Airfoil Design

Outline of this Chapter

The chapter is divided into several sections. The first of these consist of an introduction to airfoils: some history and basic ideas. The latter sections deal with simple results that relate the airfoil geometry to its basic aerodynamic characteristics. The latter sections deal with the process of airfoil design.

History of Airfoil Development

The earliest serious work on the development of airfoil sections began in the late 1800's. Although it was known that flat plates would produce lift when set at an angle of incidence, some suspected that shapes with curvature, that more closely resembled bird wings would produce more lift or do so more efficiently. H.F. Phillips patented a series of airfoil shapes in 1884 after testing them in one of the earliest wind tunnels in which "artificial currents of air (were) produced from induction by a steam jet in a wooden trunk or conduit." Octave Chanute writes in 1893, "...it seems very desirable that further scientific experiments be be made on concavo-convex surfaces of varying shapes, for it is not impossible that the difference between success and failure of a proposed flying machine will depend upon the sustaining effect between a plane surface and one properly curved to get a maximum of 'lift'."

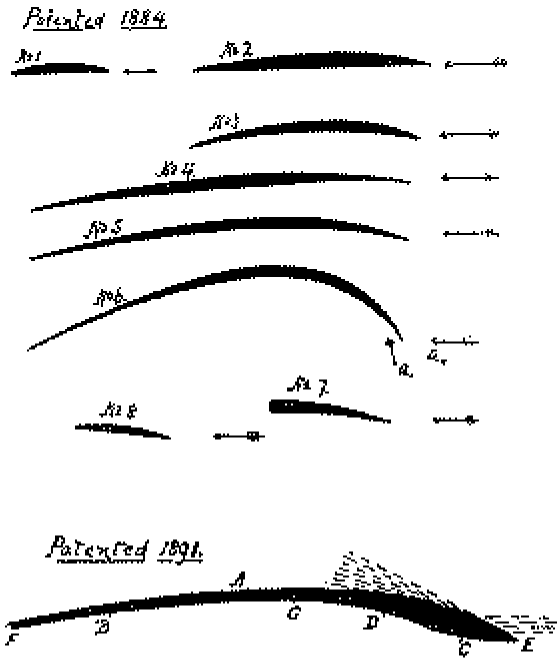

At nearly the same time Otto Lilienthal had similar ideas. After carefully measuring the shapes of bird wings, he tested the airfoils below (reproduced from his 1894 book, "Bird Flight as the Basis of Aviation") on a 7m diameter "whirling machine". Lilienthal believed that the key to successful flight was wing curvature or camber. He also experimented with different nose radii and thickness distributions.

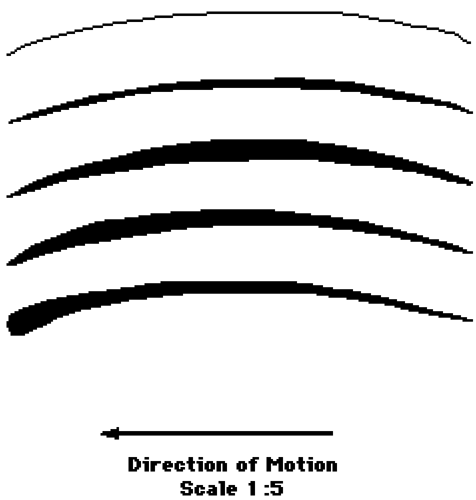

Airfoils used by the Wright Brothers closely resembled Lilienthal's sections: thin and highly cambered. This was quite possibly because early tests of airfoil sections were done at extremely low Reynolds number, where such sections behave much better than thicker ones. The erroneous belief that efficient airfoils had to be thin and highly cambered was one reason that some of the first airplanes were biplanes.

The use of such sections gradually diminished over the next decade.

A wide range of airfoils were developed, based primarily on trial and error. Some of the more successful sections such as the Clark Y and Gottingen 398 were used as the basis for a family of sections tested by the NACA in the early 1920's.

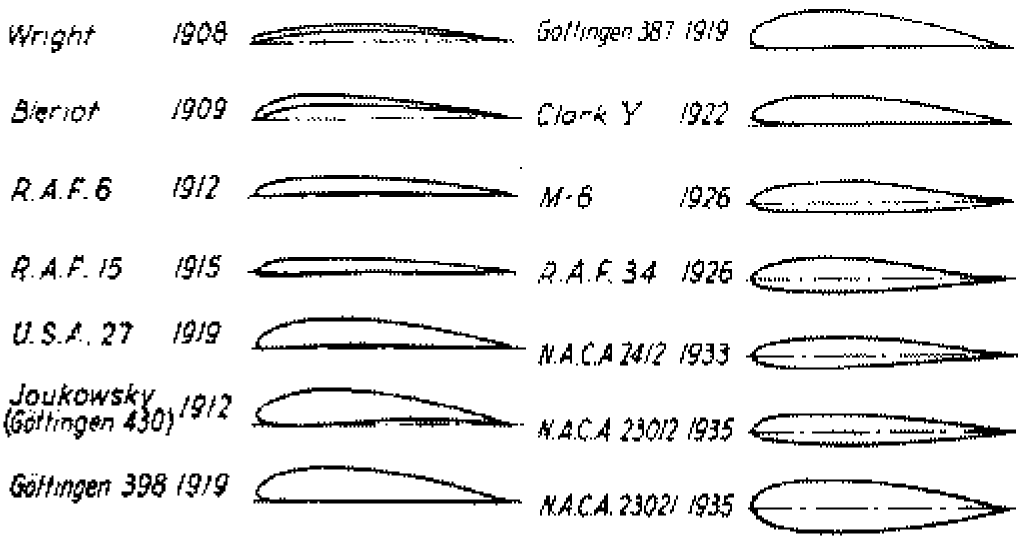

In 1939, Eastman Jacobs at the NACA in Langley, designed and tested the first laminar flow airfoil sections. These shapes had extremely low drag and the section shown here achieved a lift to drag ratio of about 300.

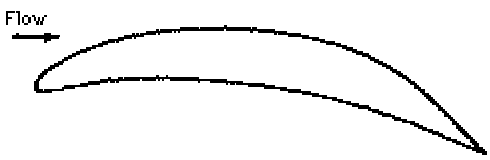

A modern laminar flow section, used on sailplanes, illustrates that the concept is practical for some applications. It was not thought to be practical for many years after Jacobs demonstrated it in the wind tunnel. Even now, the utility of the concept is not wholly accepted and the "Laminar Flow True-Believers Club" meets each year at the homebuilt aircraft fly-in.

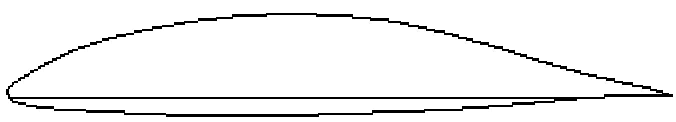

One of the reasons that modern airfoils look quite different from one another and designers have not settled on the one best airfoil is that the flow conditions and design goals change from one application to the next. On the right are some airfoils designed for low Reynolds numbers.

At very low Reynolds numbers (<10,000 based on chord length) efficient airfoil sections can look rather peculiar as suggested by the sketch of a dragonfly wing. The thin, highly cambered pigeon wing is similar to Lilienthal's designs. The Eppler 193 is a good section for model airplanes. The Lissaman 7769 was designed for human-powered aircraft.

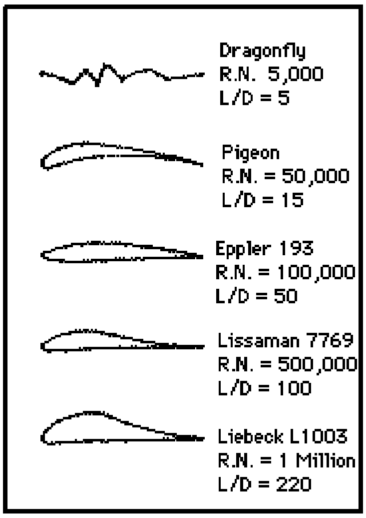

Unusual airfoil design constraints can sometimes arise, leading to some unconventional shapes. The airfoil here was designed for an ultralight sailplane requiring very high maximum lift coefficients with small pitching moments at high speed. One possible solution: a variable geometry airfoil with flexible lower surface.

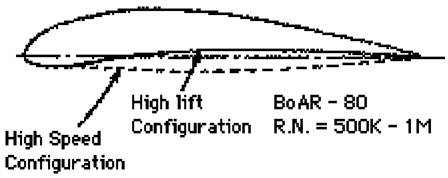

The airfoil used on the Solar Challenger, an aircraft that flew across the English Channel on solar power, was designed with an totally flat upper surface so that solar cells could be easily mounted.

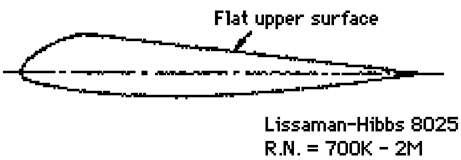

The wide range of operating conditions and constraints, generally makes the use of an existing, "catalog" section, not best. These days airfoils are usually designed especially for their intended application. The remaining parts of this chapter describe the basic ideas behind how this is done.

Airfoil Geometry

Airfoil geometry can be characterized by the coordinates of the upper and lower surface. It is often summarized by a few parameters such as: maximum thickness, maximum camber, position of max thickness, position of max camber, and nose radius. One can generate a reasonable airfoil section given these parameters. This was done by Eastman Jacobs in the early 1930's to create a family of airfoils known as the NACA Sections.

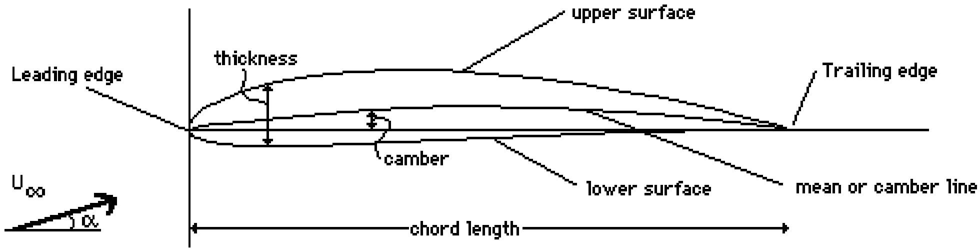

The NACA 4 digit and 5 digit airfoils were created by superimposing a simple meanline shape with a thickness distribution that was obtained by fitting a couple of popular airfoils of the time:

$$\pm y = (t/0.2) * (.2969*x^{0.5} - .126*x - .3537*x^2 + .2843*x^3 - .1015*x^4)$$

The camberline of 4-digit sections was defined as a parabola from the leading edge to the position of maximum camber, then another parabola back to the trailing edge.

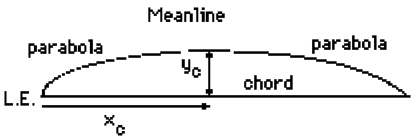

NACA 4-Digit Series:

4	4	1 2
max camber in % chord	position of max camber in 1/10 of c	max thickness in % of chord

After the 4-digit sections came the 5-digit sections such as the famous NACA 23012. These sections had the same thickness distribution, but used a camberline with more curvature near the nose. A cubic was faired into a straight line for the 5-digit sections.

NACA 5-Digit Series:

2	3 0	1 2
approx max	position	max thickness

camber	of max camber	in % of chord
in % chord	in 2/100 of c	

The 6-series of NACA airfoils departed from this simply-defined family. These sections were generated from a more or less prescribed pressure distribution and were meant to achieve some laminar flow.

NACA 6-Digit Series:

6	3,	2	-	2	1	2
Six-	location	half width		ideal Cl	max thickness	
Series	of min Cp	of low drag		in tenths	in % of chord	
	in 1/10 chord	bucket in 1/10 of Cl				

After the six-series sections, airfoil design became much more specialized for the particular application. Airfoils with good transonic performance, good maximum lift capability, very thick sections, very low drag sections are now designed for each use. Often a wing design begins with the definition of several airfoil sections and then the entire geometry is modified based on its 3-dimensional characteristics.

Airfoil Pressure Distributions

The aerodynamic performance of airfoil sections can be studied most easily by reference to the distribution of pressure over the airfoil. This distribution is usually expressed in terms of the pressure coefficient:

$$C_p = \frac{p - p_\infty}{\frac{1}{2} \rho U_\infty^2}$$

Cp is the difference between local static pressure and freestream static pressure, non-dimensionalized by the freestream dynamic pressure.

What does an airfoil pressure distribution look like? We generally plot C_p vs. x/c.

x/c varies from 0 at the leading edge to 1.0 at the trailing edge. C_p is plotted "upside-down" with negative values (suction), higher on the plot. (This is done so that the upper surface of a conventional lifting airfoil corresponds to the upper curve.)

The C_p starts from about 1.0 at the stagnation point near the leading edge...

It rises rapidly (pressure decreases) on both the upper and lower surfaces...

...and finally recovers to a small positive value of C_p near the trailing edge.

Various parts of the pressure distribution are depicted in the figure below and are described in the following sections.

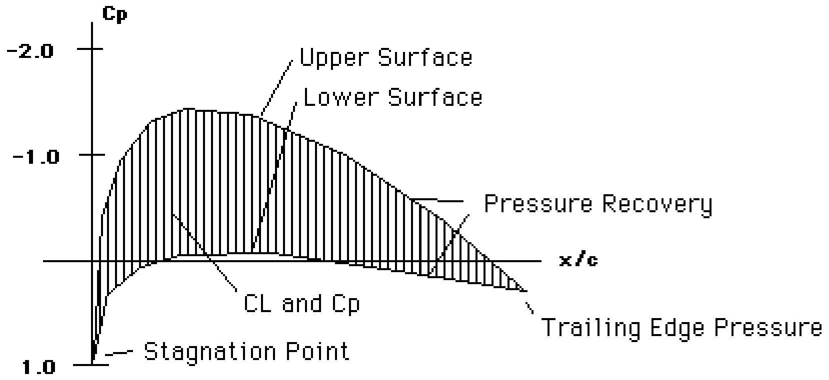

- Upper Surface
  The upper surface pressure is lower (plotted higher on the usual scale) than the lower surface Cp in this case. But it doesn't have to be.

- Lower Surface
  The lower surface sometimes carries a positive pressure, but at many design conditions is actually pulling the wing downward. In this case, some suction (negative Cp -> downward force on lower surface) is present near the mid-chord.

- Pressure Recovery
  This region of the pressure distribution is called the pressure recovery region.
  The pressure increases from its minimum value to the value at the trailing edge.
  This area is also known as the region of adverse pressure gradient. As discussed in other sections, the adverse pressure gradient is associated with boundary layer transition and possibly separation, if the gradient is too severe.

- Trailing Edge Pressure
  The pressure at the trailing edge is related to the airfoil thickness and shape near the trailing edge. For thick airfoils the pressure here is slightly positive (the velocity is a bit less than the freestream velocity). For infinitely thin sections $C_p = 0$ at the trailing edge. Large positive values of C_p at the trailing edge imply more severe adverse pressure gradients.

- C_L and C_p
  The section lift coefficient is related to the C_p by: $C_l = \text{int}\ (C_{pl} - C_{pu})\ dx/c$
  (It is the area between the curves.)
  with C_{pu} = upper surface C_p
  and recall C_l = section lift / (q c)

- Stagnation Point
  The stagnation point occurs near the leading edge. It is the place at which V = 0. Note that in

incompressible flow C_p = 1.0 at this point. In compressible flow it may be somewhat larger.

We can get a more intuitive picture of the pressure distribution by looking at some examples and this is done in some of the following sections in this chapter.

Airfoil Pressures and Performance

The shape of the pressure distribution is directly related to the airfoil performance as indicated by some of the features shown in the figure below.

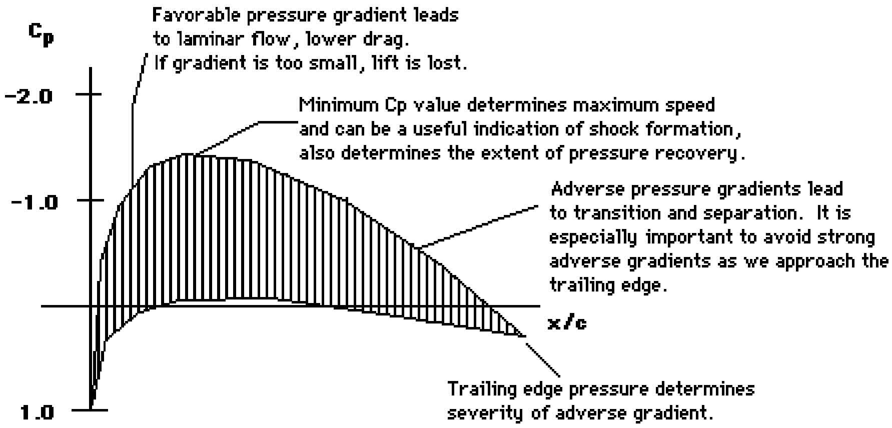

Most of these considerations are related to the airfoil boundary layer characteristics which we will take up later, but even in the inviscid case we can draw some conclusions. We may compute the maximum local Mach numbers and relate those to lift and thickness; we can compute the pitching moment and decide if that is acceptable.

Whether we use the inviscid pressures to form qualitative conclusions about the section, or use them as input to a more detailed boundary layer calculation, we must first investigate the close relation between the airfoil geometry to these pressures.

Relating Airfoil Geometry and Pressures

The relationship between airfoil geometry and airfoil pressure distributions can be predicted numerically solving the relevant field equations. But it can also be understood in a ratrher intuitive way.

We first look at the effect of changes in surface curvature.

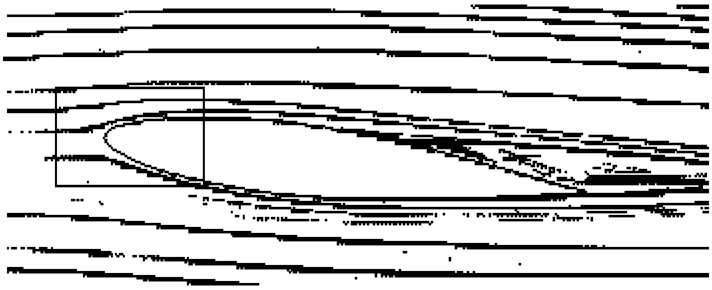

An intuitive view of the C_p-curvature relation

For equilibrium we must have a pressure gradient when the flow is curved.

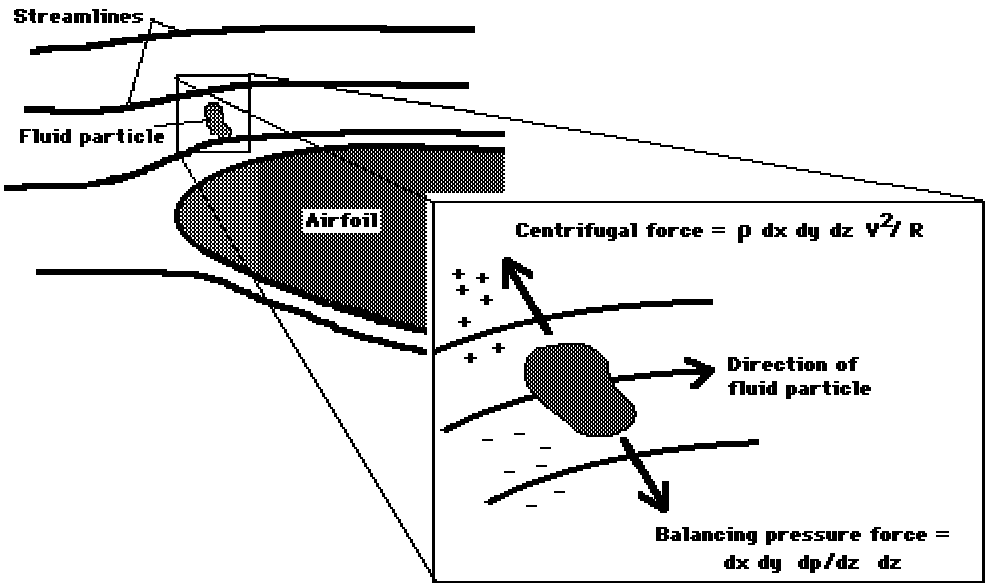

In the case shown here, the pressure must increase as we move further from the surface. This means that the surface pressure is lower than the pressures farther away. This is why the C_p is more negative in regions with curvature in this direction. The curvature of the streamlines determines the pressures and hence the net lift.

The figure below shows how the airfoil pressures vary with angle of attack. Note that the "nose peak" becomes more extreme as the angle increases.

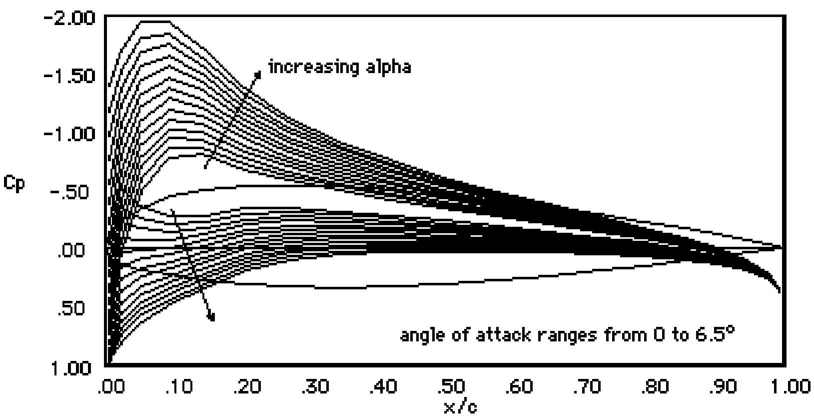

Let's consider, in more detail the relationship between airfoil geometry and airfoil pressure distributions. The next few examples show some of the effects of changes in camber, leading edge radius, trailing edge angle, and local distortions in the airfoil surface.

A reflexed airfoil section has reduced camber over the aft section producing less lift over this region. and therefore less nose-down pitching moment. In this case the aft section is actually pushing downward and C_m at zero lift is positive.

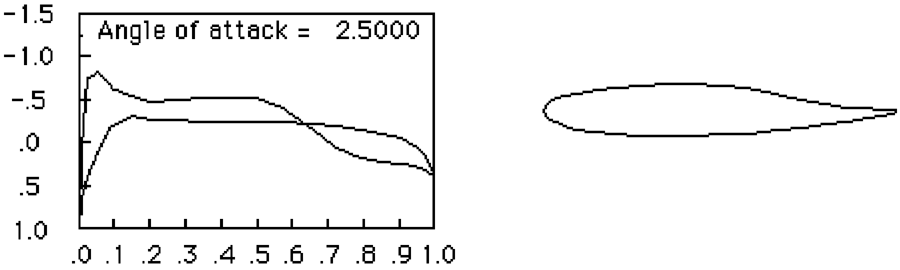

A natural laminar flow section has a thickness distribution that leads to a favorable pressure gradient over a portion of the airfoil. In this case, the rather sharp nose leads to favorable gradients over 50% of the section.

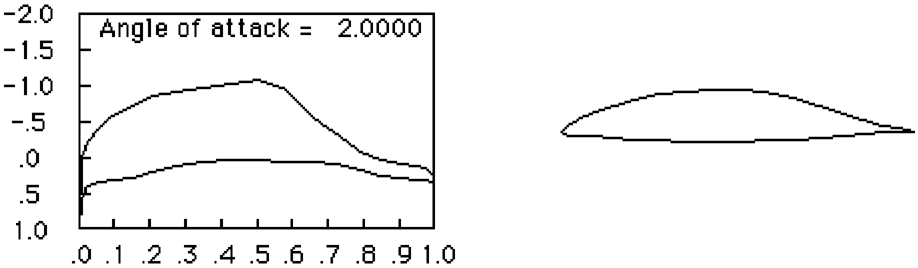

This is a symmetrical section at 4° angle of attack.
Note the pressure peak near the nose. A thicker section would have a less prominent peak.

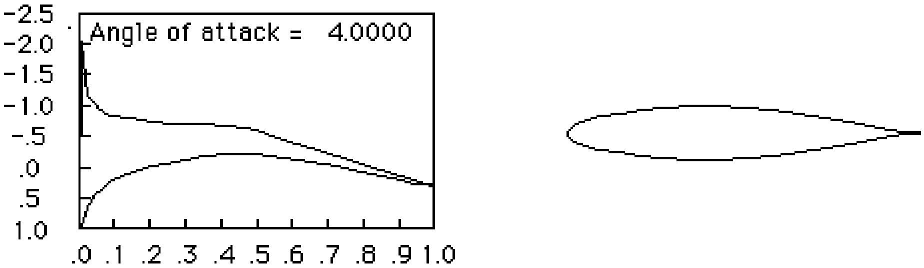

Here is a thicker section at 0°. Only one line is shown on the plot because at zero lift, the upper and lower surface pressure coincide.

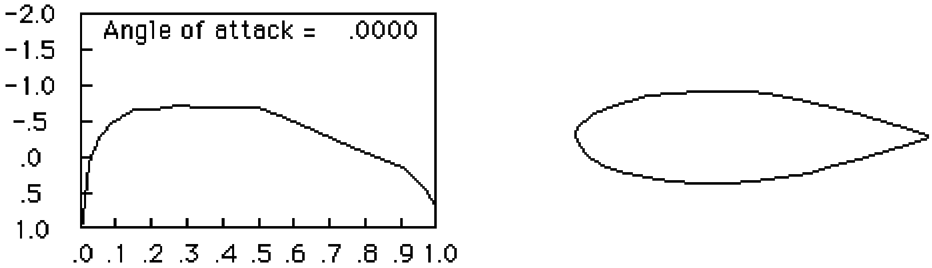

A conventional cambered section.

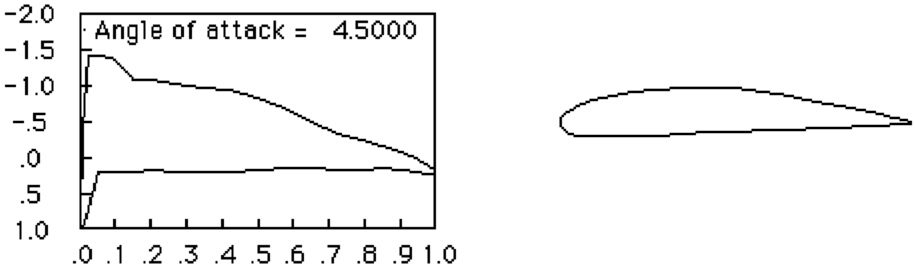

An aft-loaded section, the opposite of a reflexed airfoil carries more lift over the aft part of the airfoil. Supercritical airfoil sections look a bit like this.

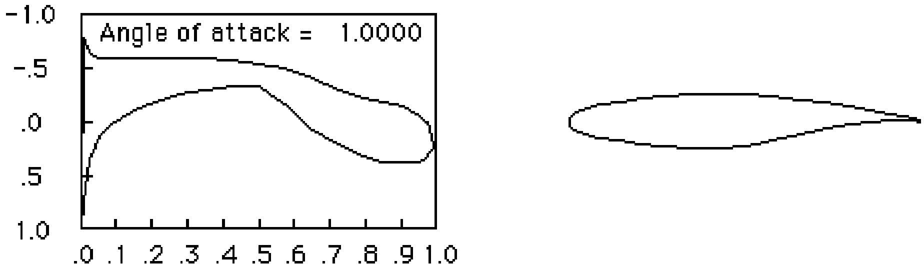

Airfoil Design Methods

The process of airfoil design proceeds from a knowledge of the boundary layer properties and the relation between geometry and pressure distribution. The goal of an airfoil design varies. Some airfoils are designed to produce low drag (and may not be required to generate lift at all.) Some sections may need to produce low drag while producing a given amount of lift. In some cases, the drag doesn't really matter - it is maximum lift that is important. The section may be required to achieve this performance with a constraint on thickness, or pitching moment, or off-design performance, or other unusual constraints. Some of these are discussed further in the section on previous section of historical examples.

One approach to airfoil design is to use an airfoil that was already designed by someone who knew what he or she was doing. This "design by authority" works well when the goals of a particular design problem happen to coincide with the goals of the original airfoil design. This is rarely the case, although sometimes existing airfoils are good enough. In these cases, airfoils may be chosen from catalogs such as Abbott and von Doenhoff's Theory of Wing Sections, Althaus' and Wortmann's Stuttgarter Profilkatalog, Althaus' Low Reynolds Number Airfoil catalog, or Selig's "Airfoils at Low Speeds".

The advantage to this approach is that there is test data available. No surprises, such as a unexpected early stall, are likely. On the other hand, available tools are now sufficiently refined that one can be reasonably sure that the predicted performance can be achieved. The use of "designer airfoils" specifically tailored to the needs of a given project is now very common. This section of the notes deals with the process of custom airfoil design.

Methods for airfoil design can be classified into two categories: direct and inverse design.

Direct Methods for Airfoil Design

The direct airfoil design methods involve the specification of a section geometry and the calculation of pressures and performance. One evaluates the given shape and then modifies the shape to improve the performance.

The two main subproblems in this type of method are

1. the identification of the measure of performance
2. the approach to changing the shape so that the performance is improved

The simplest form of direct airfoil design involves starting with an assumed airfoil shape (such as a NACA airfoil), determining the characteristic of this section that is most problemsome, and fixing this problem. This process of fixing the most obvious problems with a given airfoil is repeated until there is no major problem with the section. The design of such airfoil does not require a specific definition of a scalar objective function, but it does require some expertise to identify the potential problems and often considerable expertise to fix them. Let's look at a simple (but real life!) example.
A company is in the business of building rigid wing hang gliders and because of the low speed requirements, they decide to use a version of one of Bob Liebeck's very high lift airfoils. Here is the pressure distribution at a lift coefficient of 1.4. Note that only a small amount of trailing edge separation is predicted. Actually, the airfoil works quite well, achieving a C_{lmax} of almost 1.9 at a Reynolds number of one million.

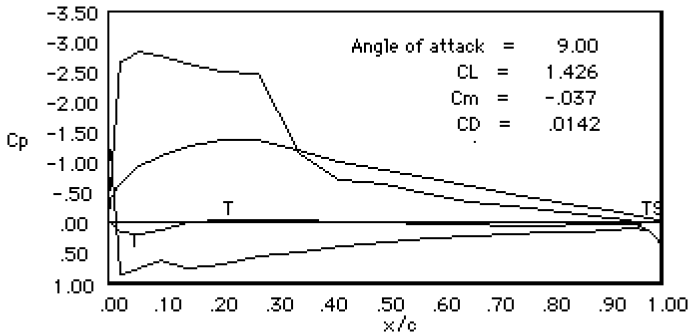

This glider was actually built and flown. It, in fact, won the 1989 U.S. National Championships. But it had terrible high speed performance. At lower lift coefficients the wing seemed to fall out of the sky. The plot below shows the pressure distribution at a C_l of 0.6. The pressure peak on the lower surface causes separation and severely limits the maximum speed. This is not too hard to fix.

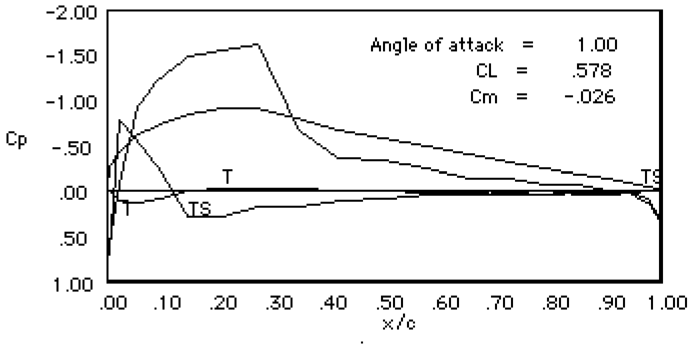

By reducing the lower surface "bump" near the leading edge and increasing the lower surface thickness aft of the bump, the pressure peak at low C_l is easily removed. The lower surface flow is now attached, and remains attached down to a C_l of about 0.2. We must check to see that we have not hurt the C_{lmax} too much.

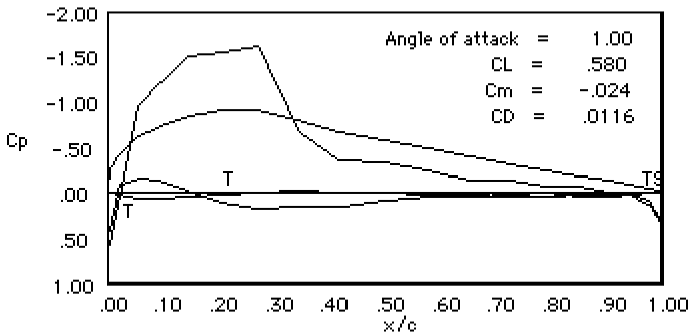

Here is the new section at the original design condition (still less than C_{lmax}). The modification of the lower surface has not done much to the upper surface pressure peak here and the C_{lmax} turns out to be changed very little. This section is a much better match for the application and demonstrates how effective small modifications to existing sections can be. The new version of the glider did not use this section, but one that was designed from scratch with lower drag.

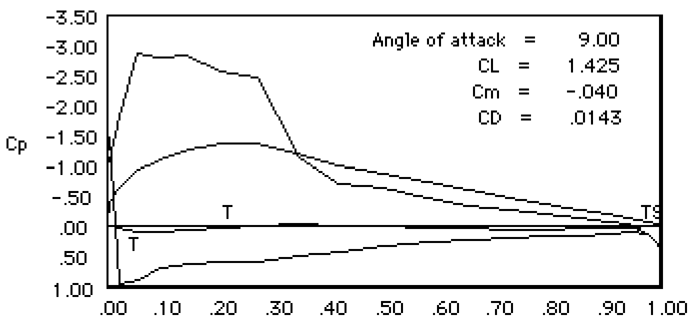

Sometimes the objective of airfoil design can be stated more positively than, "fix the worst things". We might try to reduce the drag at high speeds while trying to keep the maximum C_L greater than a certain value. This could involve slowly increasing the amount of laminar flow at low C_l's and checking to see the effect on the maximum lift. The objective may be defined numerically. We could actually minimize Cd with a constraint on C_{lmax}. We could maximize L/D or $C_l^{1.5}/C_d$ or C_{lmax} / C_d@$C_{ldesign}$. The selection of the figure of merit for airfoil sections is quite important and generally cannot be done without considering the rest of the airplane. For example, if we wish to build an airplane with maximum L/D we do not build a section with maximum L/D because the section C_l for best C_l/C_d is different from the airplane C_L for best C_L/C_D.

Inverse Design

Another type of objective function is the target pressure distribution. It is sometimes possible to specify a desired C_p distribution and use the least squares difference between the actual and target C_p's as the objective. This is the basic idea behind a variety of methods for inverse design. As an example, thin airfoil theory can be used to solve for the shape of the camberline that produces a specified pressure difference on an airfoil in potential flow.

The second part of the design problem starts when one has somehow defined an objective for the airfoil design. This stage of the design involves changing the airfoil shape to improve the performance. This may be done in several ways:
1. By hand, using knowledge of the effects of geometry changes on C_p and C_p changes on performance.
2. By numerical optimization, using shape functions to represent the airfoil geometry and letting the computer decide on the sequence of modifications needed to improve the design.

Typical Airfoil Design Problems

Regardless of the design goals and constraints, one is faced with some common problems that make airfoil design difficult. This section deals with the common issues that arise in the following design problems:

Thick Airfoil Design

The difficulty with thick airfoils is that the minimum pressure is decreased due to thickness. This results in a more severe adverse pressure gradient and the need to start recovery sooner. If the maximum thickness point is specified, the section with maximum thickness must recover from a given point with the steepest possible gradient. This is just the sort of problem addressed by Liebeck in connection with maximum lift. The thickest possible section has a boundary layer just on the verge of separation throughout the recovery.

The thickest section at Re = 10 million is 57% thick, but of course, it will separate suddenly with any angle of attack.

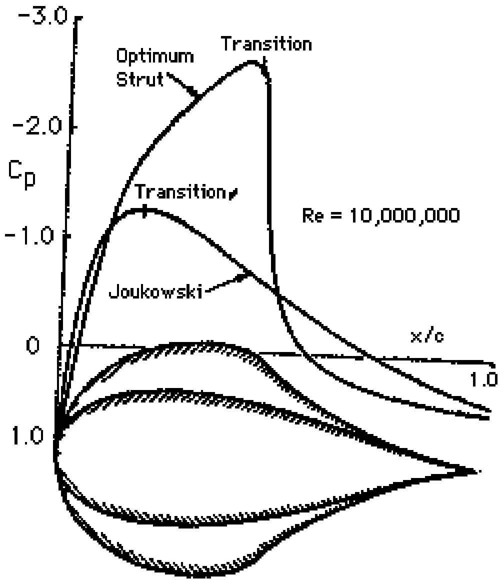

High Lift Airfoil Design

To produce high lift coefficients, we require very negative pressures on the upper surface of the airfoil. The limit to this suction may be associated with compressibility effects, or may be imposed by the requirement that the boundary layer be capable of negotiating the resulting adverse pressure recovery. It may be shown that to maximize lift starting from a specified recovery height and location, it is best to keep the boundary layer on the verge of separation*. Such distributions are shown below for a Re of 5 million. Note the difference between laminar and turbulent results.

The thickest section at Re = 10 million is 57% thick, but of course, it will separate suddenly with any angle of attack.

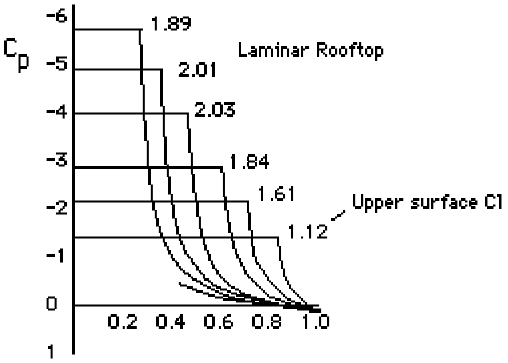

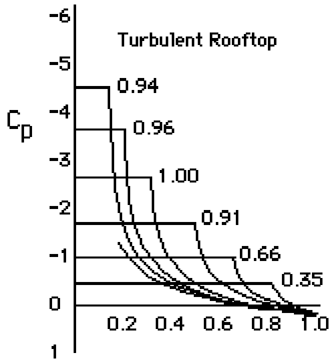

For maximum airfoil lift, the best recovery location is chosen and the airfoil is made very thin so that the lower surface produces maximum lift as well. (Since the upper surface Cp is specified, increasing thickness only reduces the lower surface pressures.)

Well, almost. If the upper surface Cp is more negative than -3.0, the perturbation velocity is greater than freestream, which means, for a thin section, the lower surface flow is upstream. This would cause separation and the maximum lift is achieved with an upper surface velocity just over 2U and a bit of thickness to keep the lower surface near stagnation pressure.

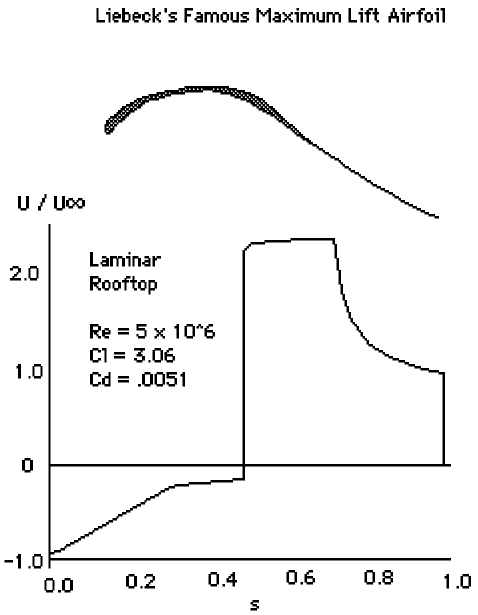

Laminar Airfoil Design

Laminar flow may be useful for reducing skin friction drag, increasing maximum lift, or reducing heat transfer. It may be achieved without too much work at low Reynolds numbers by maintaining a smooth surface and using an airfoil with a favorable pressure gradient. The section below shows how the pressures may be tailored to achieve long runs of laminar flow on upper and/or lower surfaces.

Again, the Stratford-like pressure recovery is helpful in achieving the maximum run of favorable gradient on either upper or lower surfaces.

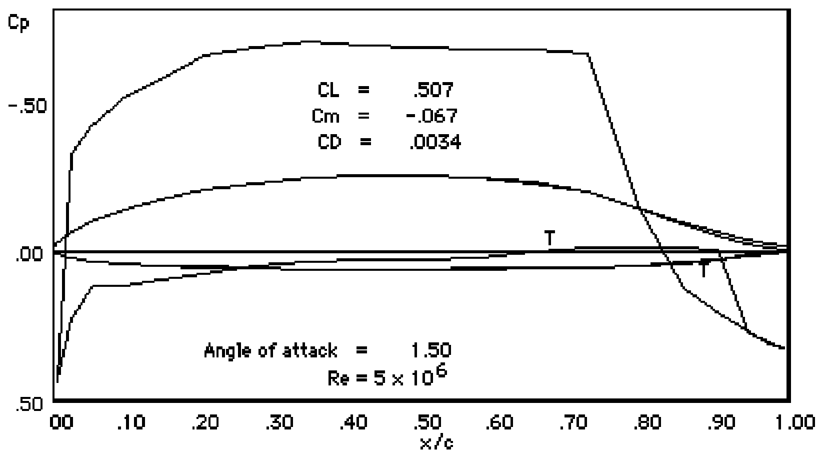

Transonic Airfoil Design

The transonic airfoil design problem arises because we wish to limit shock drag losses at a given transonic speed. This effectively limits the minimum pressure coefficient that can be tolerated. Since both lift and thickness reduce (increase in magnitude) the minimum C_p, the transonic design problem is to create an airfoil section with high lift and/or thickness without causing strong shock waves. One can generally tolerate some supersonic flow without drag increase, so that most sections can operate efficiently as "supercritical airfoils". A rule of thumb is that the maximum local Mach numbers should not exceed about 1.2 to 1.3 on a well-designed supercritical airfoil. This produces a considerable increase in available C_l compared with entirely subcritical designs.

Supercritical sections usually refer to a special type of airfoil that is designed to operate efficiently with substantial regions of supersonic flow. Such sections often take advantage of many of the following design ideas to maximize lift or thickness at a given Mach number:

- Carry as much lift as is practical on the aft potion of the section where the flow is subsonic. The aft lower surface is an obvious candidate for increased loading (more positive C_p), although several considerations discussed below limit the extent to which this approach can be used.
- Make sure that sufficient lift is carried on the forward portion of the upper surface. As the Mach number increases, the pressure peak near the nose is diminished and without additional blunting of the nose, possible extra lift will be lost in this region.
- The lower surface near the nose can also be loaded by reducing the lower surface thickness near the leading edge. This provides both lift and positive pitching moment.
- Shocks on the upper surface near the leading edge produce much less wave drag than shocks aft of the airfoil crest and it is feasible, although not always best, to design sections with forward shocks. Such sections are known as "peaky" airfoils and were used on many transport aircraft.
- The idea of carefully tailoring the section to obtain locally supersonic flow without shockwaves (shock-free sections) has been pursued for many years, and such sections have been designed and tested. For most practical cases with a range of design CL and Mach number, sections with weak shocks are favored.

One must be cautious with supercritical airfoil design. Several of these sections have looked promising initially, but led to problems when actually incorporated into an aircraft design. Typical difficulties include the following.

- Too much aft loading can produce large negative pitching moments with trim drag and structural weight penalties.
- The adverse pressure gradient on the aft lower surface can produce separation in extreme cases.
- The thin trailing edge may be difficult to manufacture.
- Supercritical, and especially shock-free designs often are very sensitive to Mach and C_L and may

perform poorly at off-design conditions. The appearance of "drag creep" is quite common, a situation in which substantial section drag increase with Mach number occurs even at speeds below the design value.

The section with pressures shown below is typical of a modern supercritical section with a weak shock at its design condition. Note the rooftop C_p design with the minimum C_p considerably greater above C_p*.

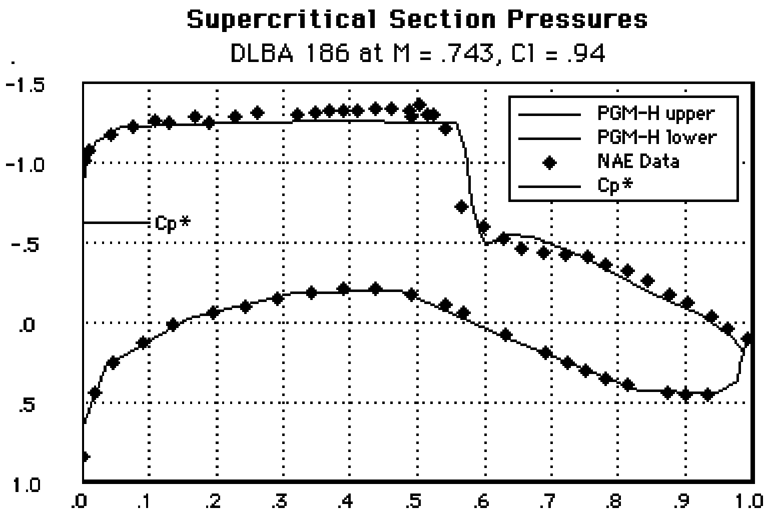

Low Reynolds Number Airfoil Design

Low Reynolds numbers make the problem of airfoil design difficult because the boundary layer is much less capable of handling an adverse pressure gradient without separation. Thus, very low Reynolds number designs do not have severe pressure gradients and the maximum lift capability is restricted.

Low Reynolds number airfoil designs are cursed with the problem of too much laminar flow. It is sometimes difficult to assure that the boundary layer is turbulent over the steepest pressure recovery regions. Laminar separation bubbles are common and unless properly stabilized can lead to excessive drag and low maximum lift.

At very low Reynolds numbers, most or all of the boundary layer is laminar. Under such conditions the boundary layer can handle only gradual pressure recovery. Based on the expressions for laminar separation, one finds that an all-laminar section can generate a C_L of about 0.4 or achieve a thickness of about 7.5%.

Low Moment Airfoil Design

When the airfoil pitching moment is constrained, it is not always possible to carry lift as far back on the airfoil as desired. Such situations arise in the design of sections for tailless aircraft, helicopter rotor blades, and even sails, kites, and giant pterosaurs. The airfoil shown here is a Liebeck section designed to perform well at low Reynolds numbers with a positive C_{m0}. Its performance is not bad, but it is clearly inferior in C_{lmax} when compared to other sections without a C_m constraint. (C_{lmax} = 1.35 vs. 1.60 for conventional sections at Re = 500,000.)

The thickest section at Re = 10 million is 57% thick, but of course, it will separate suddenly with any angle of attack.

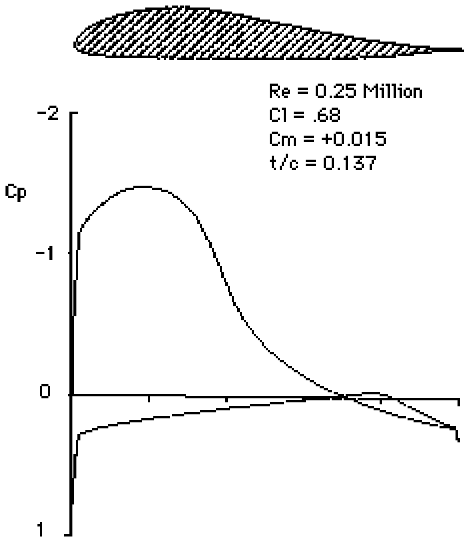

Multiple Design Point Airfoils

One of the difficulties in designing a good airfoil is the requirement for acceptable off-design performance. While a very low drag section is not too hard to design, it may separate at angles of attack slightly away from its design point. Airfoils with high lift capability may perform very poorly at lower angles of attack.

One can approach the design of airfoil sections with multiple design points in a well-defined way. Often it is clear that the upper surface will be critical at one of the points and we can design the upper surface at this condition. The lower surface can then be designed to make the section behave properly at the second point. Similarly, constraints such as Cmo are most effected by airfoil trailing edge geometry.

When such a compromise is not possible, variable geometry can be employed (at some expense) as in the case of high lift systems.