

Future Prospects for Low-Pollution Combustion Engines

by FRANCIS H. CLAUSER

During the last three decades the crucial problems in the design of combustion engines have changed radically. Because of air pollution, exhaust emissions from combustion engines have become a dominant factor in engine design. It was Caltech's Arie Haagen-Smit who showed that sunlight produced the photochemical reactions of hydrocarbons and oxides of nitrogen that created the irritating oxidants in our southern California smog. And these, along with carbon monoxide, came principally from the combustion of vast amounts of hydrocarbons here in the Los Angeles basin.

Soon after the mechanism of smog formation was unraveled, it became clear that automobiles were major contributors. Early attempts to clean up the emissions were relatively easy, principally because cars in those days were so bad. The sixties and seventies have seen great progress in legislation governing air pollution, in the organizational problems of managing air quality, and in technological improvements to automobiles that have reduced their emissions by more than tenfold.

Now we have reached a point where improvements are more difficult to come by. Instead of tenfold improvements the struggle is to achieve improvements of tens of percent, and this raises the question of what the future holds. A few years ago optimists believed that the worst was over and that with continued improvements in combustion engines the smog problem would be solved. But the last few summers, with their long stretches of first- and second-stage alerts, have shown that we are barely holding our own.

Has our progress been slowed because we are approaching natural lower limits on the emissions from combustion processes? Or are the natural limits so low that we can expect to make improvements of tenfold several times over in the future? If the latter is the case then the optimists may yet be proven right. No one really knows what the

lower limits are on exhaust emissions, but I believe that improvements of tenfold several times over are indeed still possible.

The business of studying combustion emissions and combustion engines is still largely an empirical one. Great emphasis has been placed on reducing emissions by advancing or retarding the spark, making the fuel-air mixture leaner or richer, getting better carburetion, introducing exhaust gas recirculation, adding catalytic converters, and similar fixes. When I became interested in the problems, it seemed to me it would be foolhardy for a university to compete with a vast industry whose basic approach was to look only a few years ahead to the regulations they would have to meet. Inevitably, their emphasis was on the near-term future. What I could do — and I have found it a matter of keen interest — was to undertake research to ascertain what the fundamental limits of the problem are and how those limits can be approached.

As far as the thermodynamic side of combustion engines is concerned, we know the basic limits and how close we are to them. In fact, it was in elucidating these fundamental limits that Sadi Carnot got classical thermodynamics off to a start more than 150 years ago.

What is the scope of the emissions problem as we now see it? There are five pollutant emissions under legislative control — unburned hydrocarbons (UHC), carbon monoxide (CO), oxides of nitrogen (NO_x), particulates (smoke, or soot), and sulfur dioxide (SO₂). At present, only the first three are of significance here in southern California, and I will focus my attention on them. Sulfur dioxide is of concern in areas where fuels containing a significant amount of sulfur are burned, and particulates will probably be an increasing factor in the future.

To understand the emissions that come from a typical modern internal combustion engine, we first need to know a little of the technical jargon. Whenever any hydrocarbon

fuel such as methane, propane, or gasoline burns in air, there is a critical ratio of fuel to air such that each hydrocarbon molecule can find just enough oxygen partners to form CO_2 and H_2O with nothing left over. Chemists call this the stoichiometric ratio. Different hydrocarbons have different stoichiometric ratios, but for many purposes their behavior is best compared by using for each the relative fuel-air ratio, that is, the actual fuel-air ratio divided by the stoichiometric ratio. A rich mixture is one having excess fuel, and its relative fuel-air ratio is thus greater than one. A lean mixture correspondingly has a relative fuel-air ratio of less than one.

When a lean mixture burns, it does pretty much what is expected — the fuel burns to form principally CO_2 and H_2O , and the excess O_2 appears in the exhaust as just that, excess O_2 . However, when a rich mixture burns, the principal result is not that the excess fuel appears in the exhaust. Instead, there is a rearrangement of partners, and CO and H_2 appear in the exhaust along with H_2O and decreasing amounts of CO_2 . Only for very rich mixtures will significant amounts of fuel and carbon appear in the exhaust.

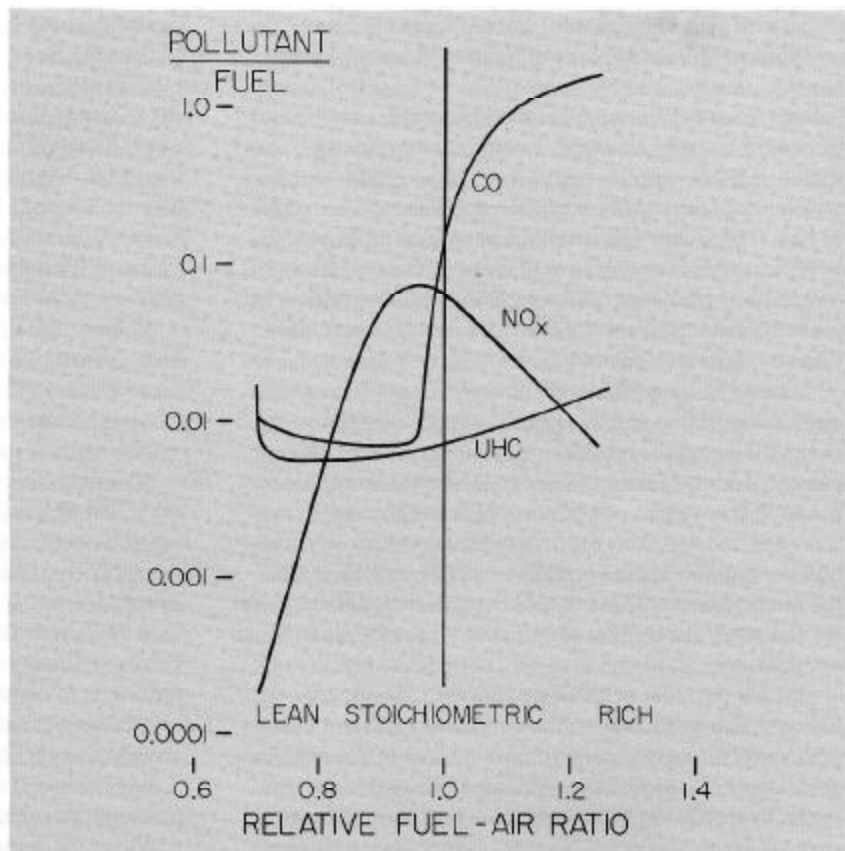
Of the three pollutants CO , UHC , and NO_x , only CO is a primary exhaust product and this only when the mixture is rich. With this one exception, all of the exhaust pollutants are only secondary products, and they occur in

surprisingly small quantities. It is this fortunate fact that permits us to burn hundreds of thousands of tons of hydrocarbons each day and have only hundreds of tons of pollutants discharged into our atmosphere.

In Figure 1, which depicts the emission characteristics of a typical gasoline engine, we have plotted the concentrations of UHC , CO , and NO_x in the exhaust, expressed as ratios to the amount of fuel entering the intake. Each varies greatly as the relative fuel-air ratio is changed. The only one that rises to large values is CO , which — as we remarked earlier — becomes a primary product when the mixture is rich. For lean mixtures, CO rapidly drops to low values and then abruptly levels off. The curve for NO_x rises to a peak on the lean side of stoichiometric and drops to quite low values for both very rich and very lean mixtures. In contrast, the UHC has a minimum fairly well into the lean region. It rises somewhat for rich mixtures, but it abruptly goes up almost vertically for mixtures so lean that the engine misfires. This happens because when firing does not occur, raw fuel becomes a primary exhaust product, appearing as unburned hydrocarbons.

In order to obtain a feeling for what constitutes a "good" engine as far as pollution is concerned, let us relate the levels shown in Figure 1 to the permissible federal and California levels for passenger cars. The federal standards set limits of 3.4 grams per mile for CO , and .4

Figure 1. The exhaust emissions of a typical gasoline engine are plotted here, with the horizontal axis representing the fuel-air ratio divided by the stoichiometric ratio. Along the vertical axis are the concentrations of the pollutants (unburned hydrocarbons, carbon monoxide and oxides of nitrogen) divided by the concentration of intake fuel; this scale is logarithmic (in powers of 10) and thus covers a very wide range of values.



grams per mile for UHC. Here in California we also have limits of .4 grams per mile for NO_x.

How are we to compare these values for vehicles with those we used in quoting engine emissions? For the engine we used the ratio pollutant/fuel, whereas vehicle levels were expressed as pollutant/distance. We can of course write the engine index, pollutant/fuel, as pollutant/distance times distance/fuel, and we recognize that distance/fuel is the familiar gasoline mileage, or miles/gallon of the vehicle. Unfortunately, we have here a strange conglomeration of units — miles, grams, and gallons. Inserting the proper factors to take this into account, we have

$$\frac{\text{(engine)} \text{ grams pollutant}}{\text{grams fuel}} = \frac{1}{3000} \times \frac{\text{(vehicle)} \text{ grams pollutant}}{\text{mile}} \times \frac{\text{miles}}{\text{gallon}}$$

Let us consider two extreme cases: a gas guzzler that gets 10 mpg and an economy car that gets 40 mpg. Using the above formula, we can translate the legislative vehicle limits into permissible engine limits as follows:

Gas Guzzler	Economy Car
$\frac{\text{UHC}}{\text{Fuel}}$ and $\frac{\text{NO}_x}{\text{Fuel}} = .0013$	$\frac{\text{UHC}}{\text{Fuel}}$ and $\frac{\text{NO}_x}{\text{Fuel}} = .0053$
$\frac{\text{CO}}{\text{Fuel}} = .011$	$\frac{\text{CO}}{\text{Fuel}} = .045$

If we apply Figure 1 directly to these numbers, we see immediately that a typical engine, when used in the economy car, can — by operating with lean mixtures — have acceptably low values for all three pollutants. But the gas guzzler hasn't a chance. It would need all the help it can get from a catalytic converter, exhaust gas recirculation, and any other "fix" available. For an engine to be good enough to get the gas guzzler within the standards it would have to be some four times better than the typical engine.

The pollutant scale of Figure 1 covers a range of more than 1000 to 1. It is thus a matter of some wonderment that engine performance and legal standards are currently so neatly matched that the line between pass and fail falls between the gas guzzler and the economy car, which differ in mpg by only a factor of four. This is no accident. Each time engineers have devised a technical improvement, legislators have lowered the limits in such a way that only the best pass.

Now let us look at the problem from a quite different vantage point. A simple hydrocarbon flame (Figure 2) performs fundamentally the same task as combustion in the gasoline engine — burning hydrocarbon. But here the process is far less complex. The flame burns smoothly and steadily in a uniform well-mixed mixture contained in a cylindrical quartz tube. At first glance the flame appears to be urged along by a glowing disk. Actually both the flame and the disk are standing still. It is the mixture that is moving in the tube, and the disk creates a flow field be-




Figure 2. This free-standing flame, produced by a well-mixed mixture of propane and air entering a vertical quartz tube from below, provides an excellent setup for measuring the fundamental characteristics of combustion emissions. In the experiment, the curved flame is positioned by a disk-shaped flame holder. The flow is smooth and steady so the flame stands motionless, striking the quartz tube in a wavy ring around the boundary.

hind the flame that anchors it in place. Such a setup is well suited to laboratory experiments because the free-standing flame can be probed in detail to reveal what reactions are taking place.

How do the emissions from this free flame compare with those from combustion in an engine cylinder? Analogous to the emission levels in Figure 1, we show in Figure 3 the corresponding values for the free flame. A number of startling differences are immediately apparent. The flame has a UHC level that is 1000-fold lower than that of the engine, and this low value extends over a much wider range of fuel-air ratios, including a significant portion of the rich-mixture region.

If we probe the free flame to determine what happens to the hydrocarbon fuel as it passes through the flame, we are immediately struck with what a wonderful thing it is. In front of the flame the hydrocarbon concentration is several percent, or several tens of thousands parts per million (ppm). Immediately behind the flame this drops to a level of a few tenths of a part per million. This abrupt drop of 100,000-fold shows how efficient the flame is in consuming the hydrocarbons.

To put this in perspective, here in southern California the air we breathe usually contains one or two ppm of


hydrocarbons. In contrast, the exhaust from the free flame has a concentration one-tenth as great. If the engine were as efficient as the flame, our cars would cleanse the air of hydrocarbons instead of contributing to the pollution. Clearly something significantly different is happening in the two cases.

For NO_x emissions, the flame is also significantly better than the engine. Its peak values are approximately one-tenth those for the engine. For CO, in the rich region where CO is a primary product, the flame and engine are, as expected, nearly the same. As the mixture becomes lean, the CO from the flame does not drop as rapidly as for the engine, but instead of leveling off it continues to drop to quite low values.

These comparisons raise a host of questions. Why is the engine worse than the free flame? Does this flame represent the best that can be done; or is it possible to go to even lower levels? Might it be possible to build a useful engine that could approach the levels of the free flame, or achieve even lower values?

In seeking answers to these questions we have to look at the mechanisms that lead to the formation of each of the pollutants individually. First we examine the situation for CO. In the free flame, there is a luminous front that is about one millimeter thick, where the primary reactions take place. First the hydrocarbon molecules break up; almost immediately the carbon atoms unite with oxygen atoms to form CO. If the mixture is rich (that is, oxygen deficient), many of these CO molecules cannot find other oxygen partners, and CO emerges as a primary exhaust product. If the mixture is lean, the CO molecules at a more leisurely pace seek additional oxygen partners (usually stealing them from OH radicals) to form CO_2 . However this second union doesn't always stay stuck. Temperature has a large effect, both on how fast this second union takes place and how fast it comes apart. Both reactions slow down as the temperature decreases, so much so that for ordinary temperatures the time constants are measured in acons. Furthermore, the ratio of the two rates shifts, favoring CO_2 over CO, so that at ordinary

Figure 3. Exhaust emissions from the free flame shown in Figure 2 are plotted here on the same scales used in the Figure 1 graph for the gasoline engine. Striking differences are immediately apparent.



temperatures CO should virtually disappear.

With this picture in front of us we can formulate general principles for obtaining low levels of CO. A prime requisite is that all combustion should be with lean mixtures. At first glance it would appear that the leaner the better. This seems to be borne out by the CO curve in Figure 3, which continues to drop as the mixture becomes leaner. But a closer analysis reveals that the principal effect of leanness here is to lower the flame temperature, shifting the equilibrium from CO to CO₂. Any other means for reducing the flame temperature would also lower the CO level.

This immediately suggests a way of obtaining even lower concentrations of CO than those in the free flame. In almost every combustion process, after combustion has taken place, the temperature of the products is reduced. In some engines this is done by expanding the gases to extract mechanical energy. In others, the heat is transferred to other media (water, for instance). This lowering of the temperature can be used to reduce the CO concentrations, but to be effective it must be done in a carefully prescribed way. As the temperature drops, not only does the equilibrium shift but also the rate of approach to equilibrium drops as well. At elevated temperatures, equilibration times are measured in milliseconds, but as the temperature falls they become seconds, then hours, and finally centuries. If at any point in the cooling process the rate of cooling gets ahead of the equilibration rate, the CO will be "frozen in" and the benefits of further cooling lost. Below 1200K equilibration rates become impossibly slow. For temperatures such that the times are reasonable, however, CO concentrations of the order of five ppm can be reached. On the scales of Figures 1 and 3 this is a level of approximately .0001 for CO/fuel, which is below any of the values for either the engine or the free flame. Here we have a case where it should be possible to build an engine that would do even better than the free flame.

Now let us turn our attention to the problem of NO_x emissions. In the air around us, nitrogen and oxygen have existed together without detectable interaction for billions of years. It's not that they won't react, but rather that such reactions, like the CO reactions, are extremely temperature dependent. Here, very little happens until temperatures of the order of 2000K are reached, and then O₂ and N₂ dissociate into N and O, and a set of reactions gets under way that leads to the formation of NO and NO₂. As the temperature goes up, the pace of these reactions increases rapidly, so rapidly that it appears as though a critical threshold has been passed.

In the case of the free flame, this critical temperature is exceeded only for mixtures with relative fuel-air ratios between .75 and 1.5. In Figure 3, we see that lean and rich flames both produce very little NO_x (that is, either NO or NO₂). But the NO_x curve rises to a peak near a relative fuel-air ratio of one.

For the engine a similar behavior occurs, but there the

peak of the NO_x curve is some ten times higher. This is caused by the much higher temperatures that exist in the engine for two reasons. First, the mixture is compressed before burning, which raises its temperature significantly. Second, the combustion takes place at nearly constant volume rather than constant pressure, which causes an additional increase in temperature.

At this point we see that a portion of the strategy we outlined for reducing CO will also apply in the case of NO_x, namely that of keeping the combustion temperatures low. In the first case, after the rapid initial formation of CO, the slower oxidation of CO to CO₂ is favored by lower temperatures. In the second, lower temperatures inhibit the formation of NO_x.

It is in the cooling portion of the strategy that a discrepancy in goal occurs. For CO, a properly paced cooling rate will cause the CO to be oxidized to CO₂. In the case of NO, the attack of atomic O on the N₂ molecule promotes the formation of NO_x. It is the attack of atomic N on NO_x, however, that causes the destruction of NO_x. Unfortunately, upon cooling atomic N disappears considerably faster than atomic O. This means that it is virtually impossible to get rid of NO_x by a program of cooling. For NO_x, the best strategy is simply never to let the temperature get high or, if it does, to cool it off before the relatively slow NO_x formation reactions have had time to occur. Fortunately, the procedures for obtaining low CO and low NO_x are not completely contradictory. The useful range of temperatures for oxidizing CO lie between 1200K and 2000K, whereas NO_x forms at a significant rate only above 2000K. Hence a proper procedure is to keep the burning temperature below 2000K and then to cool at a properly tailored rate down to 1200K. In this way it should be possible to obtain concentration levels of the order of five ppm or less for both CO and NO_x. On the scales of Figures 1 and 3 this translates into values of approximately .0001 for CO/fuel and NO_x/fuel. This number is more than 100 times lower than those required to meet the CO standard for the gas guzzler we considered earlier, and 13 times less than the NO_x standard. For the economy car, the numbers are even more impressive.

Now we come to the last of the three major pollutants — unburned hydrocarbons. We have seen that in a typical engine UHC levels are 1000 times greater than in a simple flame. Why is this so? For the UHC to appear in the exhaust, some of the incoming fuel must escape being burned, not only during combustion itself, but also during the subsequent expansion and expulsion from the cylinder.

In the early 1960s Wayne Daniel of the General Motors Research Laboratories proposed two ways in which this might occur. In the first the fuel-air mixture enters into crevices such as that around the top of the piston above the piston rings. This takes place during the intake and compression strokes. After combustion, these unburned hydrocarbons emerge during the expansion stroke and are carried out in the exhaust. The second mechanism proposed

by Daniel was that of wall quenching. Since the time of Sir Humphry Davy it has been known that a flame will not pass through a wire screen if the mesh is fine enough. The flame is quenched by heat being drawn away by the screen. Daniel similarly reasoned that a flame near the cooled wall of the piston and cylinder would be quenched and thus unable to burn a thin layer of combustible mixture lying along the surface. This unburned layer would at a later stage be swept off the surface and carried into the exhaust.

Tests by Wentworth of G.M. and by others showed that in fact the crevice around the top of the piston above the rings could be a major source of unburned hydrocarbons. When this crevice was eliminated in tests, the exhaust hydrocarbons decreased by more than half. And then a strange thing happened. Attention became focused on the wall-quenching problem, and nothing was done about the piston crown device. Cars in production today still have that same crevice.

A few years ago two thoughts occurred to me, since my background is in fluid mechanics, and I had come to believe that these quenching phenomena were of a boundary layer kind that can involve significant dynamic behavior. The first thought was that there might be a third mechanism for producing unburned hydrocarbons, namely, the quenching of the flame at the boundary between the fresh incoming charge and the residual charge left over from the preceding cycle. The second was that fundamental research on wall quenching might be worthwhile because the evidence for its existence was largely inferential and no one seemed to have examined in detail what was actually taking place. So I undertook research on both of these subjects.

In the wall-quenching experiments, I set up a situation not unlike that shown in Figure 2 except that I used a water-cooled metal wall instead of a quartz tube. With small probes I was able to explore in detail the regions both ahead of and behind the flame, and in particular the region along the wall. Now according to the wall-quenching hypothesis there should be a layer about a millimeter thick of unburned hydrocarbons next to the wall behind the point where the flame strikes the wall.

Much to my surprise there was no such UHC layer. As mentioned earlier, the flame reduces the hydrocarbon level from tens of thousands of parts per million in the oncoming stream to tenths of parts per million in the region after the flame. Contrary to expectations, this state of affairs was found to hold right up to the wall. In spite of whatever quenching that may be taking place, the flame is able to consume the hydrocarbons near the wall to the same low level that it does in the free stream.

Upon taking stock, I realized that in my original experiment both the wall and the flame were stationary in space, with the combustible mixture flowing through the flame front. In an engine the flame sweeps across the cylinder and moves with respect to the wall. To determine if the



Francis Clauser sets up his laboratory experiments on flame quenching as a source of unburned hydrocarbons.

movement of the flame relative to the wall was affecting the result, I curved the flow channel in such a way that the cooled wall could be part of a large aluminum wheel. Here, with the flame striking the wall, I could have the wall moving in either direction with any speed, essentially duplicating the relative velocities of wall and flame that occurred in an engine cylinder. For small speeds of the wall relative to the flame, there was still no layer of quenched hydrocarbon at the wall. As the speed of the wall increased, however, either with or against the flow direction, such a layer did begin to appear, and at large speeds it was comparable in size to that predicted by the empirical theory (which contained no notion that speed would be a factor).

These experiments show that wall quenching does play a role in producing UHC, but less of a role than formerly believed. Perhaps more importantly, they show that wall quenching can be eliminated in engines that are designed in such a way that the flame front is stationary, or nearly so, with respect to the surrounding wall (as occurs in steam engines, stirling engines, and gas turbines). This almost certainly will prove to be an important factor affecting the design of engines of the future.

In the second series of experiments, those on UHC pro-




Figure 4. In this experiment the setup is similar to that in Figure 2 (the rod is the same size) but with the central stream surrounded by an annulus of incombustible gas also rising vertically. The flame in the combustible cone spreads outward and upward until it strikes the incombustible stream where it is quenched, leaving in its wake a surprisingly large amount of fuel that emerges as unburned hydrocarbons.

duced by incombustible gas quenching, the setup shown in Figure 4 was employed. Here the combustible mixture coming from below is surrounded not by a quartz tube as in Figure 2, but by an annulus of incombustible gas (such as exhaust gas or air) that is also moving upward. When the flame strikes the boundary between the combustible core and the surrounding annulus, it is extinguished. The question is how much UHC will be left over along the boundary after the combustion is over. The results were dramatic. Hydrocarbons were found to exist in large quantities along the boundary in the wake of the flame. All of this indicated that incombustible gas quenching can in fact be a major source of UHC in combustion engines. A second part of this experiment is now under way. We have designed and built an engine that almost completely scavenges the residual charge after each cycle. We hope to find out soon if this can lead to a significant reduction in exhaust hydrocarbons.

As yet we do not know the full story of the origin of all of the exhaust hydrocarbons. Some as yet unknown source mechanisms may be found. In the laboratory, however, we have built combustion chambers that yield UHC levels comparable with the values of the free flame, indicating that it may be possible to build an engine having these

same remarkably low values.

Again, why is the typical gasoline engine so much worse than the free flame? The gasoline engine, like all combustion engines, performs three basic thermodynamic acts. It compresses its working medium, here a mixture of fuel and air; it adds heat to this medium, in this case by combustion of the medium; and finally it expands the medium to extract useful mechanical work. The gasoline engine, in company with the diesel, does all three acts in the same chamber. This puts rather severe time constraints on the combustion and post-combustion processes. At an engine speed of 3000 rpm, the combustion itself must be completed in about two milliseconds and the subsequent expansion in about ten. When flame temperatures are high, combustion is fast, but as temperatures are reduced, by leaning the mixture or dilution with exhaust gas, combustion slows to rates well below those needed to meet the two-millisecond requirement. This constitutes a major obstacle in reducing temperatures enough to insure very low concentrations of NO_x . Further, the restricted expansion times militate against obtaining the low concentrations of CO that are possible in a properly paced cooling program.

Because of the repetitive nature of combustion in the compression chamber of the gasoline engine, ignition must be initiated each cycle, and the flame must sweep across the chamber at relatively fast speeds. As a result there appears to be no way to avoid UHC produced by wall quenching since the flame cannot be held stationary relative to the cooled walls. With the present configuration of gasoline engines it may be found that the intermixing of the residual gases from one cycle with the fresh mixture of the next leads to quenching that cannot be avoided. If so, this too will prevent the achievement of very low levels of UHC. In addition, there is evidence that both wall quenching and residual gas quenching cause the CO to freeze in before it has been oxidized to CO_2 .

All of the above stand as significant obstacles that make it difficult for the gasoline engine to achieve very low pollution levels. What are the prospects for other combustion engines? The whole class of "external" combustion engines, which includes the steam engine and the Stirling engine, has a significant advantage as far as low pollution is concerned. Since the combustion process takes place independently of the compression and expansion processes, much greater freedom exists in design to meet the requirements for low pollution. If this is so, why haven't these engines demonstrated such potential? Recently I measured the emissions from the large steam boilers here at Caltech. I found that both the CO and the UHC were in fact reasonably close to the very low values I have indicated above should be possible. In contrast, it was the NO_x that was unacceptably high. In these boilers, combustion takes place at a leisurely rate. Neither wall quenching nor incombustible gas quenching is a problem. Heat is extracted over a long flow path. As a result, both the hydrocarbons




Figure 5. Setting a bad example, this flame is typical of a large number of current combustion devices that, because of ease and simplicity, mix and burn at the same time. Here the air and fuel enter separately from below and begin to burn as they begin to mix. In the central bright stream, the mixture is extremely rich (almost entirely fuel), releasing a great deal of carbon as smoke; the burning mixture in the surrounding sheath is quite lean because little fuel has diffused out that far. In between these two regions, burning is occurring with a broad range of fuel-air ratios, producing large amounts of exhaust pollution at the unfavorable ratios (where the curves reach high levels in Figure 3).

and the CO are oxidized to concentrations of a few parts per million.

I believe the reason for the high NO_x lies in an affliction that plagues a wide variety of combustion devices, namely, the act of mixing and burning at the same time. This affliction has a long history, dating back to the camp fires of earliest man and continuing today in the vast majority of combustion devices we use.

The problem lies in the wide range of fuel-air ratios that occur in a mix-and-burn flame. In the emissions for the free flame shown in Figure 3 we see that for each of the pollutants there are certain fuel-air ratios that produce unusually high values of the pollutant. Clearly these fuel-air ratios should be avoided if at all possible. But consider a flame such as that shown in Figure 5. Here the fuel and the air are introduced separately, and they burn as they mix. If we examine what is happening locally, we find that near the core of the fuel stream, combustion is taking place at very rich mixtures. Out in the airstream combus-

tion is taking place at very lean mixtures. In between, all other fuel-air ratios can be found. For each pollutant there are local regions where that pollutant is being produced in unacceptably large quantities. Clearly it does no good to keep pollution levels low elsewhere if some regions are producing excessive quantities.

Now let us reverse the thrust of the question and ask why mix-and-burn is so widely used if it is so bad. The alternative, of course, is to mix the fuel and air thoroughly and then, and only then, to burn it. The difficulty is that, given half a chance, the flame will either strike back into the mixing region or, worse, be blown out. For the premixed flame, control becomes a major issue. In contrast, the mix-and-burn configuration is a highly stable one that requires little if any control. In nature's scheme of things, mix-and-burn is the lowest of combustion denominators, the configuration of last resort.

The Caltech steam boilers also have burners of the mix-and-burn type. Within the main fire box large quantities of all the pollutants, including soot, are produced. Fortunately for all except NO_x , the process of extracting heat and cooling the mixtures is so leisurely and the intensity of mixing so great that they are eventually oxidized to very low levels. Initial local hot spots, however, produce relatively large amounts of NO_x . As we saw earlier, cooling and mixing do not reduce NO_x significantly. If the Caltech boilers are a reasonably good case study, then the indications are that "external" combustion power plants have good chances of achieving very low pollution levels if they can master the problems of premixed burning.

Finally, let us take a brief look at the problems of jet engines and gas turbines. Here too combustion takes place as a separate and specialized function, so these engines should enjoy many of the advantages of the external combustion engines. Further, the temperature limitations of the turbine are low enough to make the achievement of very low values of NO_x quite favorable. However, all conventional gas turbines and jet engines employ mix-and-burn systems that hinder the achievement of these potentially low levels of NO_x . Experimental jet engine combustors using burners with premixed fuel and air have demonstrated that significantly lower levels of CO, UHC, and NO_x are possible. Here again, the problem is one of preventing strike back and blowout over the wide operating range of the jet engine.

And what of the engine of the future? Almost certainly we shall soon be able to build combustion systems that use premixed mixtures having precise fuel-air ratios and that can operate stably over a wide range of conditions. With this accomplished, exhaust emissions should experience a dramatic decrease. Engine designers can then employ a wide variety of compressors and expanders to assemble engines that can be tailored for many different uses. All should be able to enjoy exhaust emissions levels powers of ten less than our present engines. And we should be able to have our cars and breathe clean air too. □