

Vortex Generators for Cooling

A simple fix reduced CHTs in a Velocity by 55 degrees.

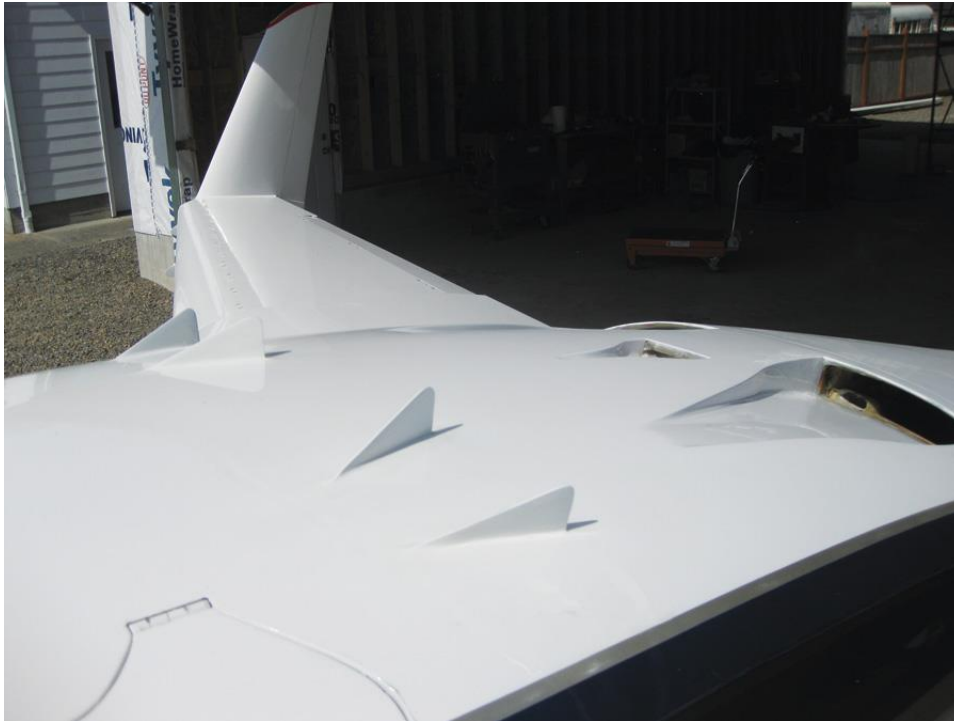
By

[David Ullman](#)

-

October 15, 2015

[2](#)



The cooling air for the IO-360 engine in my Velocity SEFG comes through two NACA style ducts in the top of the fuselage. The air then downflows through the cylinders and is exhausted out the rear of the fuselage about 2 inches in front of the propeller. I have flown the plane for about three years, and it has always run hot.



I added external scoops on the rear edge of the NACA ducts and that helped, but looked crude and not very elegant as the NACA ducts were supposed to be low-drag, internal scoops. When I painted the airplane, after three years in primer (a color I called “blotch white”), I took the scoops off. The combination of no external scoops and a smooth paint surface made the NACA ducts very ineffective. On climb-out and in cruise, my engine was overheating with cylinder head temperatures (CHTs) above 425° F.

This led me to study different methods to get more air through the engine. First, I tried to see how much air was going through the NACA ducts by putting smoke oil on the top and flying around the pattern. Traces showed that air was indeed flowing into the ducts, but this gave no indication of how much.

Early efforts to make the engine run cooler were based on suggestions from flying colleagues; latter efforts were based on studying literature to find a good engineering solution. I call these the “hacking” phase and the “engineering” phase, respectively. I will document what I tried and also detail the engineering solution. To give away the ending, I knocked 55 F off the CHTs with a simple fix. It just took a while to find it.

Understanding the Problem

One check to learn more about the airflow through the engine was to measure the pressure drop across the cylinders. Lycoming specifications state that there should be a drop of 6 inches H₂O for adequate cooling. To measure this, and to better understand the pressures created by the NACA ducts and the changes, two manometers were used. These were made of long lengths of clear tubing and some water with red dye and a drop of dish soap in it (Figure 1).

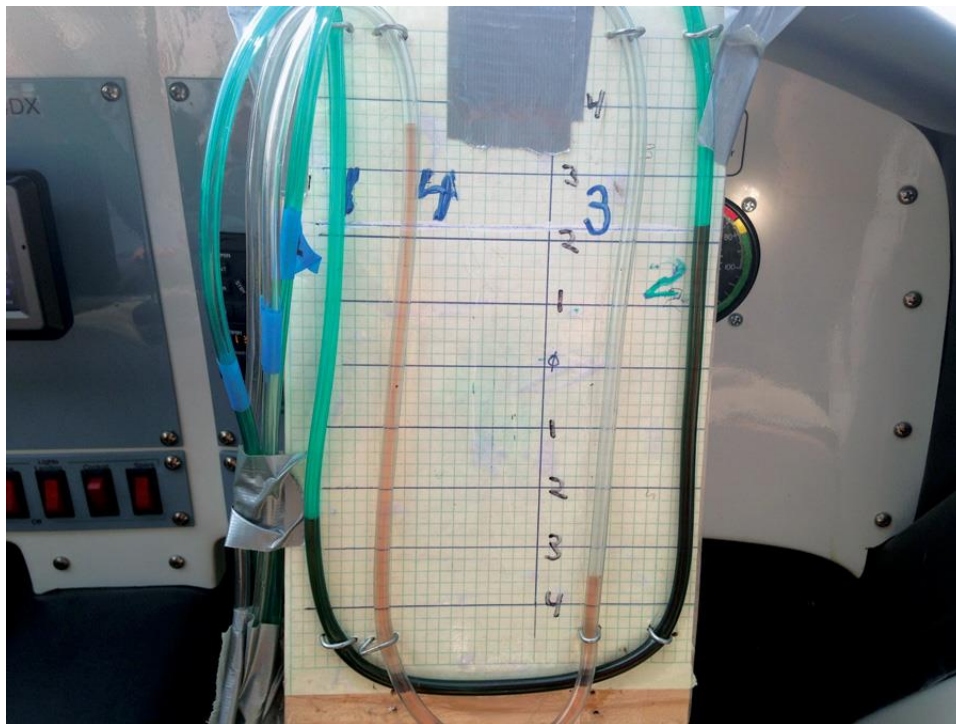


Figure 1: Two manometers were used for testing. One measured the pressure difference between the static pressure and the plenum above cylinders 3 and 4. The other measured the difference in pressure across cylinders 1 and 2.

The first manometer measured the pressure difference between the static pressure and the plenum above the cylinders labeled 3 and 4 in the photo. One end was plumbed directly into the static system in the airplane. The other end was secured to the fuel injection spider in the plenum above the cylinders. The end of that tube was blocked off and holes were drilled around the periphery of the last inch of the tube so that it was clearly sensing static pressure.

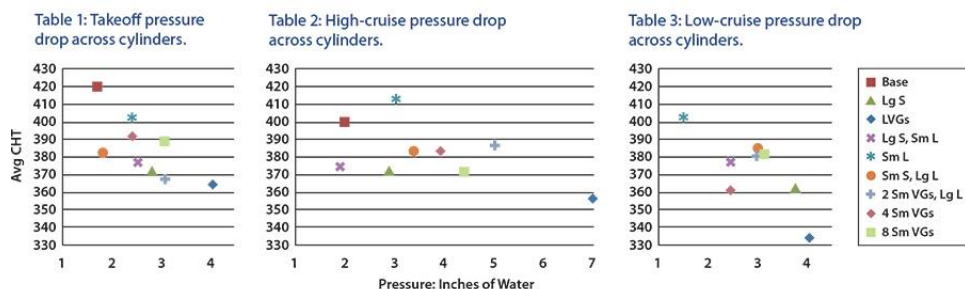
The second manometer measured the difference in pressure across the cylinders labeled 1 and 2 in the photo. One end was mounted in the upper plenum next to the tube from the first manometer. It too was plugged and drilled. The other end was mounted just below the cylinders under the engine and was plugged and drilled.

The manometers themselves were mounted on a board in the cockpit so that a copilot could photograph them for later data reduction. The example shown is from late in the experiments and at high velocity. It shows 4.6 inches of H₂O across the cylinders and 7.2 inches of static pressure in the plenum above the cylinders. Initially, before any additions, the pressures were 2.0 inches H₂O and 3.0 inches H₂O.

The Hacking Phase

There were two schools of thought on what to do: Push more air into the NACA ducts or pull more air out the back of the fuselage. Most of the advice came from builders with front-engine experience, where often the problem is that not enough air is being pulled.

A variety of hacks were tried: Scoops were used to push more air into the NACA ducts. Louvers were added to the bottom of the cowl to pull more air out. Small vortex generators (0.43 inches high) left over from the wing installation were placed in front of the NACA ducts to force more air in. Sometimes a variety of ideas were used in combination.



The results of testing can be seen in Tables 1, 2, and 3. Three test conditions were used:

Climb	115 kts (104 ft/sec). Note that VX is 95 knots, but I generally climb out at this higher speed.
Low Cruise	125 kts (114 ft/sec)
High Cruise	170 kts (155 ft/sec)

Data was taken from the manometers and the average cylinder head temperature. The engine was run full rich for all test points to be consistent. All temperatures were corrected for the outside air temperature (OAT) by normalizing them to a 60 F day. The results are plotted in Tables 1, 2 and 3. The points are for:

Lg S	Large Scoop—stuck up 3 inches above the fuselage surface.
------	---

Sm S	Small Scoop—stuck up 1.5 inch above the fuselage surface.
Sm L	Small Louvers
Lg L	Large Louvers
Sm VGs	Small Vortex Generators—these were the same vortex generators used on the wings and canard. They are 0.43 inch tall.
L VGs	Large Vortex Generators—these worked! They are 2 inches tall.

The results are shown for each of the three conditions. Data for the base condition, just the NACA scoops as built, is in the upper left corner of the takeoff and high-speed cruise plots. No base data was taken at the low-speed cruise condition. These points are all worse than they appear as the temperature was still climbing when I throttled back. Note that besides high temperatures, the pressure drop across the cylinders was only 2 inches or less—no wonder the engine was overheating!

The results in the lower right corner are for the final configuration with large vortex generators. Here the temperatures are acceptable and pressure drops 4-7 inches, much closer to the Lycoming 6-inch spec.

The other configurations are scattered between these two extremes. Some key points that can be taken away from these are listed below.

- The Large Scoop (Lg S) helped on takeoff and high-speed cruise, but not as well on low-speed cruise. But this scoop was both ugly and increased the drag (no firm data on this).
- The louvers on the bottom did not make significant difference.
- Small VGs helped some, but it was unclear how many to use and where to put them.

Note that not all combinations were tested, as that would have been too many runs. This was also trial and error, so not all options were known beforehand.

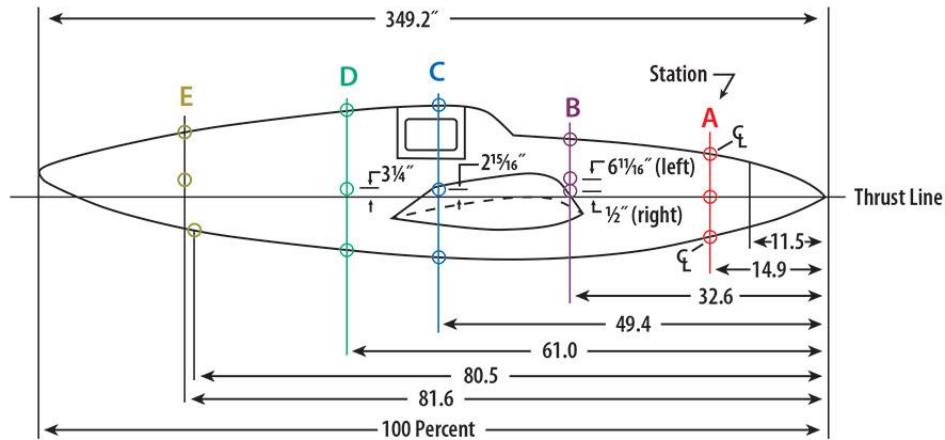
The Engineering Solution

Parallel to the hacking phase, I worked to understand the physics of what was happening. It became clear that even though air was flowing into the NACA duct as shown with the oil traces, there was not enough. The boundary layer, which increases in thickness on the fuselage, was keeping air out of the NACA ducts.

To explain what the boundary layer was doing and why it is important, here are some basics. These are all worded from the viewpoint of the surface, with the air moving past it, as it makes easier reading.

The boundary layer is the region of air near the surface. At the surface it's not moving at all, and at some distance out, it's moving at the speed of the air flowing over the body. We usually think of the boundary layer as quite thin. It isn't!

The actual thickness of the boundary layer can be seen from the results of an experiment described in a [NACA Technical Note¹](#). In this note, the authors measured the velocity of the air near the fuselage of an unidentified fighter (Figure 2). They had removed the propeller, antennas and other protuberances, and sealed all ducts. They measured the velocities in the boundary layer on the top, bottom, and sides at various angles of attack.



(a) Locations of pressure rakes on the fuselage of the model.

Figure 2: Details of pressure-rake locations and fuselage contours. (Courtesy of National Advisory Committee for Aeronautics [NACA])

Typical of what they found is shown in Figure 3. Here the vertical axis is the ratio of the speed of the air in the boundary layer divided by the speed of the air in the free stream (u/U) for the various stations along the fuselage bottom. This bottom image is the clearest in the report, so it is used here. It is typical of the airflow on the top and sides. The edge of the boundary layer is generally defined as when $u/U = 0.99$ (air moving at 99% of the free stream velocity). So here the boundary layer on the bottom, at station E (81.6% the length of the fuselage), is about 5 inches thick!

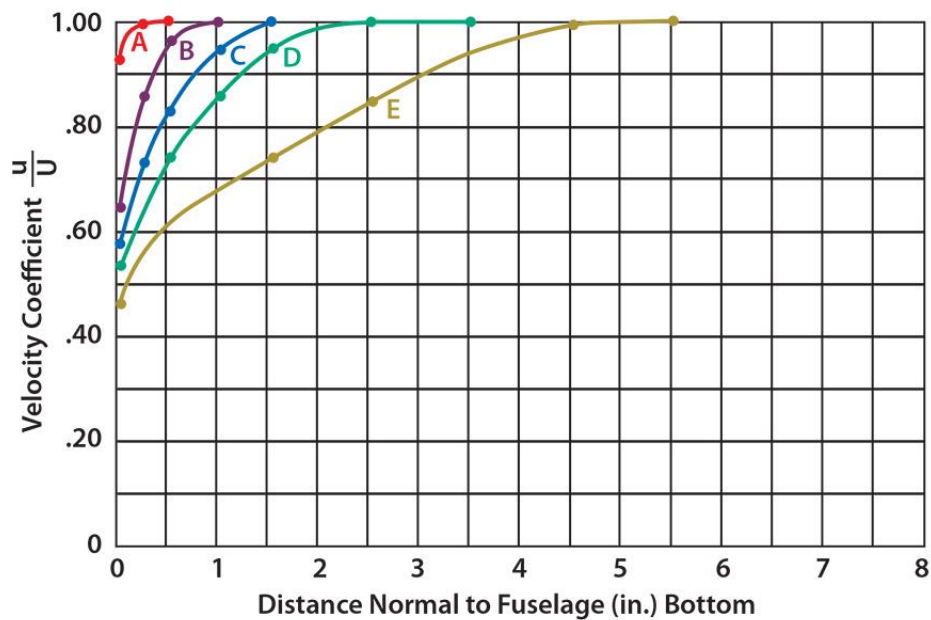


Figure 3: Boundary layer thickness at various locations on the fuselage shown in Figure 2.

Note further that the air near the surface is moving at only 50% of the free stream velocity; the NACA engineers could not get all the way to the surface with their pitot tube where the velocity actually goes to zero.

What is important here is that on top of the test plane, behind the cockpit, the boundary layer thickness (d) was measured at 3.0 inches with the airplane in a dive, 4.0 inches in cruise, and 5.5 inches in a climb. No wonder my NACA ducts didn't work as they should on the Velocity.



Small scoops protruded 1.5 inches above the fuselage surface. Large scoops (not shown) were 3 inches high.

To make sure that these results make sense, consider a simple explanation of the boundary layer theory. Theoretically, boundary layers start off laminar and, after a distance, become turbulent. Think of smoke coming off a match that has just been blown out. The smoke leaves the match as a smooth column and then, after a few inches, becomes a turbulent jumble. The first part is called laminar and the second, turbulent. On a fuselage, with its long distance, most of the boundary layer is turbulent. For a turbulent boundary layer, the thickness over a flat plate is:

$$d = x * 0.16 / (Re)^{1/7}$$

where:

x = the distance from the front in feet

Re = Reynolds number which for standard conditions is $= 6350 * U * x$ (U is in ft/sec).

The Reynolds number is a non-dimensional number that can be used to determine if the flow is laminar or turbulent. If below about 1×10^6 , the flow is laminar and above this value, turbulent.

The formula above is for a smooth, flat plate. The shape of the fuselage and the surface smoothness affects this in complex and second-order ways. Assuming this is adequate, then for the fighter in the NACA report, $x = 23.7$ feet (81.6%), and the tests were run with $U = 63$ mph or 92 feet per second. Thus,

$$d = 23.7 * 0.16 / (6350 * 92 * 23.7)^{1/7} = 0.36' (4.34 \text{ "})$$

This result is close enough to that measured at the cruise condition to give comfort that it is OK to use on the Velocity.

Then, for the Velocity, the NACA ducts are about 11 feet from the nose and thus:

$$d = 11 * 0.16 / (6350 * U * 11)^{1/7} = 0.36 / U^{1/7}$$

As can be seen, the speed has little effect, and the boundary layer is about 2 inches for all conditions. It is then no surprise that the small vortex generators tried earlier (0.43 inches tall) had so little impact. They were only stirring the bottom layers of air, those with low velocity.



Small louvers on the bottom rear of the cowl helped draw air out.



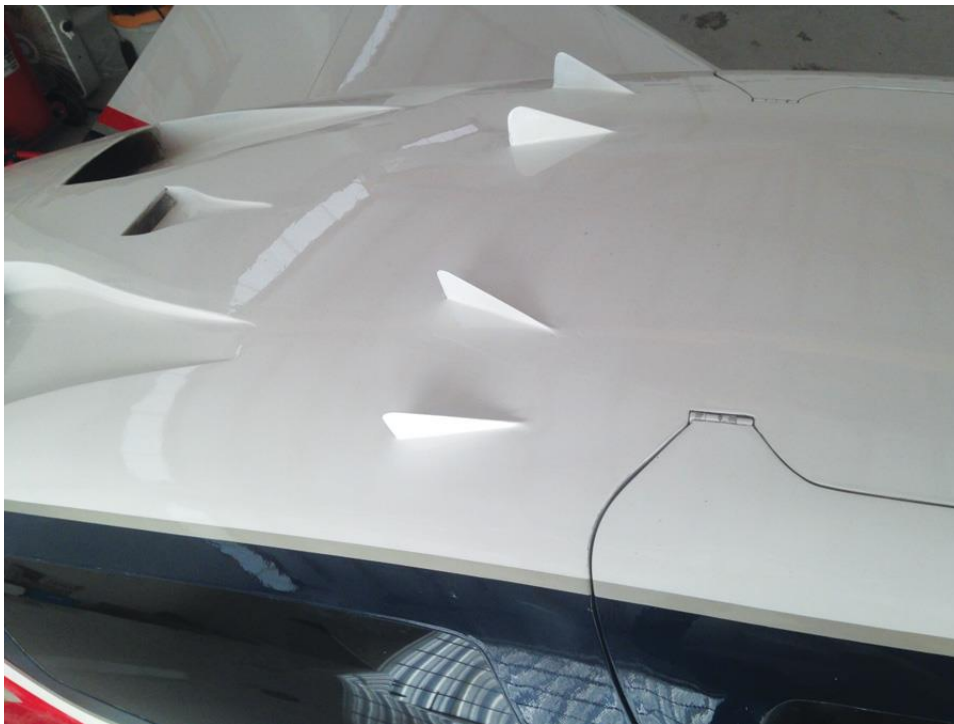
Large louvers placed over the same holes as the small louvers, but with much more projection into the slipstream.

[Adding Vortex Generators](#)

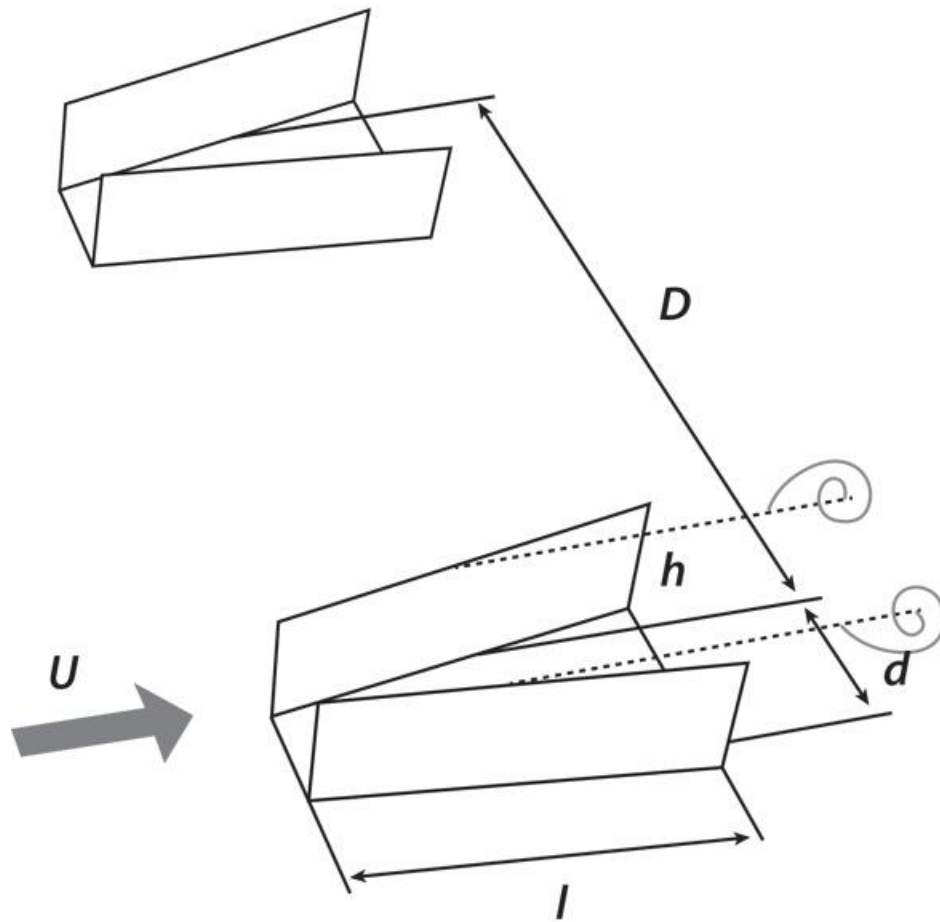
I put vortex generators on the wings and canard of my Velocity from the beginning. My test pilot strongly encouraged this, as I was a low-time pilot. They make the handling very docile. I know this because just before I painted my plane, I took off the inner of the VGs on both the wing and canard. The plane was much looser at low speed. After painting I put all of them back on, but I did not really understand what they did to the airflow.



Two small VGs, 0.43 inch tall, were tried first in front of each NACA duct. They helped, but not enough. The small VGs in the foreground were added later to help bring flow in when at a high angle of attack.



Large VGs that are 2 inches tall and 5 inches long solved the problem. They are placed in pairs 15 inches in front of the start of each NACA scoop.



The VGs I used on the lifting surfaces were 0.43-inch tall and at about 22% of the chord. Since the airfoils on the Velocity have the maximum thickness at 35%, the boundary layer is still laminar ($Re < 1 \times 10^6$).

A really good article on vortex generators on certificated airplanes is on AVweb ("[Vortex Generators: Band-Aids or Magic?](#)"). This gives a good overview of the basics for use on wings and tails.

A well-designed VG stirs the free stream into the boundary layer. This brings higher energy air (more velocity) into the boundary layer at the cost of a slight increase in drag. There are two design variables: the height of the VG relative to the height of the boundary layer, and orientation of the VG to the free-stream air and adjacent VGs.

To be effective the VG must reach into the free-stream air or near to it. For the Velocity cooling problem, the boundary layer is a little over 2 inches thick. Thus, the VGs need to be nearly that high to stir in free-stream air. The VGs tried during hacking were only 0.43-inch tall and, even at that height, they did some good, but there was more to be had. Note that some literature claims that VGs that reach 20% into the boundary layer are just as effective as those reaching into the free stream. It will be shown that this was not the case here.

The position and orientation of the VGs is also important. Typically VGs are oriented at 15-20 degrees from the flow direction. Thus they are like little wings at high angle of attack with a vortex rolling off of them.

Counterrotating VGs

There are a variety of styles of Vortex generators available. After researching the various options, I decided to try counterrotating vanes, which also happen to be the most common. Counterrotating vanes reinforce each other by driving air from the free stream down into the area between the VGs.

The design rules for counterrotating VGs are generally accepted to be:

- $h = 0.95 * \text{boundary layer height}$
- $D = 10 * h$ (Distance between set of VGs)
- $d = D / 4$ (Center distance between a pair of VGs)
- $l = 2.5 * h$

Thus I designed the VGs for cooling to be:

- $h = 2$ inches
- $l = 5$ inches
- $D = 20$ inches
- $d = 5$ inches

I put them about 15 inches in front of the start of the scoop and at 15 degrees from the centerline. I would have liked them farther forward, but wanted to stay away from the door opening.

At first I bent some aluminum VGs and pop-riveted them on for testing. When the data showed them effective, I replaced them with fiberglass. Table 4 shows the results. The values in the table are the same as in the earlier plots with two exceptions. First, data is also shown for the change in the static pressure in the plenum. Second, I did not take data for low cruise in the base condition. Thus I have used the small louver data instead, as it was near (actually slightly better than) base for the other conditions.

	Delta P across cylinders, inches H ₂ O			Plenum static pressure relative to static port, inches H ₂ O			CHT corrected to an OAT of 60° F		
	Base	Final	Difference	Base	Final	Difference	Base	Final	Difference
Takeoff	1.6	4.0	2.2	2.1	4.4	2.3	420+	365	55
High Cruise	2.0	7.0	5.0	3.0	9.2	6.2	400	357	43
Low Cruise	1.6	4.0	2.4	2.0	5.2	3.2	403	334	69

Table 4: Final results with large, 2-inch VGs placed in front of the NACA scoops.

The plenum static pressure is an indication of how well the NACA ducts are working. As can be seen, the addition of the VGs increased the pressure there dramatically (between 2.3 and 6.2 inches). Even the large scoop only increased the plenum static pressure to values about 70% as high as did the VGs.

Note that all data was taken with a payload of 480-520 pounds (pilot, copilot and 10-15 gallons of fuel).

Conclusion

The VGs work well. The 6 inches of H₂O across the cylinders is only achieved at high cruise, but the 4 inches at takeoff and low cruise is more than double the base values. Most importantly, the average CHT is down an average of 55 F, just by the addition of four VGs. The data in the table above is even better than I hoped for. I probably could improve on this further by moving the VGs to another location, but this is good enough.

I like this solution. It is elegant, simple, and effective. Other Velocity builders have had similar results, and the factory is adding them on some aircraft. There's no reason large VGs won't work on similar types of airplanes, too.

I don't know for certain, but as best I can measure, there is no speed penalty. Another plus is that it gives yet another area for people to ask questions about. Now it's off to the next thing I want to improve.