Cooling: The Inside Story

Part 2 of this series describes how your aircraft cooling system works for you.

BY BILL WELCH

In the first installment of this three-part series, we looked at the basics that should be considered when designing a cooling system for your homebuilt aircraft. But now that you have the cooling air onboard, what do you do with it?

That's a troublesome question for any designer, but thinking in terms of the basics helps you sort out your options. First, you must place the air on one side of the passage through which you want it to go, and it must be under higher pressure than the air on the other side. Next, the difference in pressures must be enough to move the required mass flow of air through the passage .

Obviously, if you have a moving stream of air pointed directly at the hole you want it to go through, it will. But in the case of a typical aircraft cooling system, you have a large number of small passages generally aligned in one direction, but scattered over an awkwardly large area. Rarely do they face the oncoming external airstream. In addition, there are a few odd passageways that stand off by themselves in no particular order. The problem is to maintain a flow of air through the entire maze of cooling channels.

The classic response to this problem is some form of *plenum* chamber. A plenum chamber is essentially a box where the air velocity has been reduced as much as possible to provide a reservoir of high pressure that is connected to all those little nooks and crannies that require cooling. The most common location of the plenum chamber is in the engine cowling just above the horizontal cylinders. It isn't the optimum location in all cases, but a detailed look at it illustrates the basic principles of cooling dynamics that apply across the board.

The first step in obtaining the increased static pressure you need in the plenum is to slow the entering air. Once you've captured a stream of air and confined it in a duct, the same amount moves past every cross section in a given time. Obvious? Of course it is. If you increase the width of the duct, it becomes a diffuser, and the air slows down as it moves, doesn't it? Yes, but not necessarily, at least not necessarily where you want it, and not without a substantial loss of the increased pressure you're looking for.

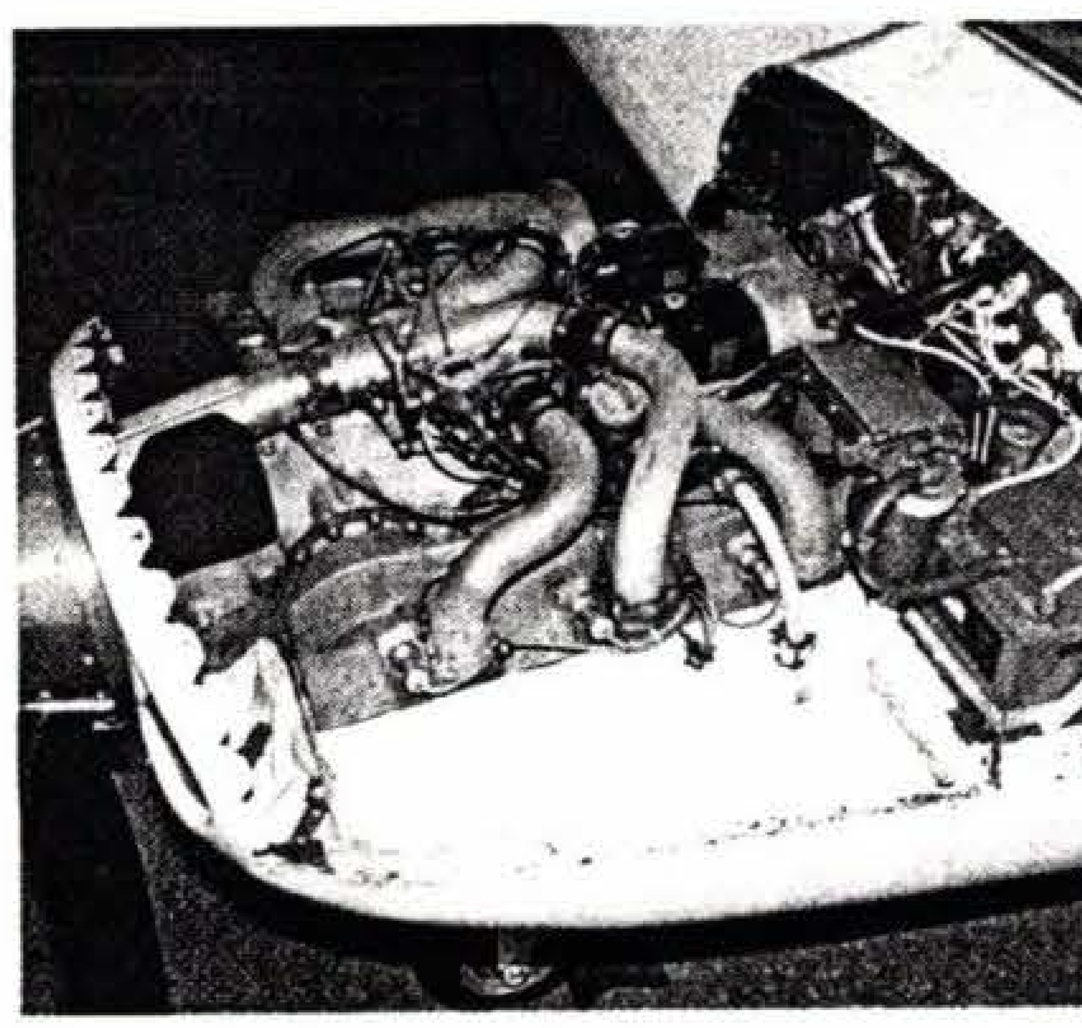

Questair Venture designers specified an impregnable, flexible sheet as part of the engine baffle. The system minimizes air leaks in the plenum.

How you slow the air controls the pressure available in the plenum. If you expand the duct too quickly, the airflow separates from the walls and it stalls. Much of its energy is squandered in useless vortices and other turbulence. Thus, pressure recovery is reduced by such features, the loss easily amounting to more than half of the free-stream dynamic pressure.

In aerodynamic tests, conical diffusers expanding faster than about a 7° half-angle show increasing pressure loss, a figure quoted by most writers. However, Dr. Sighard Hoerner, in his 1951 book, *Aerodynamic Drag*, pointed out that in locations where no appreciable boundary layer has formed ahead of the entrance, more rapid expansion of a conical diffuser can be used. The locations he is referring to are those where the boundary layer is actually diverted from the intake by design.

He illustrated the point with a diffuser that expanded to six times the intake area in about twice the diameter of the intake. The same expansion with a 7° half-angle would require a length 10 times the diameter of the intake.

In their 1957 book, *Applied Hydro- and Aeromechanics*, L. Prandtl and O. G. Tietjens pointed out that as early as 1904 Prandtl had suggested that airflow could be prevented from separating by removing the boundary layer. Photographs indeed showed that fully attached airflow could be maintained in sharply diverging experimental diffusers.

Boundary layer control, or even diversion, is a more complex problem than most builders would care to tackle, so they rely on simpler ways to maximize pressure recovery. One way is by controlling the duct area and expanding it as gradually as the physical dimensions of your design will allow. It's important to realize that expansion can be horizontal as well as vertical, and the less control you provide, the less pressure you'll have to work with.

One critically important step that is often neglected is the sealing of the baffles to prevent air from leaking out of the plenum through unwanted paths. Usually, the worst offender is the flexible seal between the removable cowl and the diaphragm baffle aft of, or alongside the cylinders. Attention to detail and the use of a backup to the seal, such as an angle on the removable part of the cowling, goes a long way toward maintaining the pressure. Other leaks occur at the attachments or joints in the baffles and the usually overlooked little holes that may or may not serve any purpose at all. As a rule, anything penetrating a baffle at least deserves a tight-fitting grommet to help seal it.

When the design is not attentive enough to detail to provide smooth transitions, the intentional exits from the plenum are another common source of pressure loss. The main path is the area between cylinder fins and barrels, with some flow also going around the

COOLING

continued

20% of its diameter allows it to draw air from all directions and eject a coherent stream that continues as a jet for many diameters. This difference between the properties of intake and exiting jets is the reason you don't suffocate in your own breath.

At this point, some acknowledgement of conflicting interests is necessary. Turbulence generally wastes energy in a flowing stream. However, in cooling a surface, if the stream of coolant (air or fluid) is laminar, only a thin boundary layer picks up heat directly from the surface. Small-scale turbulence induces a continuous turnover of the fluid in contact with the surface, increasing the amount of heat that is transferred and dissipated. Therefore, a certain amount of turbulence in the cooling air is good—even essential—to the proper functioning of the system.

The roughness of many surfaces, particularly sand-cast cylinder fins, creates some beneficial turbulence. But the majority of turbulence—usually far more than is necessary—is caused by the design of the cooling system itself. In general, it's better to use design features that create minimal turbulence.

For example, those simple sheet metal baffles atop the fins function better with rolled edges than with the customary raw edges. In most cases, the exit from cylinder fins probably should be baffled in order to maintain as much of the flow as possible in the desired path through the fins and around the barrel.

At this point, another factor enters the picture. Discounting the negligible effect caused by the differences in static pressure, when it is heated, air expands. While the mass of the airflow remains constant, the volume increases in proportion to the absolute temperature. Ideally, we could use some of the added heat energy to help move the air or even gain some useful thrust from it. In principle, that's what a ramjet engine does. Some carefully designed liquid-cooling systems have been able to use the added heat energy to produce a small amount of useful thrust. Even though it might do no more than offset part of the cooling drag, it is enough recovery to be useful.

A typical four-cycle gasoline engine is only about 30% efficient because between the exhaust and the cooling air, it dissipates into the environment twice as much energy as it delivers in propulsive thrust. Two-stroke engines are even less efficient—part of the reason for their high fuel consumption.

To recover heat energy as thrust, the expanded air must be ejected from the system at a velocity higher than its entry velocity. The increased kinetic energy is equal to the recovered heat, minus internal losses, and thrust is created as the sum of excess static pressure components acting forward on all the surfaces involved. This occurs primarily in the exhaust diffuser, where the area is adjusted to control the exhaust velocity, but it may occur in the form of reduced pressure on a forward-facing bellmouth.

If you like number-crunching, you can estimate the division of heat between exhaust and cooling (including the oil cooler), by computing the mass of the airflow through the engine from the displacement and rpm. Fuel consumption reflects the amount of heat released by burning. An engine that is running with the proper mixture for maximum power will dump its exhaust with a blue flame at about 1500°F, which begins to cool a few inches from the exhaust stack. However, some of the heat goes into the main cooling stream, depending on the length of exhaust stacks immersed in the stream. At cruise power and the correct fuel-air mixture, the first several inches of the stack will typically run cherry-red.

Absolute temperature on the Rankine scale is measured from −459.7°F. Standard sea-level air density is 0.0764 pounds/cubic feet (weight), 0.002377 slug/cubic feet. Another number needed for any calculations on cooling is the heat equivalent of power: One horsepower equals 0.70697 BTU/second (British Thermal Units). Heat added to a volume of air increases its total pressure in direct proportion to the absolute temperature before and after the heat is added. The temperature change can be computed from the specific heat of the gas. For air, this is approximately 0.17 BTU/°F/pound mass (slug) at constant volume, or 0.24 at constant pressure.

Another important consideration in the design of a cooling system is the supply of waste heat to the cabin or cockpit. It is almost universal practice to duct ventilating air around a hot exhaust stack to heat it, and then route the heated air directly to the cabin. This involves the obvious risk of introducing carbon monoxide from any leak in the exhaust stack, as the pressure in an exhaust stack is higher than in the surrounding air. For metals exposed to heat and the corrosive properties of combustion products—especially welded assemblies—this is a serious risk.

It is actually quite simple to eliminate the carbon monoxide threat. The pressure in an exhaust stack is usually well above the static pressures in the surrounding area. The FAA allows 2 inches of mercury above ambient pressure on type-certificated installations. By taking ventilation air in where most of the dynamic pressure can be recovered and passing it through a duct in a zone of reduced static pressure, such as the duct below the cylinders on a typical engine, you can ensure that any leakage of the ventilating duct is conducted outward, not into the system and the cabin.

Also, you can duct some of the cooling air already warmed by the cylinders around an exhaust stack to heat it further and then around the ventilating duct to heat the cabin air without risking carbon monoxide poisoning. The cost of this insurance is a small increase in the weight and complexity of your cooling system. But for anyone who is safety-conscious, it is a worthwhile price to pay to eliminate one of the major hazards of flying.

Although I've confined this discussion of cooling systems to the typical air-cooled, horizontally opposed engine installation, the same principles apply to any other arrangement and liquid-cooled systems as well. Some of the older inline engines are cooled by recovering dynamic pressure in a plenum on one side and passing the air sideways through the fins and out the rear on the other side.

Though somewhat heavier, liquid-cooled engines usually offer more streamlined nose and cowling configurations than air-cooled engines can, resulting in slightly better propeller efficiency. Further, they are easier to install for efficient cooling flow and some useful thrust recovery. The net result can be an advantage over the more common air-cooled engine installation.

With the increase in the popularity of two-cycle engines among homebuilders, a new tradeoff has become necessary. The increased weight of fuel required for a given flight range/duration has to be balanced against the greater weight of a four-stroke engine and any difference in propeller efficiency. Making the correct choice is not always obvious and is a matter of personal preference for each builder. □

cylinder heads. To limit the exits, most manufacturers and builders place simple sheet metal baffles against the cylinder fins. The exit area in each case should approximate the net area of the path around the cylinder between fins. Remember, the exit areas from the plenum control the allocation of airflow among all possible exit paths, so they should be adequate to facilitate the required distribution of air.

Air moving from an area of inactivity through an opening with sharp edges will detach (stall) from the edges and form vortices that drift downstream, greatly restricting the effectiveness of the opening and causing the stream to mix with the surrounding, relatively still air. Providing a simple bellmouth intake with a radius that's as small as

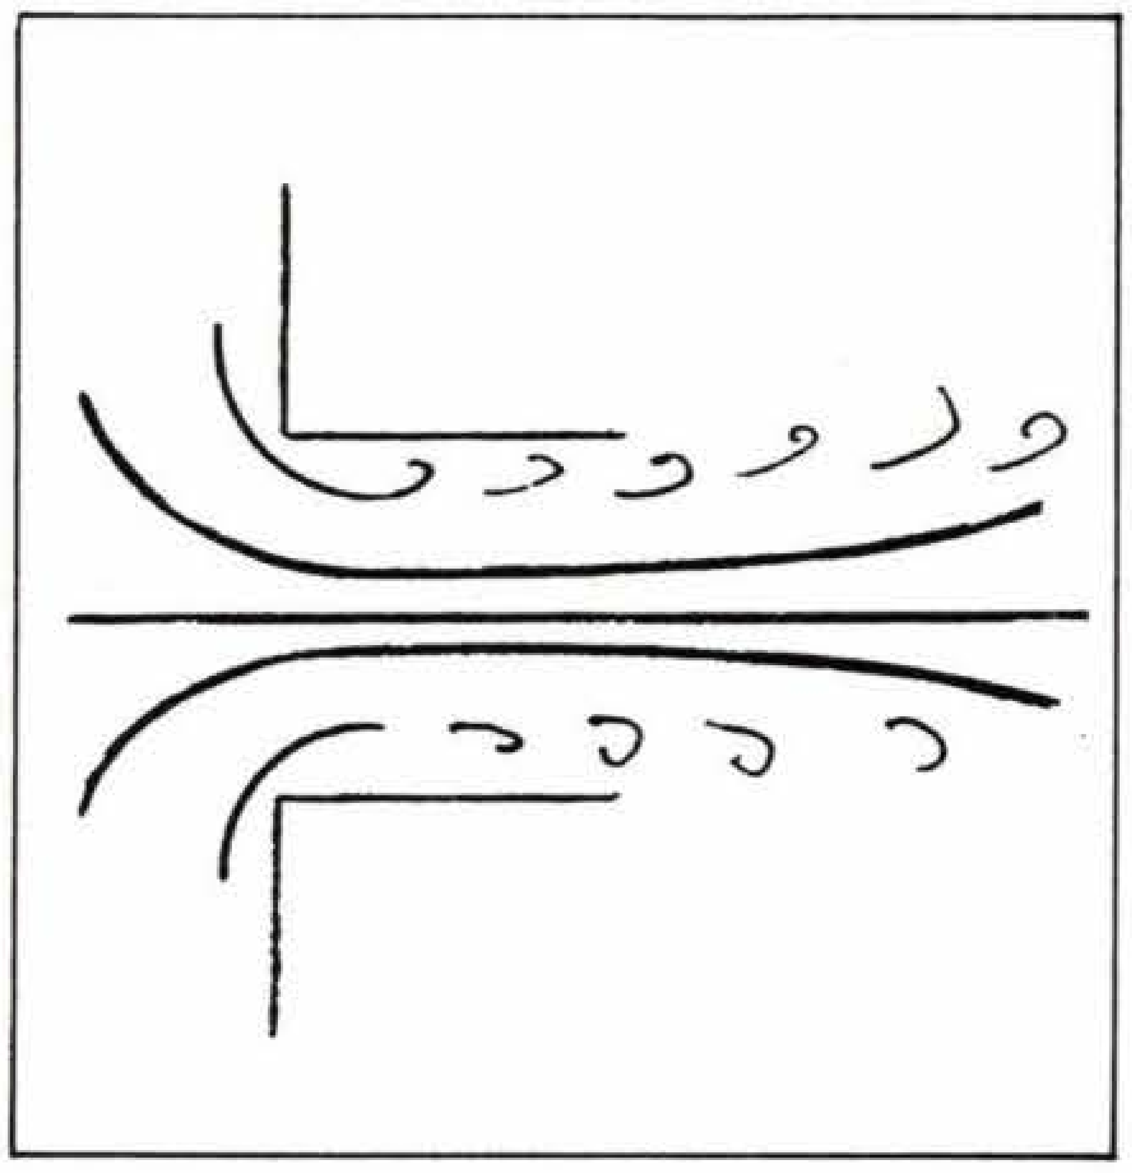

Figure 1. A sharp edge on an intake causes airflow to separate and mix rapidly with surrounding air, reducing the effectiveness of the inlet size.

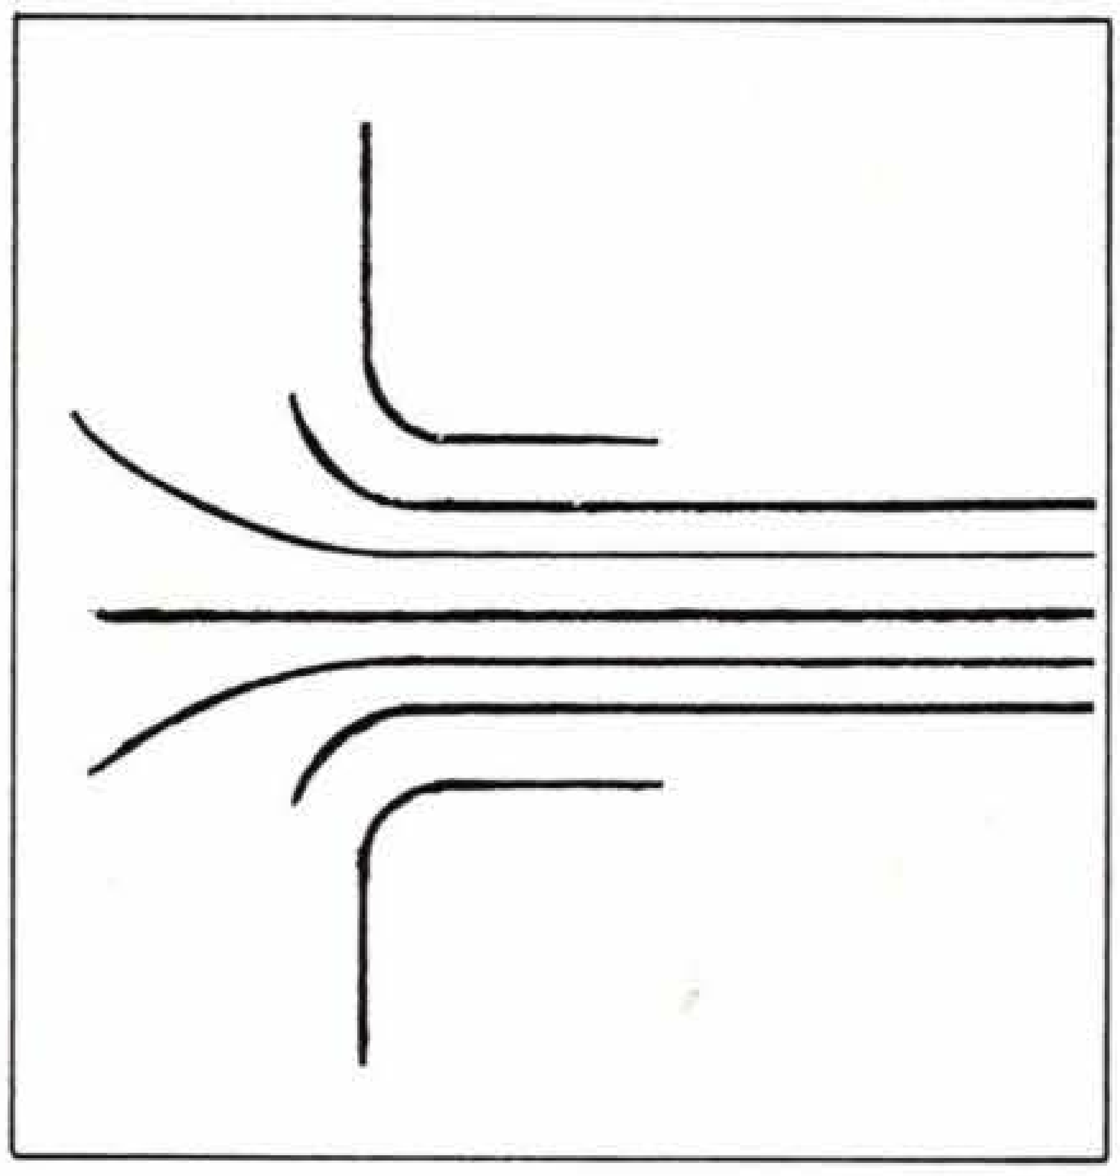

Figure 2. Rounding the intake edge allows full discharge through the ducting and enables the exiting air to flow freely with little mixing or burbling.

Figure 3. Attention to detail helps avoid excessive cooling drag of a typical lightplane. Cylinder baffles and diaphragm should fit and seal tightly. Rolled edges improve airflow and reduce pressure losses. Bottoms of fins should be baffled so that air flows as close as possible to the cylinder barrels.