

## Performance

Since I've been talking about "performance" for some thousands of well chosen (I hope) words, it might not be a bad idea for me to explain, at this point, exactly what I mean by the word.

The object of a rocket motor is to produce thrust—a force. This it does by ejecting a stream of gas at high velocity. And the thrust is dependent upon two factors, the rate at which the gas is being ejected, in, say, kilograms per second, and the *velocity* at which it is ejected. Multiply rate by velocity and you get thrust. Thus, kilograms per second times meters per second gives the thrust in Newtons. (That is, if you're a man of sense and are working with the MKS *Système Internationale* of units.) If you want to increase your thrust you can do it either by increasing the mass flow (building a bigger motor) or by increasing the jet velocity, which generally means looking for a better propellant combination. The performance of a propellant combination is simply the jet velocity it produces.

Sometimes people not in the rocket business ask what is the "power" of, say, the Saturn V rocket. Power isn't a very useful concept in rocketry, since what you're trying to give your vehicle is momentum, which is proportional to the thrust times the time it is exerted. But if you define the power as the rate at which thermal or chemical energy is being converted to kinetic energy in the exhaust stream, a meaningful figure can be dug out. The kinetic energy of a given mass of exhaust gas (relative to the rocket, that is, not to the Earth or the Moon or Mars) is  $Mc^2/2$ , where  $M$  is the mass, and  $c$  is the velocity (again, relative to the rocket). And the power, or rate of energy conversion,

is  $\dot{M}c^2/2$ , where  $\dot{M}$  is the mass flow—kilograms per second, say. But, as we saw above  $\dot{M}c = F$ , the thrust. So, putting these together,  $\text{Power} = Fc/2$ . Nothing simpler. Let us now proceed to Saturn V.

Saturn V has a thrust of 7,500,000 pounds *force*. Not mass, mind you; the distinction is important. That is equal to  $33.36 \times 10^6$  Newtons. (One pound force = 4.448 Newtons, the MKS unit of force. That's a nice thing about MKS—there's no confusion between mass and force!) I don't remember the exact exhaust velocity of the Saturn engines, but it can't be very far from 2500 meters per second. So, multiply  $33.36 \times 10^6$  by  $2.5 \times 10^3$  and divide by two—and out comes the power, neatly in watts.

And the power so calculated is

$$41.7 \times 10^9 \text{ Watts}$$

or  $41.7 \times 10^6$  Kilowatts

or  $41.7 \times 10^3$  Megawatts,

which amounts to some 56 *million* horsepower. For comparison, the nuclear powerplant of the *Enterprise*, the most powerful afloat, generates some 300,000 HP. And the mass flow of propellants into the engines and exhaust gases out of the nozzles is some fifteen tons a second. Considered as the through-put of a chemical reactor—which it is—the figure is impressive.

So far, everything has been simple. But now things begin to get a little sticky. For the question arises, "How do you calculate the exhaust velocity,  $c$ , that you can get out of a given pair of propellants, burned at a definite chamber pressure, and properly expanded through a nozzle?" As we saw above, the energy of a given mass of exhaust gas,  $E = \dot{M}c^2/2$ . Rearranging this,  $c = (2E/\dot{M})^{1/2}$ . As all of the propellant injected into a motor comes out as exhaust gas (we hope!), the " $\dot{M}$ " in that equation is also the mass of the propellant which produced the mass of the exhaust gas that we're considering. But the  $E$  is *not* equal to the thermal energy,  $H$ , in the exhaust gas before it was expanded. So, actually,  $c = (2H/\dot{M} \times \eta)^{1/2}$ , where  $\eta$  is the efficiency of conversion of thermal to kinetic energy. And  $\eta$  depends upon the chamber pressure, upon the exhaust pressure, and upon the nature of the exhaust gas, both as it exists unexpanded in the chamber and as it changes during expansion.

So, obviously, we have to know the chemical composition of the gas in the chamber. That's the first step. And you can't take it by using simple stoichiometry. If you put two moles of hydrogen and one of oxygen into the chamber you do *not* come out with two of water. You

will have  $H_2O$  there, of course. But you will also, because of the high temperature, have a lot of dissociation, and the other species present will be  $H$ ,  $H_2$ ,  $O$ ,  $O_2$ , and  $OH$ . Six species in all, and you can't know, a priori, in what proportions they will appear. And to solve for six unknowns you need six equations.

Two of these are simple. The first is derived from the atomic ratio between hydrogen and oxygen, and simply states that the sum of the partial pressures of all the hydrogen-bearing species, each multiplied by the number of hydrogen atoms in it, all divided by the sum of the partial pressures of all of the oxygen bearing species, each multiplied by the number of oxygen atoms in it, is a certain value upon which you have already decided, in this case two. The second equation states that the sum of the partial pressures of all the species present shall equal the chamber pressure which you have chosen. The other four equations are equilibrium equations of the type  $(H)^2/(H_2) = K_1$  where  $(H)$  and  $(H_2)$  represent the partial pressures of those species, and  $K_1$  is the constant for the equilibrium between them at the chamber temperature. This is a very simple case. It gets worse exponentially as the number of different elements and the number of possible species increases. With a system containing carbon, hydrogen, oxygen, and nitrogen, you may have to consider fifteen species or more. And if you toss in some boron, say, or aluminum, and perhaps a little chlorine and fluorine—the mind boggles.

But you're stuck with it (remember, I didn't *ask* you to do this!) and proceed—or did in the unhappy days before computers. First, you make a guess at the chamber temperature. (Experience helps a lot here!) You then look up the relevant equilibrium constants for your chosen temperature. Devoted and masochistic savants have spent years in determining and compiling these. Your equations are now before you, waiting to be solved. It is rarely possible to do this directly. So you guess at the partial pressures of what you think will be the major constituents of the mixture (again, experience is a great help) and calculate the others from them. You add them all up, and see if they agree with the predetermined chamber pressure. They don't, of course, so you go back and readjust your first guess, and try again. And again. And eventually all your species are in equilibrium and you have the right ratio of hydrogen to oxygen and so on, and they add up to the right chamber pressure.

Next, you calculate the amount of heat which would have been evolved in the formation of these species from your propellants, and compare that figure with the heat that would be needed to warm the combustion products up to your chosen chamber temperature. (The same devoted savants have included the necessary heats of formation

and heat capacities in their compilations.) And, of course, the two figures disagree, so you're back on square one to guess another chamber temperature. And so on.

But all things come to an end, and eventually your heat (enthalpy) all balances, your equilibria all agree, your chamber pressure adds up, and you have the right elemental ratios. In short, you know the chamber conditions.

The next morning (the procedure described above has probably taken all day) you have to make a decision. Shall you make a frozen equilibrium calculation, or shall you make a shifting equilibrium calculation? If the first, you assume that the composition of the gas and its heat capacity remains unchanged as it is expanded and cooled in the nozzle. If the latter, you assume that as the gases cool and expand the equilibria among the species shift in accordance with the changing pressure and temperature, so that neither the composition nor the heat capacity of the exhaust gas is identical with what it was in the chamber. The first assumption amounts to a statement that all reaction rates are zero, the second to a statement that they are infinite, and both assumptions are demonstrably false.

If you want a conservative figure, you choose to make a frozen equilibrium calculation. (It gives a lower value than a shifting equilibrium calculation.) And you plug the data from the chamber calculations into the following horrendous formula.

$$c = \left\{ 2 \frac{R\gamma}{\gamma - 1} \frac{T_c}{\bar{M}} \left[ 1 - \left( \frac{P_e}{P_c} \right)^{\frac{\gamma-1}{\gamma}} \right] \right\}^{1/2}$$

Here,  $R$  is the universal gas constant,  $\gamma$  is the ratio of specific heats,  $C_p/C_v$  of the chamber gases.  $\bar{M}$  is their average molecular weight.  $T_c$  is the chamber temperature.  $P_e$  and  $P_c$  are the exhaust and chamber pressures respectively. This formula looks like a mess, and it is, but it can be simplified to

$$c = \left[ 2H/\bar{M} \right]^{1/2} \left[ 1 - \left( \frac{P_e}{P_c} \right)^{R/C_p} \right]^{1/2}$$

where  $H$  is the sum of the enthalpies of all the species present. (The reference state of zero enthalpy is taken to be the perfect gas at absolute zero.) " $\bar{M}$ ," of course is the mass of propellants which produced them. And the efficiency,  $\eta$ , is

$$1 - \left( \frac{P_e}{P_c} \right)^{R/C_p}$$

If you feel optimistic—and energetic—you make a shifting equilibrium calculation. This is based on the assumption that although the gas composition will change during the expansion process, the entropy will not. So your next step is to add up the entropies of all the species present in the chamber, and put the figure on a piece of paper where you won't forget it. (Entropies are in the compilations, too.) Then, you guess at the exhaust temperature, at the exhaust pressure you have decided upon. And then you determine the composition of the exhaust gas, just as you did the chamber composition. And add up the entropies, there, and compare it with the chamber entropy. And try another exhaust temperature, and so on. Finally you have the exhaust conditions, and can calculate the enthalpy per unit mass there. And then, finally,  $c = \left[ \frac{2(H_c - H_e)}{M} \right]^{1/2}$ ,  $\eta = (H_c - H_e)/H_c$ .

Solid and liquid exhaust products complicate the process somewhat when they appear, but that's the general idea. There is nothing complicated about it, but the execution is insufferably tedious. And yet I know people who have been doing performance calculations for twenty years and are still apparently sane!

The time and labor involved in an "exact" performance calculation had two quite predictable consequences. The first was that those calculations which *were* made were cherished as fine gold (for shifting equilibrium calculations read "platinum"), circulated, compiled, and squirreled away by anyone who could get his hands on them. The second consequence was that everybody and his uncle was demanding an approximate, or short method. And these were forthcoming, in considerable variety.

The most elaborate of these took the form of Mollier charts of the combustion products of various propellant combinations. These usually plotted enthalpy versus entropy, with isotherms and isobars cutting across the chart. A typical set of charts would be for the combustion products of jet fuel with various proportions of oxygen. Another, the decomposition products of 90 percent peroxide, another, ammonia and oxygen, at various O/F ratios. Some were more general, applying to a defined mixture of carbon, oxygen, hydrogen, and nitrogen atoms, without specifying what propellants were involved. These charts were easy to use, and gave results in a hurry, but they seldom applied to *exactly* the combination you had in mind. They were also very difficult to construct, involving, as they did, dozens of calculations. The Bureau of Mines, with its extensive experience with combustion phenomena, was a leader in this field.

A more general but less informative method was developed in 1949

by Hottel, Satterfield, and Williams at MIT. This could be used for practically any combination in the CHON system, but using it for any chamber pressure other than 300 psia, or any exhaust pressure other than 14.7, was an involved and messy procedure. I later modified and streamlined the method, and made some provision for other elements, and published it in 1955 as the "NARTS Method of Performance Calculation."

These, and similar graphical methods, involve, essentially, an interpolation between accurately calculated systems, and they gave a fairly good approximation of the results of a shifting equilibrium calculation.

The other group of methods gave, generally, results that approximated those of a frozen equilibrium calculation, and were based on the equation  $c = (2H/M \times \eta)^{1/2}$ . The usual procedure was to determine H by ignoring any minor products (pretending that there wasn't any dissociation). The products in the CHON system were assumed to be CO<sub>2</sub>, H<sub>2</sub>O, CO, H<sub>2</sub>, and N<sub>2</sub>. Once the water-gas equilibrium was determined (that was done by using the equilibrium constant at some arbitrary temperature, such as 2000 K, or at the whim of the operator – it didn't matter too much) \* H could be determined by simple arithmetic. As for  $\eta$ , with a little experience you could make a pretty good guess at it, and any error would be halved when you took the square root of your guess! Or, if you wanted to be fancy, you could determine the average C<sub>p</sub> of your gases at somewhere near what you thought your chamber temperature ought to be, and plug that into the efficiency term. Tom Reinhardt's 1947 method included curves of temperature *vs* enthalpy for various exhaust gases, as well as C<sub>p</sub> *vs* temperature. You determined your temperature from your enthalpy, and the C<sub>p</sub> from the temperature. The temperature, of course, was much too high, since dissociation was ignored. Ten years later I modified the method, eliminating the curves, devising a fast and easy way of getting an R/C<sub>p</sub> averaged over the whole temperature range, and providing a nomograph for calculating  $\eta$  from that and the pressure ratio. It was called the NQD – NARTS Quick and Dirty – method. The thing worked astoundingly well, giving results agreeing with complete shifting equilibrium calculations (I suppose that the averaged R/C<sub>p</sub> helped there) to something like 1 percent. And you could make a calculation in fifteen minutes. It worked best, too, when you postulated the simplest – in fact the most simple minded – set of

\* Consider the case where one O<sub>2</sub>, one H<sub>2</sub> and one C react. If the reaction went to H<sub>2</sub>O + CO, the performance would vary by only 2.5 percent from the performance if it went to CO<sub>2</sub> and H<sub>2</sub>. And this is the worst possible case!

products imaginable. And it was adaptable. When a man from Callery Chemical Co. came in one day and told me for the first time about the BN system I learned that. In this system the exhaust products are hydrogen and solid BN. I hauled out my tables when he told me about it, and letting two atoms of carbon (graphite) pretend that they were one of molecule of BN, made a fast estimate. And lit on the nose. My value was within half a percent of the one he had obtained from a fancy machine calculation. The only trouble with the method was that I never could keep a copy for myself. Some character was always mooching my last copy, and I'd have to run off another fifty or so.

There were other approximate methods developed, some as late as 1963, but they were all similar to those I've described. But the day of the shorthand method is gone – as is, thank God! – the complete hand calculation.

The computers started getting into the act in the early 50's, although considerable chemical sophistication was needed to make the most of their initially somewhat limited capabilities. At Bell Aerosystems they were considering fluorine as an oxidizer, and a mixture of hydrazine and methanol as the fuel, and demanded performance calculations. The programmer protested that he couldn't handle that many elements, and Tom Reinhardt retorted, "The carbon and the oxygen will go to CO, and you just tell the little man who lives inside that box to treat it exactly like nitrogen." End of problem.

All the compilations of thermodynamic data are on punch cards, now, versatile programs, which can handle a dozen or so elements, are on tape, and things are a lot simpler than they were. But the chemical sophistication is still useful, as is a little common sense in interpreting the print-out. As an example of the first, calculations were made for years on systems containing aluminum, using thermodynamic data on gaseous  $\text{Al}_2\text{O}_3$  calculated from its assumed structure. And the results didn't agree too well with the experimental performances. And then an inconsiderate investigator proved that gaseous  $\text{Al}_2\text{O}_3$  didn't exist. Red faces all over the place. As an example of the second, consider the case of a propellant combination that produces a lot of solid carbon, say, in the exhaust stream. The machine makes its calculations on the assumption that the carbon is in complete thermal and mechanical equilibrium with the gaseous part of the exhaust. A bit of common sense suggests that this will not be so, since heat transfer is not an infinitely rapid process, and that the carbon may well be exhausted considerably hotter than the surrounding gas. So you look at the print-out with considerable pessimism – and wait for experimental results before committing yourself.

A great deal of effort, in recent years, has gone into attempts to develop programs which will take things like heat transfer from solid to gas into account, and which will allow for the actual velocity of the change in the exhaust composition during expansion. These are called "kinetic" programs, as opposed to the frozen or shifting equilibrium programs, and only the big computers make them possible. There is only one trouble with them. Reliable kinetic data are as hard to come by as honest aldermen – and when you feed questionable data into the machine, questionable results come out at the other end. As the computer boys say, "Garbage in – garbage out."

And there is one disconcerting thing about working with a computer – it's likely to talk back to you. You make some tiny mistake in your FORTRAN language – putting a letter in the wrong column, say, or omitting a comma – and the 360 comes to a screeching halt and prints out rude remarks, like "ILLEGAL FORMAT," or "UNKNOWN PROBLEM," or, if the man who wrote the program was really feeling nasty that morning, "WHAT'S THE MATTER STUPID? CAN'T YOU READ?" Everyone who uses a computer frequently has had, from time to time, a mad desire to attack the precocious abacus with an axe.

Rocket performance is not usually reported in terms of exhaust velocity, although the early workers wrote in those terms. Instead, it is reported as "specific impulse," which is the exhaust velocity divided by the standard acceleration of gravity, 9.8 meters or 32.2 feet per second<sup>2</sup>. This practice gives figures of a convenient size in the range of 200 to 400 or so, but it has led to some rather tortuous, if not ludicrous definitions. The most common one is that specific impulse is the thrust divided by the *weight* flow of propellant, and it comes out in seconds. Putting the acceleration of gravity into the equation did that, but specifying the performance of a rocket, whose whole job is to get away from the earth, in terms of the acceleration of gravity on the surface of that planet, seems to me to be a parochial, not to say a silly procedure. (The Germans, during World War II, used an even sillier measure of performance, "specific propellant consumption," which was the reciprocal of specific impulse. This didn't even have the virtue of producing figures of a convenient size, but gave things like 0.00426 per second.)

Probably the best way of thinking of specific impulse is as a velocity expressed, not in meters or feet per second, but in units of 9.8 meters (or 32.2 feet) per second. That way you retain the concept of mass flow, which is relevant everywhere, and doesn't depend upon the local properties of one particular planet, and at the same time lets Euro-



pean and American engineers understand each other. When he hears  $I_s = 250$ , the European multiplies by 9.8 to get the exhaust velocity in meters per second, while the American does the same with 32.2 and comes out with feet per second. (*When will the U.S. ever change over to MKS?!*)

I've told you what performance is, and I've described the way you go about calculating it. But now comes the practical problem of picking a propellant combination which will give you a good one. Here it will be helpful to go back to the velocity equation,  $c = [2H/M]^{1/2} [1 - (P_e/P_c)^{R/C_p}]^{1/2}$  and to consider the  $H/M$  term and the efficiency term separately. Obviously, you want to make  $H/M$  as large as possible. And to do this, it is useful to consider the exhaust gases you hope to get. The energy contributed by a molecule of combustion products equals the heat of formation of that molecule from its elements at 25°C, plus its sensible heat above absolute zero (this is a very small item) *minus* the energy required to break down to their elements, at 25°C, the propellants which formed it. This last term is generally much smaller than the first—otherwise we wouldn't have useful propellants. And sometimes it is negative; when a mole of hydrazine breaks down to hydrogen and nitrogen we get some twelve kilocalories as a free bonus. But the important item is the heat of formation of the product molecule. That we want as big as possible. And, obviously, to maximize  $H/M$ , we must minimize  $M$ . So, to get a good energy term, we need an exhaust molecule with a high heat of formation and a low molecular weight.

So far so good. But now let's look at the efficiency term. Obviously, we want to get it as close to 1.0 as possible, which means that we want to beat  $(\frac{P_e}{P_c})^{R/C_p}$  down as far as we can.  $P_e/P_c$  is, of course smaller than one, so to do this we must raise the exponent  $R/C_p$  as high as we can. Which, of course, means that we want exhaust products with as low a  $C_p$  as we can find. And so we are hunting for exhaust products which have:

- a. A high heat of formation.
- b. A low molecular weight.
- c. A low  $C_p$ .

Alas, such paragons among exhaust products are hard to come by. Generally, if you have a good  $H/M$  term, the  $R/C_p$  term is bad. And vice versa. And if both are good, the chamber temperature can get uncomfortably high.

If we consider specific exhaust products, this is what we find:  $N_2$

and solid C are practically useless as energy producers. HCl, H<sub>2</sub>, and CO are fair.\* CO<sub>2</sub> is good, while B<sub>2</sub>O<sub>3</sub>, HBO<sub>2</sub>, OBF, BF<sub>3</sub>, H<sub>2</sub>O, and HF, as well as solid B<sub>2</sub>O<sub>3</sub> and Al<sub>2</sub>O<sub>3</sub>, are excellent. When we consider the R/C<sub>p</sub> term, the order is quite different. The diatomic gases, with an R/C<sub>p</sub> above 0.2, are excellent. They include HF, H<sub>2</sub>, CO, HCl, and N<sub>2</sub>. (Of course a monatomic gas has an R/C<sub>p</sub> of 0.4, but finding a chemical reaction which will produce large quantities of hot helium is out of the range of practical politics.) The triatomic gases, H<sub>2</sub>O, OBF, and CO<sub>2</sub>, with an R/C<sub>p</sub> between 0.12 and 0.15 are fair. The tetratomic HBO<sub>2</sub> and BF<sub>3</sub>, at about 0.1, are poor, and B<sub>2</sub>O<sub>3</sub>—well, perhaps it should be passed over in silence. As for the solids, C, Al<sub>2</sub>O<sub>3</sub>, and B<sub>2</sub>O<sub>3</sub>, their R/C<sub>p</sub> is precisely zero, as would be the thermal efficiency if they were ever the sole exhaust products.

Faced with this situation, all the rocket man can do is hunt for a reasonable compromise. He would, if he could, choose pure hydrogen as his exhaust gas, since at any given temperature one gram of hydrogen has more heat energy in it than a gram of any other molecule around (one gram of H<sub>2</sub> at 1000 K has almost ten times the energy of one of HF at the same temperature), and its excellent R/C<sub>p</sub> makes it possible to use a large fraction of that energy for propulsion. So hydrogen is the ideal working fluid, and you always try to get as much of it as possible into your mix. For it has to be a mix (in a chemical rocket, anyway) since you need an energy source of some sort to heat that hydrogen up to 1000 K or 3000 K or whatever. And the only available energy source is the combustion of some of the hydrogen. So you bring some oxygen or fluorine into the picture, to burn part of the hydrogen to H<sub>2</sub>O or HF, bringing the temperature up to 3000 K or so, and your exhaust gas is the mixture of H<sub>2</sub>O or HF with the excess hydrogen. When hydrogen is the fuel, it is always used in excess, and never burned completely to water or HF. If it were, the chamber temperature would be uncomfortably high, and the R/C<sub>p</sub> of the mixture would be lowered and the performance would drop. Hydrogen is so light that a considerable excess of it won't harm the H/M term appreciably, and you get the maximum performance, generally, when you use only enough oxygen or fluorine to burn perhaps half of your fuel.

If you're burning a hydrocarbon with oxygen, or if you're working

\* The classification of hydrogen, as a fair contributor of energy even though it, naturally, has a zero heat of formation, is explained by the fact that the molecule is so light. At 25° it has a sensible heat, or heat content of 2.024 kilocalories per mole above absolute zero, and since the molecular weight is only 2.016, its H/M, even at room temperature, is 1.0 Kcal/gm.

with the CHON system in general, you generally get the maximum performance from a mixture ratio which gives a 1.05 to 1.20 ratio of reducing to oxidizing valences in the chamber—that is, you work a little on the rich side of stoichiometric to get some CO and H<sub>2</sub> into the mixture and improve R/C<sub>p</sub>. (“Rich” and “lean” in the rocket business mean exactly what they do in a carburetor.)

If you're using a halogen oxidizer with a storable fuel, the best results generally show up if your mixture ratio makes the number of fluorine atoms (plus chlorine atoms, if any) exactly equal the number of hydrogen atoms. If there is any carbon in the combination, it's a good idea to get enough oxygen into the system to burn it to CO, so you won't have any solid carbon in the exhaust. And if your energy-producing species is a solid or liquid at the exhaust temperature—BeO, Al<sub>2</sub>O<sub>3</sub> are examples—the thing to do, of course, is to cram as much hydrogen as possible into the combination.

These are just a few of the things that the propellant chemist has to consider when he's looking for performance. And coming up with propellant combinations which will perform as the engineers want them to is what he's paid for. Inadequately.

This is how he goes about it: The engineering group have been given the job of designing the propulsion system of a new surface-to-air missile—a SAM. It is specified by the customer that it must work at any temperature likely to be encountered in military operations. The maximum dimensions are fixed, so that the missile will fit on existing launchers. It must be a packaged job, loaded at the factory, so that propellants won't have to be handled in the field. It must not leave a visible trail, which would make countermeasures easier. And, of course, it must have a much higher performance than the present system, which burns acid-UDMH. (The customer probably makes a dozen more demands, most of them impossible, but that will do for a starter.)


The engineers, in turn, before sitting down to their drawing boards, demand of the propellant chemist that he produce a combination that will make the missile do what the customer wants it to do. They also add some impossible demands of their own.

The chemist crawls into his hole to consider the matter. What he'd like to recommend is the hydrazine-chlorine pentafluoride (for historical reasons, ClF<sub>5</sub> is generally called “compound A”) combination. It has the highest performance of any practical storable combination known (all the exhaust products are diatomic, and 2/3 of them are HF), and it has a nice fat density, so you can stuff a lot of it into a small tank. But he remembers that all-weather constraint, and reminds

himself that you can never tell where you might have to fight a war, and that the freezing point of hydrazine is somewhat incompatible with the climate of Baffin Land. So—the next best bet is, probably MHF-3, a 14–86 mixture of hydrazine and methyl hydrazine with the empirical formula  $C_{0.81}H_{5.62}N_2$ . Its freezing point is down to the magic  $-54^\circ$ . (There are other possible fuels, but they may be somewhat dangerous, and he *knows* that MHF-3 is safe, and works.) But, with  $ClF_5$ , MHF-3 would leave a trail of black smoke leading right back to the launcher—definitely undesirable if the crew of the latter want to live to fire another round. Also, his professional soul (it's the only soul he has left after all these years in the business) is revolted by the thought of that free carbon and its effect on the  $R/C_p$  term and what it will do to his performance.

So he decides to spike his oxidizer with a bit of oxygen to take care of the carbon. Which means spiking it with an oxygen containing storable oxidizer. The only one of these which can live with compound A is perchloryl fluoride, "PF." So PF it will be.

He knows that when you have carbon and hydrogen in your system, along with oxygen and fluorine and chlorine, you generally get the best performance when the oxygen and carbon balance out to CO, and the hydrogen and the halogens balance to HF and HCl. So he doodles around a bit, and comes up with the equation:



That looks good—lots of HF and hence a lot of energy. And there's nothing but diatomic gases in the exhaust, which means a good  $R/C_p$ , which means, in turn, that a gratifyingly large fraction of that energy will go into propulsion. To find out what that fraction will be, he packs up his notes and pays a call on the IBM 360. The results of the consultation are pleasing, so he converts his mole fractions into weight percentages, and calls on the engineers.

"Your fuel is MHF-3," he announces, "and your oxidizer is 80 per cent 'A' and 20 PF. And your O/F is 2.18. And Muttonhead says—" "Who's Muttonhead?" "Muttonhead's the computer. He says that the performance, shifting, at 1000/14.7 pounds is 306.6 seconds, and *I* say that if you can't wring out 290 on the test stand you're not half as good as you say you are. But watch your O/F. If you're lean the performance will drop off in a hurry, and if you go rich you'll smoke like crazy: The density is 1.39, and the chamber temperature is 4160 K. If you want it in Fahrenheit, convert it yourself!"

He then retreats hurriedly to his lair, pursued by the imprecations of the engineers, who, (a) complain that the density is too low, and, (b) that the chamber temperature is much too high and who ever heard of anybody operating that hot anyway? (c) demand that he do something about the toxicity of  $\text{ClF}_5$ . To which he replies that (a) he'd like a higher density himself, but that he's a chemist and not a theologian and that to change the properties of a compound you have to consult God about it; (b) to get high performance you need energy, and that means a high chamber temperature, and unless they're satisfied with RFNA and UDMH they'll have to live with it, and for (c) see the answer to (a).

And then, for the next six months or so he's kept busy telling them, in response to complaints:

"No, you can't use butyl rubber O-rings with the oxidizer! Do you want to blow your head off?"

"No, you can't use them with the fuel either. They'll go to pieces."

"No, you can't use copper fittings with the fuel!"

"Of course, your mixture ratio goes off if you put five gallons of the oxidizer in a fifty-gallon tank! Most of the PF is up in the ullage, and most of the A is down in the bottom of the tank. Use a smaller tank."

"No, there isn't any additive I can put in the oxidizer that will reduce the vapor pressure of the PF."

"And no, I can't repeal the first law of thermodynamics. You'll have to talk to Congress!"

And he dreams wistfully of climbing into a cold Martini—and wonders why he ever got into this business.