How do sails work?

Article by Paul Bogataj

Sails are wings that use the wind to generate a force to move a boat. The following explanation of how this occurs can help understand how to maximize the performance achieved from sails.

Sails are Flexible Wings

It is useful to recognize what a typical sail is. They are normally built from a flexible material in order to allow the sail to work with the wind on either side to allow tacking. This is a significant restriction that prevents many shapes from being built because they would not be able to support themselves in the wind. This leads to the traditional triangular planform of sails, since the material below has to hang from the material above, which eventually is reduced to a point at the top of the mast. So, the problem becomes how to build and operate a flexible sail in the wind to produce a substantial force component to move the boat.

As the restriction that sails support themselves is diminished (full battens and stiffer materials for example), sails can evolve to be more efficient. Their appearance then becomes more wing-like and less sail-like. Analyzing how a sail works as a wing is useful, not just for modern sails that look more like wings, but also for very traditional sails that, while they look like sails, operate very much like wings.

Velocity and Pressure

Flow accelerates over the top surface of an airfoil, either because it is at an angle to the flow, or because the top has more curvature than the bottom, or both. When a fluid (like air or water) is accelerated, the pressure that it imparts on an adjoining surface decreases. This lower pressure pulling upward on the upper surface of a wing produces lift.

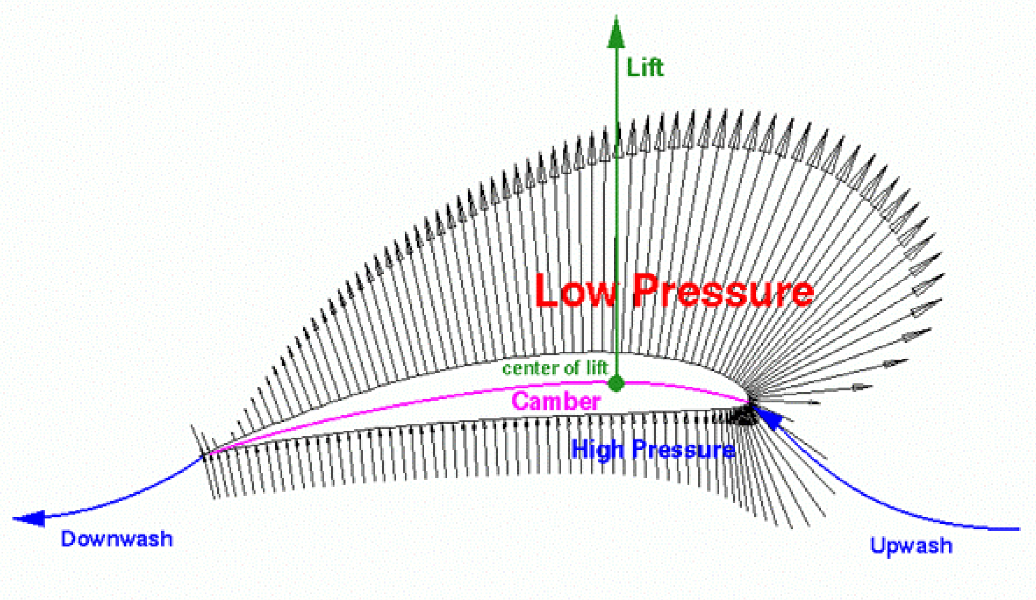

Pressure vectors and flow over a cambered section.

Camber

If the thickness of an airfoil is ignored, it can be reduced to a thin curved line defining the camber. The shape of this camber line determines the amount of lift produced at a fixed angle of attack. Since a sail has essentially no thickness, it exists only as camber. The flow over the convex leeward side has reduced pressure (through accelerated flow) and the flow over the concave windward side has increased pressure (through decelerated flow). The difference in pressure across the sail holds the flexible sail into its cambered shape and produces force to pull the boat.

Upwash

An airfoil developing lift causes the flow approaching it to bend upward. This is because the lower pressure on top of the airfoil pulls air up toward it. This upward change in flow angle is called upwash.

Planform Effects

The planform of a wing is defined by the shape of the leading (front) and trailing (back) edges.

In addition to the upwash that an airfoil causes on itself due to the lower pressure on top influencing more air to flow over it, additional upwash occurs due to changes in the planform of the wing. This is because, just as the low pressure on top of the wing influences the air some distance upstream to move upward toward it, that low pressure also influences air a similar distance away in the spanwise direction to alter its direction. This causes variations in upwash along the span of the wing on adjacent sections.

Sweep

The sweep of a wing is defined as the angle between a line perpendicular to the flow and a line (called the quarter-chord) passing through the 25% chordwise (luff to leech) positions along the span. The 25% chordwise position is chosen because, typically, the load on a section can be thought of as being centered there. This is because an airfoil generates much more lift in its forward portion than it does aft, so using the quarter-chord line as a reference is a convenient manner to characterize the sweep of a wing.

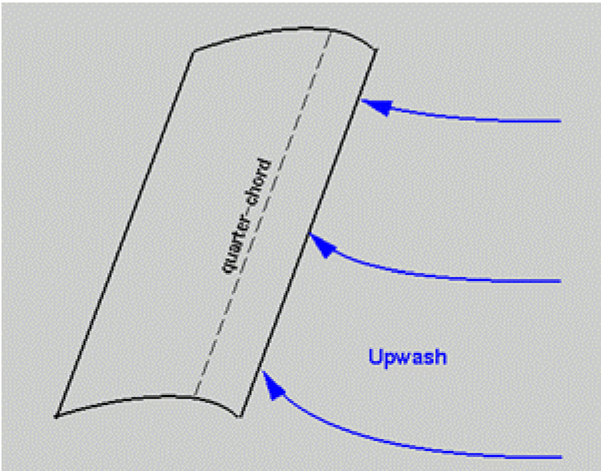

Upwash increases outboard due to sweep.

Sweep has the effect of increasing the upwash on the outboard wing sections. As a wing is angled aft, flow over the outboard sections must pass by the low pressure on top of the wing sections immediately inboard and forward. The close proximity of that low pressure to the air just outboard causes the outboard flow to turn upward more, resulting in higher upwash on the outboard wing.

Taper

Taper is defined as the ratio of the chordlength of the tip divided by the chordlength of the root. For sails, where the head tapers to nearly a point, the taper is extreme (zero), resulting in a triangular planform.

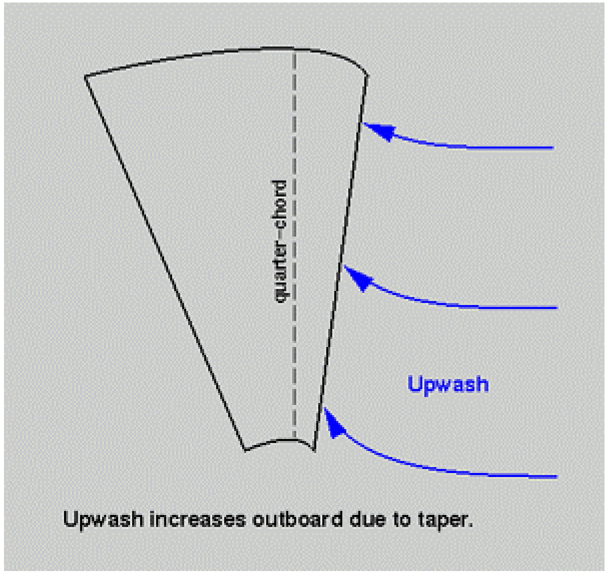

A tapered wing has a much shorter tip section than root section. As the wing tapers, lift produced by the shorter outboard sections is less because they have less surface area to support lift. Since the outboard sections are smaller than the inboard sections, they are significantly influenced by the larger wing just inboard. Air approaching the outboard portion of the wing is deflected by the low pressure on top of the larger inboard wing that is still generating a large amount of lift only a short distance away. The close proximity of that low pressure to the outboard wing causes the flow to be pulled upward additionally over the outboard wing. Hence, the smaller outboard sections operate with higher upwash. This enhances the amount of lift that they produce but does not make up for their loss of area.

Flow Conditions in Earth's Boundary Layer

Identifying the flow conditions that sails operate in is very useful for understanding how they work. The wind blows over the surface of the earth and, as with any fluid flowing over a surface, has friction with it. This friction slows the air closest to the surface and through shear causes the air immediately above it to slow some, too. This effect continues upward until at some distance above the surface the air is all moving at a similar speed. This behavior is called the boundary layer. While it occurs at a very small scale in the water flowing along the surface of hulls and keels, it occurs at quite a large scale in the air flowing over the earth. This means that the true wind speed is increasing up the entire height of a mast.

Apparent Wind

Apparent wind is the wind velocity experienced by the sails on a moving boat. This is the wind speed and direction that can be directly measured (felt) from the boat while it is moving. It is a combination of the true wind and the wind generated by the motion of the boat. The figure shows how these two wind components are added to create the apparent wind.

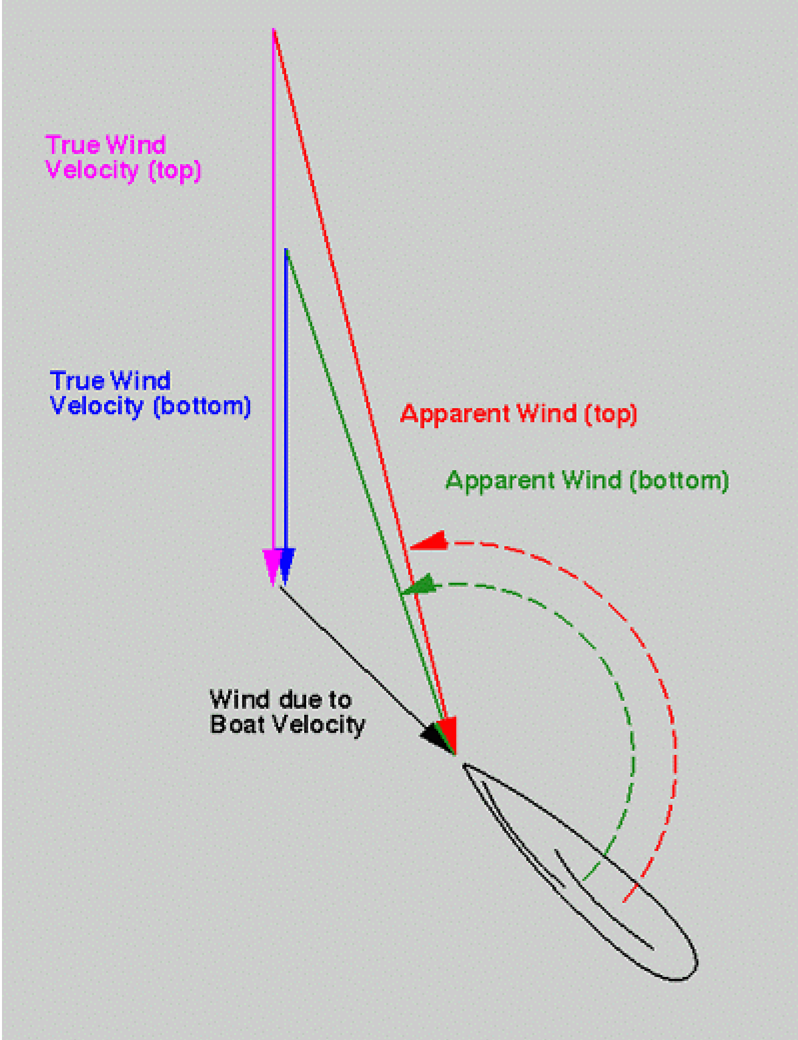

Flow twist due to vertical velocity gradient.

Notice that the apparent wind vector at the bottom of the rig, where the true wind speed is slower, is shorter (slower) and angled from a more forward direction, than the apparent wind vector at the top of the rig, where the true wind speed is faster. The true wind is coming from a single direction in this example, but varies in speed with height due to the earth's boundary layer. This variation in true wind speed not only causes the variation of apparent wind speed with height, but also its variation in angle. This is because all of the mast and sail are moving at the same speed and in the same direction as the boat across the moving air. Since the wind solely due to the movement of the boat is identical at all heights, the apparent wind speed and direction resulting from its addition to different true wind speeds at various heights is different.

While in this example the true wind velocity only varied in strength with height, it is possible that a variation in true wind direction can occur with height. In that situation, each tack will experience different apparent wind twist than the other.

Twist

Since the flowfield that a sail experiences is twisted due to the movement of the boat through the earth's boundary layer, the sail needs to incorporate some twist in order to fly in that flowfield. The increase of apparent wind angle with height is a factor that influences a sail to fly in a twisted manner, where the top is angled more off-center from the boat than the bottom. Other factors affecting how much twist is appropriate are sweep and taper as they alter the amount of upwash along the span of the sail.

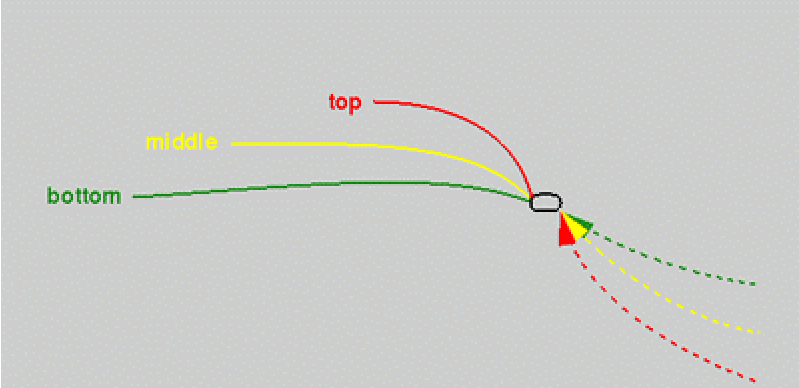

Effect of twisted apparent wind and upwash on sail's twist.

Isolated Sails

A mainsail by itself (cat rig) is tapered, but if the mast is close to vertical is actually swept forward. Recall that sweep is measured relative to the 25% chord line, which in the case of a tapered sail on an upright mast is angled forward. In this case, the forward sweep would have somewhat of a canceling effect on the increased upwash due to taper. The actual degree of upwash depends on the magnitudes of taper, sweep, and aspect ratio (height/width) of the sail. The sail still operates in the twisted flowfield caused by the boat moving through the earth's boundary layer, so an amount of twist would be appropriate. Raking the mast back increases sweep and will cause additional upwash on the top of the sail, necessitating more twist to the sail.

Genoas and jibs are very tapered and swept. Those two features, combined with the already twisted apparent wind, cause significant upwash toward the head of the sail.

Sails in Combination

Each sail by itself is much simpler than the combination of a foresail and mainsail as in the sloop rig. The sails are operating so close to each other that they both have significant interaction with the other. The most interesting feature of this is that the two sails together produce more force to pull the boat than the sum of their forces if they were each alone.

Earlier, upwash was identified as the increase in flow angle immediately upstream of a wing. There is also a corresponding change in angle, called downwash, just behind a wing, where the flow leaving the wing has been turned to an angle lower than the original flow. This is the cause of the well known "bad-air" that a boat just to windward and behind another boat experiences.

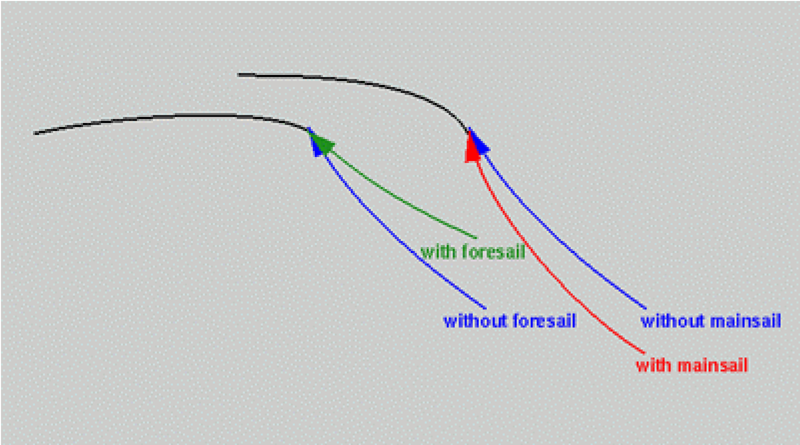

Upwash variations due to the presence of the other sail.

The mainsail of a sloop rig operates in the downwash of the forward sail, causing the flow angle approaching the mainsail to be significantly reduced from what it would be otherwise. This decreases the amount of force that the mainsail produces. The observed affect commonly referred to as "backwinding" is partially a result of downwash from the foresail, but is also due to the higher pressure on the windward side of the genoa being very close to the forward, leeward side of the mainsail, causing the flexible material of the mainsail to move away from that higher pressure.

The foresail of a sloop rig operates in the upwash of the mainsail. The wind as far upstream as the luff of a genoa is influenced by the upwash created by the mainsail. Hence, a jib or genoa in front of a mainsail has a higher flow angle than it otherwise would have by itself, causing an increase in the amount of force that the forward sail produces. So, while the mainsail is experiencing detrimental interference from the foresail, the foresail benefits from the interference of the mainsail. Notice that more air is directed around the curved leeward side of the foresail. This causes higher velocity (lower pressure) and more force. The net result is that the total force of the two-sail system is increased, with the foresail gaining more than the mainsail loses.

There is a converse affect to a windward boat receiving "bad air" (downwash) from a boat ahead and to leeward. A leeward boat gains additional upwash ("good-air"?) from a boat just to windward and slightly behind that acts like a lifting windshift until it moves ahead of the windward boat. This is the same phenomenon from which a foresail of a sloop rig benefits.

Another consequence of the difference in flow angles that the two sails experience in each others' presence is that the mainsail must be trimmed to a much closer angle with the boat's centerline than the foresail, which is able to be trimmed to a lead position well outboard. This angle represents the difference in upwash on the foresail and downwash on the mainsail due to each other.

Masthead Rig

On a masthead rig, where the forestay is attached to the top of the mast and both sails taper to basically zero chordlength at their heads in a similar fashion, the interference effects of the sails on each other are similar along the entire height of the mast. The mainsail ends up being rather tightly trimmed all the way up because of the genoa's downwash, and the genoa gains from favorable upwash all the way up.

Fractional Rig

A fractional rig has the more complicated characteristic that the top of foresail is not as high as the top of the mainsail. This means that the top of the foresail is very close to the front of the mainsail at a height where there is still an ample amount of chordlength in the mainsail. As the foresail luff approaches the mainsail luff, the upwash on the foresail due to the mainsail increases, because the low pressure behind the mainsail has more affect the closer the flow gets to it. This causes the top of the foresail to experience even more upwash and contributes to a fractional rig's foresail being trimmed more twisted than a masthead rig's foresail.

The top of the main on a fractional rig extends well above the foresail, leaving the upper portion of the mainsail free to experience the apparent wind without the downwash interference of the foresail. Apparent wind toward the top of the mast comes from a much higher angle, so the mainsail above the foresail experiences much higher wind angles than the lower portion of the mainsail where the genoa is causing substantial downwash. This change in flow angle with height on a mainsail is quite dramatic with a fractional rig and leads to trimming a fractional rig's mainsail with more twist than a masthead rig's mainsail.

Flow Angles

Reviewing all of the affects so far reveals that both sails experience increasing flow angle with height. The foresail operates in the twisted flow of the apparent wind, with upwash induced by itself due to taper and sweep, and in the upwash field of the mainsail. The mainsail is operating in the same twisted apparent wind, with additional upwash caused by its taper, but somewhat lessened by its forward sweep. It is also flying in the downwash field of the foresail, which is probably twisted because the foresail flies in a twisted fashion. This is particularly exaggerated with a fractional rig.

Sail Shape

With the flow directions established, it is now useful to consider the ramifications of sail shape. Previously, it was stated that a sail section exists solely as camber. Now it is interesting to explore the differences in camber that are possible and what would be most beneficial.

Since a sail is constructed of flexible material, its cambered shape is supported by the pressure difference that it generates. It follows that the leading edge entry angle of the sail must be reasonably aligned with the incoming flow angle. If the entry angle is too high the sail will luff, and if it is too low the sail will stall, since the flow would be required to turn an impossibly sharp corner around the luff. It is also apparent that the entry angle should increase with height to match the twisted flowfield. There are two remaining issues. Where should the trailing edge be, which defines the angle of attack at each height, thus twist? What path to take to get there, or what should the specific cambered shape of the sail be?

The trailing edge location in relation to the leading edge locations establishes the angle of attack of a particular section. Lift increases proportionally with angle of attack, so, since a sailboat is trying to extract as much force as possible from the wind (until overpowered into heeling too much), it would seem best to position the trailing edge (leech) as close to the boat's centerline as possible. This would achieve the highest angle of attack and hopefully the most lift, but unfortunately the ability to trim a sail to unlimited angles of attack is not possible.

Separation and Stall

Eventually at some angle of attack, sail sections (and airfoils) experience stall. This occurs when the air flowing around the leeward side of the sail no longer travels on the surface of the sail. The flow separates from the sail resulting in a large loss in lift. Depending on the shape, stall can occur abruptly with a small increase in angle of attack, or more gradually with some indications that the flow is separating from the surface at specific locations first. This is easily seen using telltales (tufts of yarn) that swirl erratically when the flow departs from the desired direction instead of streaming aft when the flow is attached to the sail.

Separation occurs simply because the pressure gradient that the flow is trying to pass through is too extreme. Recall that lift is generated because the flow accelerates around the convexly curved leeward surface of the sail creating low pressure. Eventually, as the flow approaches the back of the sail, the flow must slow down to near its original speed and pressure, since after it leaves the sail it will return to its original state when the sail is no longer there to influence it. This is referred to as pressure recovery.

Another way to think about this is that when air flowing over the leeward side of the sail and air flowing over the windward side of the sail reach the trailing edge, they must have the same pressure, as there will not be anything in between anymore to enable maintaining different pressures. It does not mean that two particles of air that start at the leading edge and travel along different sides of the sail will arrive at the trailing edge at the same time (a common misconception). It simply means that the pressure of air flowing off the top right at the trailing edge of the sail will be equal to the pressure of air flowing off the bottom. This must be true since they are coincident there. That pressure is generally close to the original pressure of the flow prior to being disturbed by the sail.

So, accelerated flow around the leeward side slows down toward the leech in order to provide the necessary matching at the trailing edge as the air is returned toward its original conditions. Overall, air still travels much faster around the leeward side of the sail than the windward side.

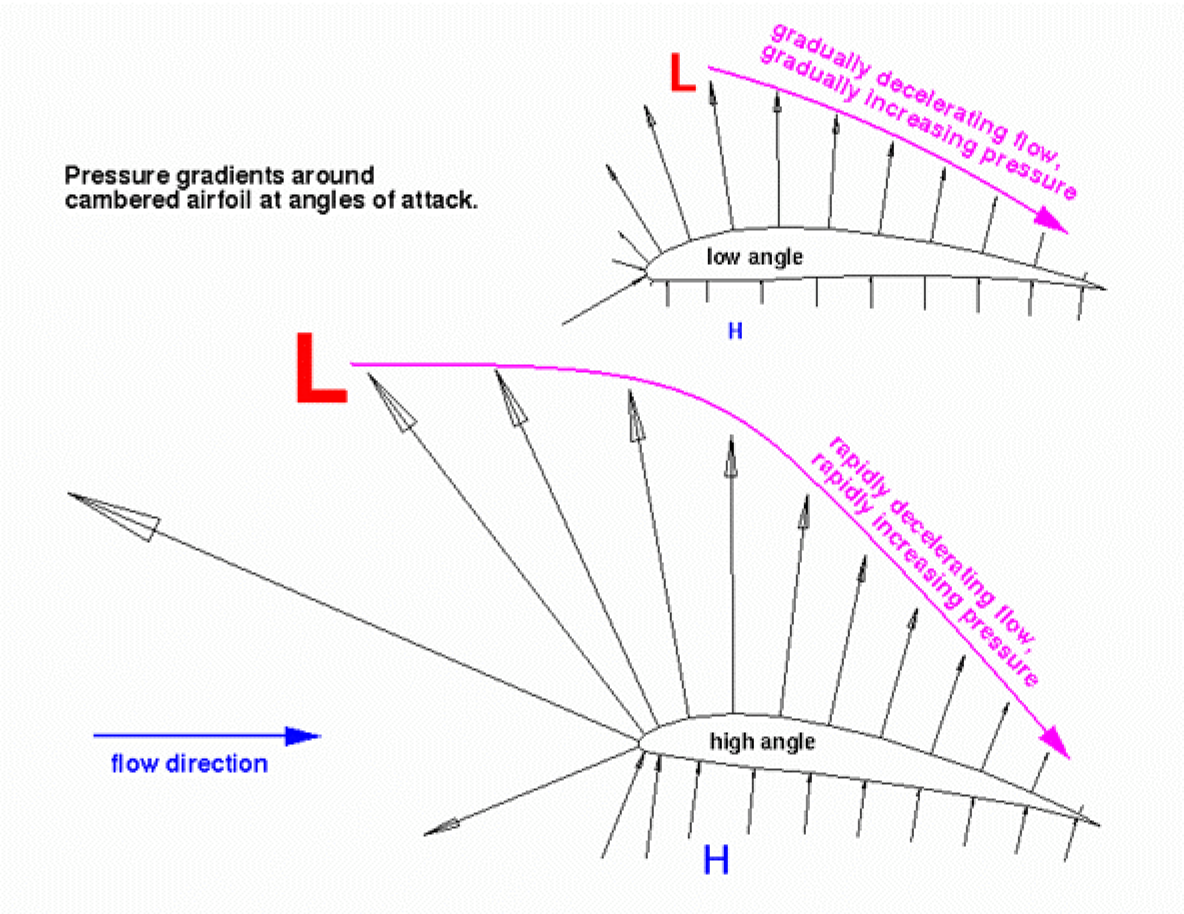

As the flow slows down from its accelerated state on the leeward side it yields a pressure gradient along the back of the sail that is increasing from very low pressure (to produce the desired lift) to a much higher pressure toward the leech. The amount of initial acceleration (dependent on angle of attack and shape) and the length of the pressure recovery determine how steep this gradient is. When the increase in pressure that the flow is experiencing becomes too extreme, the flow no longer stays attached to the surface of the sail. It is pushed away by the higher pressure and stall occurs, yielding less lift. This happens when the angle of attack is high and/or the camber is quite deep, causing very high velocity, but necessitating dramatic slowing of the flow, too, with the associated rapid increase in pressure causing separation. It is favorable to slow the flow in a smooth fashion over a longer distance so that there is no steep rise in pressure. This happens most effectively over a long, straighter shape aft, lacking in curvature that would attempt to promote higher velocity.

Pressure Distribution and Curvature

As the air flows around both sides of the sail, its pressure changes with the varying local velocity caused by the curvature of the sail. With the entry angle defined by the oncoming flow direction and the angle of attack governed by avoiding stall, there are still numerous sail shapes that can be established to connect the luff and the leech.

Since the purpose of the sail is to develop force to move the boat in a forward direction, it would be most effective to have as much of the sail as possible operating with the largest possible pressure difference across it. The way to achieve that is to accelerate the air quickly around the curved leading edge of the sail in order to generate low pressure on the leeward side close to the luff and then maintain it back over a significant portion of the sail. This is achieved by imparting high curvature to the front of the sail. Once the flow is accelerated curvature can diminish and the flow will continue quickly around the leeward side of the sail. The back of the sail needs to be flatter in order to allow the flow to gradually decelerate to avoid stall as already described. These details are the basic factors that define the rounded entry with forward draft (position of maximum camber depth) and straight leech profile that has proven fast in typical sailing applications.

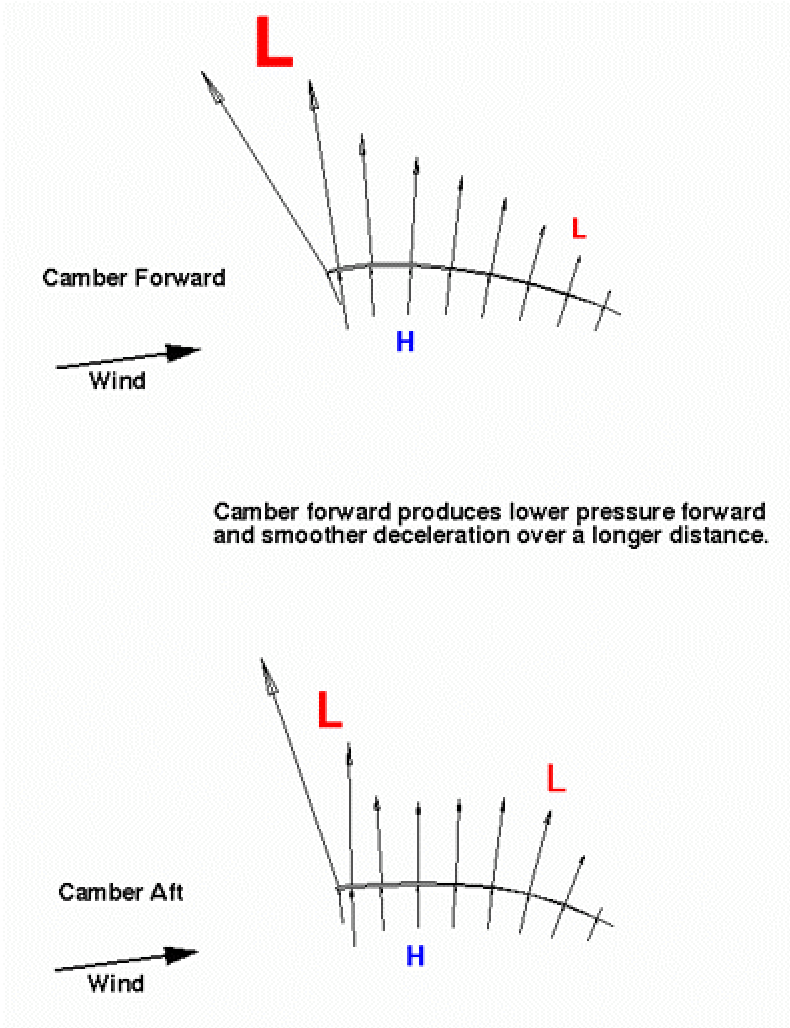

It is evident that a sail shape with more curvature aft keeps the air accelerated longer, which may produce a higher amount of total force due to a larger region of negative pressure. The problem is that the negative pressure vectors on the aft portion of the sail are angled more aft than those toward the front of the sail, so the amount of forward force that is produced in the direction that the boat travels is less, while the amount of sideways force that contributes to heeling and leeway is greater. It is also apparent that a shape with its camber aft has a shorter, steeper pressure recovery that will lead to earlier separation and stall.

One more factor to consider is that a rounder leading edge or deeper camber, while producing more force, does so at the expense of having a higher entry angle that requires a higher apparent wind angle in order to fly without luffing. This means that the boat cannot be sailed as close to the wind, which explains why spinnakers can be so full and deep, but becomes a tradeoff when setting an upwind sail. Producing more force versus sailing at a higher angle becomes a subtle optimization. It requires the proper balance between a full, curved leading edge to accelerate the flow, and a flatter, subtly curved leading edge that does not produce as much low pressure to pull the boat but does allow the boat to sail closer to the wind. Boats that sail fast with large amounts of sail area will favor flatter sails than slower boats with less sail area that need to develop more power to move the boat.

Induced Drag

Another factor that must be considered is induced drag. This is the drag that a wing generates when it creates lift. Over most of a wing, the low pressure above the wing is kept isolated from the higher pressure underneath by the physical presence of the wing. At the tip of a wing, where the wing ends, there is nothing preventing air from flowing around the wing tip from the high pressure beneath to the lower pressure above. This results in the standard tip vortex that is often seen spinning off the tips of airplane wings and flaps. When the flow takes this alternate path around the tip instead of over the airfoil surface, energy is expended that does not develop lift, but does cause drag. This is called induced drag and it increases exponentially with lift, so a wing, such as a sail, that is producing substantial lift, experiences much more induced drag than a wing that is producing a lesser amount of lift.

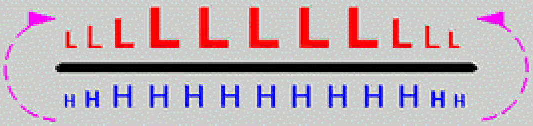

Rear view of a wing showing formation of wingtip vortices as air flows from high pressure below to low pressure above.

The most effective way to minimize induced drag is to increase span, as induced drag is inversely proportional to the span squared. Highly efficient airplanes like gliders have very high span for the amount of lift they are producing. Winglets are a way to create the effect of higher span without actually increasing the physical span. They are useful when there is an artificial constraint on wingspan (like a draft limitation on a keel).

Spanload

Beyond being heavily dependent on the amount of lift, induced drag is also dependent on how that lift is produced. It has already been explained how taper and sweep affect the upwash along the span of a wing by causing sections along the wing to have different lift levels. Also, varying the amount of camber and the angle of attack of sections along the wing will influence how much lift is generated at various spanwise locations. The distribution of lift along the length of the wing is called the spanload.

It has been found for an isolated wing in untwisted flow that the spanload with the minimum induced drag is an elliptical distribution of lift. This is achieved on an untwisted, unswept wing with an elliptical distribution of area and the same section along its entire span. A spanload can be altered in several ways.

Tapering the wing causes the lift to be reduced outboard because, while more upwash is produced and the outboard sections are loaded more, there is less area to generate lift outboard. Sweeping the wing aft increases the upwash outboard and the lift there because the area is the same and operating at higher angle of attack. Twisting the outboard wing to higher or lower angles will increase or decrease the outboard lift levels, respectively. Finally, adding camber to the outboard wing sections will increase the amount of lift that they produce.

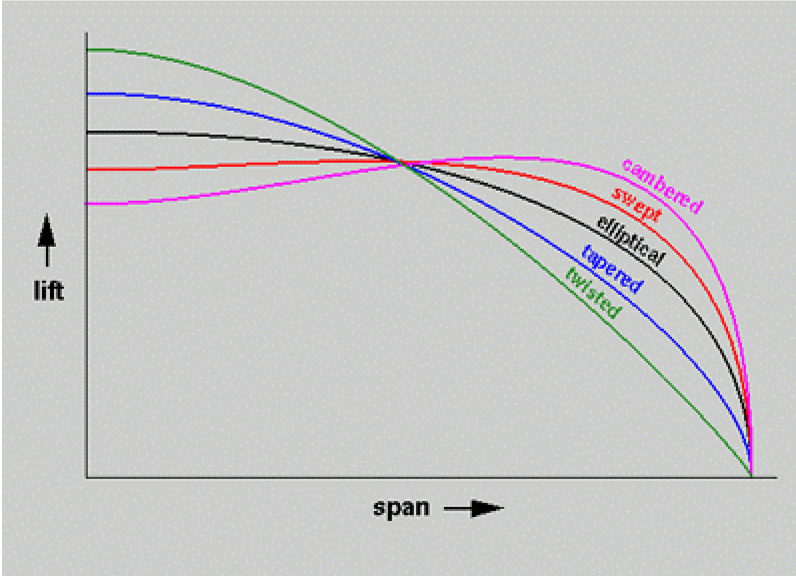

Spanload variations due to outboard feature.

All of these features can be used to modify the spanload with each of the resulting spanloads producing different amounts of induced drag for the same amount of total lift. This is because induced drag is a consequence of how the lift being produced by the wing deteriorates at the wing tip. A wing tip, since it has no more wing outboard of it, cannot sustain lift because it cannot support a pressure difference. Thus, the lift at the very end of a wing must be zero. The inboard portion of the wing produces a significant amount of lift that must diminish toward zero approaching the wing tip. The manner that the lift decreases toward the tip defines the shape of the spanload, and it is the character of that lift distribution that establishes the amount of induced drag.

The exact shape of the optimum spanload for sails in a twisted flowfield varies somewhat from the simple ideal elliptical spanload. Since the flowfield is twisted in a manner that lift toward the top of the sail is oriented more in the direction of the boat than the bottom of the sail (because lift is produced in the direction perpendicular to the local flow direction), it follows that the ideal spanload in twisted flow conditions will be even more highly loaded toward the top than the simple elliptical lift distribution that is optimal for untwisted flow. The extremely tapered planform of typical sails yield spanloads with much less lift toward the top than the elliptical spanload, so are less than optimal. While a genoa has considerable sweep to help reload its top sections, it has very little sail area near its head to develop lift. A mainsail without sweep, particularly on a fractional rig where the top of the mainsail is above the influence of the foresail, does not generate enough lift toward the top to approach the elliptical spanload.

Increasing the chordlength of the top of the sails would be an effective way to create additional lift toward the top and attain a loading closer to optimal. This has been demonstrated to be beneficial through the use of full-length battens, but is not always allowed. Another way to increase the lift levels toward the top of sails is to provide additional camber toward the top to boost the lift being produced up high. Increasing the angle of attack, through decreasing twist, would also increase the loading at the top, but the limitation of stalling the top sections must always be minded, so the sail needs to maintain a certain amount of twist because of the inherently twisted flowfield. It is highly likely that the optimum lift distribution to minimize induced drag is not achieved with typical sails.

Setting Sail

Recognize that the objective of the sails is to create force to pull the boat, but that there can also be a constraint on heel. At some point the stability of the boat or weight of the crew cannot keep the boat sailing at an angle that does not compromise performance, so just using the sails to produce the most force possible is not necessarily the fastest procedure.

In lighter winds, when the sails are struggling to extract enough force from the wind to move the boat fast, the sails should be set such that every section along the height of the sail is working to produce high lift, especially the top sections in order to minimize induced drag. When the wind builds beyond a level that the sails' force causes the boat to heel too much, the sails' characteristics must be modified. There are several options.

Reducing the amount of camber in the entire sail will decrease the amount of force the sail produces, as will decreasing the angle of attack of the entire sail. Implementing these adjustments over the entire sail may or may not be the best alternative for the windier conditions. They reduce the amount of force generated by the sail, but that force is still centered at a similar height. In order to reduce the heeling moment created by the sails to a satisfactory level, the amount of force may decrease to a level that does not pull the boat very fast anymore.

Another approach is to reduce the lift produced by the top of the sail. Reducing the camber of the top of the sail, and/or reducing the angle of attack of the top of the sail through additional twist will affect the sail's force such that the remaining force is centered lower down. A similar reduction in heeling moment as simply reducing the entire sail's force can be achieved through depowering the top of the sail, but while maintaining more total force to pull the boat. The force is centered lower as the bottom of the sail is still trimmed in a fashion that generates substantial lift. This method has the compromise of deviating further from the desired elliptical spanload, as the lift distribution diminishes much more rapidly toward the top of the sail, and causes higher induced drag. The question becomes whether the remaining higher sail force offsets the additional drag component.

A parallel situation occurs with airplanes. Airplanes are not designed to fly with the optimal spanload that yields minimum induced drag because the higher outboard load on the wing would require that the wing be made stronger, hence heavier, to carry that load. It is more efficient to build the airplane lighter and generate more lift on the inboard wing and accept a little more induced drag. This is the same tradeoff that a sailboat experiences in strong wind when heeling becomes a factor and results in a similar, less than optimal spanload in order to maximize performance.

Pointing

With all of these scenarios, the angle that the boat is able to sail to the wind is always a consideration, too. If the angle of attack of the entire sail is decreased, much of it may backwind or luff and the boat may make very little progress due to lack of power. This is a consequence of the original restriction that sails be flexible. The boat can be turned away from the wind in order to fill the sail but that causes the sail to load again and the boat to heel too much making this a less viable option (except to avoid sailing too directly into waves). The component of the boat's speed in the windward direction must be accounted for when considering course variations. Twisting or decambering the top of the sail keeps the bottom of the sail still trimmed at an effective angle of attack, continuing to produce force and allowing the boat to be sailed at a closer angle to the wind.

The correct solution is generally a combination of the various adjustment options and will vary with wind speed and sea conditions. It is also dependent on the characteristics of the boat and rig, and the trim controls available (and probably even with the time interval with which the wind velocity is varying as some adjustments are made more quickly and easily than others). Hopefully, by understanding the lift producing characteristics of the sails and how to manipulate them, a sailor can continually alter the sails in order to produce the most effective force to move the boat in the intended direction.

Summary

It is evident that sails are flexible wings, operating in a twisted flowfield, and in the presence of each other. They produce force by accelerating air over their curved leeward side causing lower pressure on that side of the sail that acts to propel the boat. Camber and angle of attack can be controlled along the span in order to produce lift in many possible ways, some more favorable than others. Stall and luffing provide an upper and lower limit on the useful angles of attack.

Taper and sweep affect the upwash of flow approaching the sails beyond the influence of the twisted apparent wind caused by the boat moving through the earth's boundary layer.

Paul Bogataj is an aeronautical engineer, specializing in sailing applications. He was responsible for appendage development for Team Dennis Conner and Young America in the 1995 and 2000 America's Cups, respectively. He approaches sail and keel design from the perspective of using advanced aerodynamic technology and methods that have been developed for the aircraft industry. He previously worked for Boeing, but currently consults independently for a variety of sailing design projects. He has employed his knowledge of how sailboats function to win the North American Championship of two different classes, and numerous fleet and district championships. Paul combines his practical understanding of sailing from his experience as a successful racing sailor with his awareness of fluid dynamic principles as an engineer to provide explanations of how sailboats work that are understandable to the average sailor. To contact him about this article, click here.