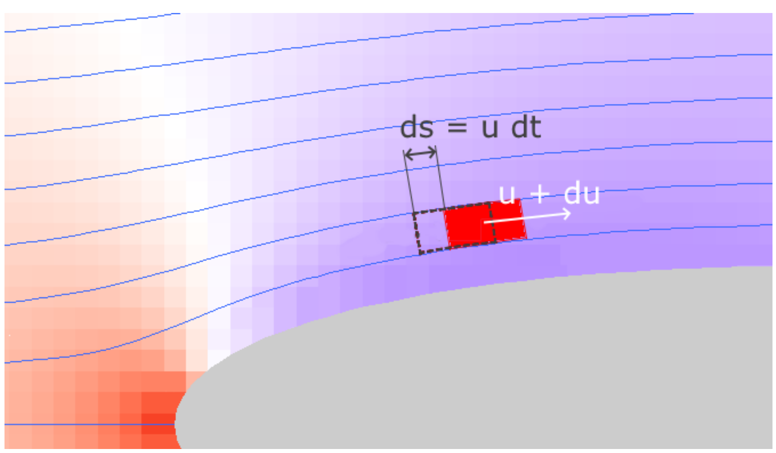

Applying Newton's second law to the red fluid element, we write:

$$Force\ = mass * acceleration$$

Where the net force, as calculated above, is $aw(p_- - p_+)$, the mass is ρabw (ρ is air density), and the acceleration is $\frac{du}{dt} = \frac{du}{ds/u} = u\frac{du}{ds}$.

Putting everything together, we have

$$aw(p_- \ - \ p_+) = \ \rho abwu\frac{du}{ds}$$

If we assume the pressure is varying linearly with distance s along the streamline (which we can do if the length b of the element is small enough), we can write

$$p_- - p_+ = b\frac{dp}{ds}$$

Replacing it in the equation gives

$$awb\frac{dp}{ds} = \ \rho abwu\frac{du}{ds}$$

And finally, simplifying all the common factors on the left and right,

$$dp = \ \rho u\ du$$

This tells us how a small pressure change will be related to a small change in velocity. But we'd like to know how pressure and velocity relate to ach other, independently of the size of the changes… That's where calculus comes to the rescue. We can't go into the details here—just let it be known that we can "integrate" the previous equation between two arbitrary position 1 and 2 along the streamline, and write

$$\int_1^2 dp = \rho\int_1^2 u\ du$$

($\int_1^2 dp$ means "take the sum of all dps between position 1 and 2")

And the rules of calculus tell us that this computes to

$$p_2 - p_1 = \frac{1}{2}\rho(u_2^2 - u_1^2)$$

Or, rearranging the terms,

$$p_1 + \ \tfrac{1}{2}\rho u_1^2 = \ p_2 + \ \tfrac{1}{2}\rho u_2^2$$

Or, since 1 and 2 are arbitrary positions, we can write that $p + \frac{1}{2}\rho u^2$ is constant along a stream tube. And since far upstream the flow is uniform, with constant pressure and velocity everywhere, we have that $p + \frac{1}{2}\rho u^2$ is constant everywhere—this is nothing else than Bernoulli's law!

So, by applying Newton's law of motion to a small particle of air along a stream tube, we actually proved Bernoulli's law. And, by combining Bernoulli's law with the pressure gradient/streamline curvature relationship we discussed previously, based on a streamline pattern, we will also be able to identify the regions of higher and lower velocities.