



Introduction to hybrid design

The following is a short introduction to hybrid rocket engines, with some examples of the current state of hybrid research around the world.

Words in **bold** appear in the Glossary at the end of the document.

In line with rocketry convention, we'll sometimes use Newton's fluxion notation: using a dot

above a letter to denote derivative with respect to time, i.e. $\dot{r} = \frac{dr}{dt}$

Introduction

A hybrid, as the name suggests, is an engine that shares components from both solid-propellant rocket and liquid-propellant rocket design.

It has a tank supplying liquid or gaseous propellant into a combustion chamber that contains a solid propellant.

The conventional 'forward' hybrid stores liquid oxidiser in the tank, and solid fuel within the chamber.

With a 'reverse' hybrid, the fuel is the liquid in the tank, and solid oxidiser is in the chamber.

Traditionally, solids tend to be called 'rocket *motors*', whereas liquids are labeled 'rocket *engines*'; you can use either term for a hybrid, though Aspire tend towards 'engine.' (or 'propulsor' if we're being pretentious.)


The Boundary Layer

In a hybrid engine, the burning takes place within the boundary layer of flow affixed to the inner wall of the central port.

The flow of fluid down a pipe/fuel grain can be *imagined* as split into two separate regions:

In the central region, the fluid is all flowing at the same flow speed, but in a concentric outer region, the flow speed drops with distance above the wall of the pipe, coming to rest as it touches the wall.

In reality, the flow speed is different right across the pipe, but it reaches 99% of the central flow speed only a few millimetres from the wall, and that is defined as the boundary between the regions.



This outer 'boundary layer' is, and appears, stuck to the inner pipe wall.

You can visualise a radial slice through the boundary layer in a very wide pipe as a pack of cards lying on a table, where successive cards in the deck represent successive plane, parallel layers of fluid at increasing distance from the wall of the pipe (from the surface of the table).


Drop the pack onto a table but with a sideways movement. Successively lower cards in the pile move forwards with lower velocity; the lowest card remains stationary on the table. This is exactly how the layers of fluid within the boundary layer behave.

This nice, ordered mental picture of parallel laminate layers, known as 'Laminar flow' needs modified somewhat:

For one thing, each layer gets thicker as it flows down the pipe, as if each card was thicker at the downstream end than at the other. This is termed 'boundary layer growth'.

Also, at some point along the length of the pipe, the boundary layer breaks up; what

determines this 'transition' to chaotic, or 'turbulent' behavior isn't desperately well known.



A turbulent boundary layer is much thicker than a laminar boundary layer, and can be thought of as a fractal series of many little eddies within eddies, the separate layers of fluid within the boundary layer get thoroughly mixed together.

If the flow in a laminar boundary layer hits a step or cliff, it will trip to turbulent flow, though the boundary layer will 'trip' turbulent all by itself simply after having had to flow over the walls for a long enough period of time. You can see this effect clearly in the smoke rising from a cigarette: the smoke is initially laminar, but after a critical distance above the cigarette it suddenly goes turbulent, breaking up into thick eddies. This distance is typically a few centimeters down a hybrid **port**.



Comparison of three main chemical rocket types

Solids

Just like a burning fuse, solid rocket motors burn simply by combusting an intimate pre-mix of fuel and oxidiser.

In essence, it's a big firework, and shares the same characteristic that once lit, it can't be extinguished until all the propellant has been consumed.

Solids suffer from the bad combination of a **regression rate** (see glossary) that is dependant on the gas pressure within the chamber, and the fact that every exposed surface of the propellant will burn.

If the flame encounters a crack in the propellant, then that's extra surface area to burn, and so the rate of gas production, and hence the combustion chamber pressure, starts to rise. The rising pressure causes the regression rate to increase, producing gas more rapidly, and hence the chamber pressure rises more. Meanwhile, the crack is eroding larger.

This positive feedback mechanism can be terminal: either the chamber pressure exceeds its safe limit, known as an 'over-pressure' (the chamber bursts) or worse, the burning rate firstly goes supersonic, known as 'detonation', and then the chamber *really* over-pressures.

Why is the regression rate pressure dependant? Because many solid motor fuels incorporate metals, and when metalised fuels burn, the burning particles are highly emissive in the infrared, whereas the gasses produced by non-metallic fuels are fairly transparent. This thermal radiation causes an extra radiative heat transfer, of a magnitude that depends upon the number of metal particles per cubic metre in the port flow, so it depends upon the port flow pressure (as it affects the density).

'Safe' solid motor design is to attempt to utterly avoid any cracks or cavities, and also the incorporation of a structural fuse: a large plug (perhaps the whole nozzle) breaks free at a set over-pressure and de-pressurises the chamber as quickly as possible. The recoil may destroy the vehicle however, and further crack the propellant.

Liquids

Liquid rocket engines pipe the liquid fuel and liquid oxidiser from their separate tanks and through different channels within the **injector**. The injector then sprays them into each other as they enter the combustion chamber.

The injector is an expensive precision device: it has to let the two propellants through at the correct rate for **Stoichiometric** combustion, which means that its multitude of small holes ('orifices') must be the correct diameter to very high tolerance to meter the flowrates accurately.

Furthermore, pairs or triplets of fuel and oxidiser streams are typically aimed to splash into each other just as they enter the combustion chamber in order to create small droplets, which requires accurately drilled and polished holes, and rows of complex alternate internal channels of fuel and oxidiser.

If the injector malfunctions, this is often terminal: perhaps the propellant feed system lost pressure, allowing back-flow of the hot combustion chamber gasses into the internal channels of the injector, or the injector cracked or leaked internally or melted for some other reason. In anycase, the fuel and oxidiser channels are physically close together within the injector, separated, due to restricted internal space, by thin walls. Any internal mixing, and a fire is possible which melts it all the more.


Once the injector goes, the feed-pipes from the separate tanks are inevitably in close proximity, and the fire consumes the pipework back towards the tanks, which may then rupture.

Liquids can also suffer from 'hard-starts' which are high chamber pressure pulses, possibly causing an overpressure. These occur if an unburnt pool of mixed fuel and oxidiser builds up within the chamber due to a delayed ignition: such a pool usually detonates, and the bigger the pool, the bigger the bang.

Hybrids

The hybrid we are most familiar with is the log fire: the **regression rate** depends upon the draught (**mass flux**) of gas passing over the surface of the solid fuel, which we can increase, akin to using bellows or opening-up the flue on the chimney, to increase performance.

In a simple forward hybrid, oxidiser, such as nitrous oxide, flows down the single central hole (or **port**) running down the center of a thick-walled pipe of solid propellant fuel, such as plastic.



Once the hybrid has been ignited, some of the hot combusted gas is carried via the turbulence within the boundary layer to the surface of the solid propellant where its heat is conducted into the top layer of this fuel. (a process known as 'forced convection'; radiation plays a minor role.) The heat then melts a layer of fuel off of the exposed surface of the fuel charge, and what happens next depends upon the type of fuel used.


'Melt' isn't quite the right phrase with plastic fuel-grains, because the long polymerised hydrocarbon molecules in plastics are too long and entangled ever to have solidified completely at low temperature like the nice, ordered atoms of a solid metal: they never really froze. Instead, they're supercooled liquids with phenomenally high viscosity as with glass, but once heated above a critical 'Glass Transition Temperature', they will flow like a liquid, albeit treacle. As the surface is heated further, the remaining liquid vaporizes to gas.

If the fuel is plastic, it may char or slough off tiny chunks before melting, and then it de-polymerises: the carbon-to-carbon bonds holding the backbone of the polymers together absorb too much heat and break apart, and the fuel turns into a mixture of several gasses and simpler hydrocarbons, this is known as pyrolysis.

The melt layer

An extremely thin 'melt layer' appears to rest on the surface of the fuel, though in fact this liquid layer is constantly being replenished from underneath, and removed from above.

If the viscosity of the melt layer is particularly low (i.e. caused by the melting of some other fuel that isn't a polymer) then it breaks up into a series of ripples and roll waves due to the flow of gas above it; an unstable process modeled for water waves by Craik in the 1960's.





Flecks of liquid then break off the crest of these waves and are carried into the main flow like the spray off whitecap waves on the sea during a storm. These droplets then pass through the flame zone where they heat-up, vaporize, and burn with the oxidiser.

This 'sea-spray' effect is a secondary transport mechanism that gets a fair percentage of the fuel and oxidiser mixed, and occurs at moderate temperature. It's also not yet terribly well known in the hybrid community because it doesn't occur for plastic or rubber fuels. Despite early promise, this secondary mechanism has proven difficult to replicate in practice; more usually a large depth of fuel grain melts into one goeey mess.

With further heating, the remaining fuel in the melt layer, or the cracked hydrocarbons in the case of polymers, vaporize to gas and get carried off into the main flow. This is the primary transport mechanism that gets the fuel and oxidiser mixed, occurs at high temperature, and is the one that most hybrid rocketeers are familiar with.

Once transported into the port flow, the vaporized fuel burns in a narrow zone within the boundary layer known as the 'flame zone' or 'flame sheet' (a 'macro diffusion flame'), which is at a depth between the fuel-grain below and the oxidiser flow above that is close enough to **stoichiometric** to support combustion.





Rocketeers have to be careful how they define combustion efficiency in the hybrid flame zone or they'll predict overly pessimistic results. *Where burning occurs*, i.e. in the zone, it is pretty much stoichiometric. Between the solid grain and the flame zone there is mainly unburnt fuel, whereas in the central core of flow it is a mix of combusted gas and unburnt oxidiser.

Whether all the fuel will get a chance to meet and burn with all of the oxidiser before exiting out the nozzle is an aerodynamic issue which we'll discuss later.

Hybrid rocket combustion is inherently safer than solids or liquids (see our 'hybrid safety' paper for the definition of safer) for several reasons.

- The liquid-propellant half of a hybrid motor's parentage gives it the great advantage over a solid motor that the oxidiser supply can be shut off in an emergency, stopping the combustion.
- The fuel and oxidiser are kept as far apart as possible, with the wall of the combustion chamber between them, and therefore should only meet within the combustion chamber.
- The injector only handles one propellant, making it much simpler and incidentally allowing cheaper tolerances. Though it can melt and burn like the rest of the pipework, it is not prone to internal fires.
- The burning zone occurs within the boundary layer and so stands free of the fuel surface. A crack in the **fuel grain** has to be large enough, has to lead somewhere (form a channel right across the grain) and have a difference in pressure at both ends, to cause a flow through it. Without a flow the liquid propellant can't get into the crack and it therefore won't burn.
- Hybrid **regression rates** are virtually insensitive to chamber pressure (*provided* that you don't dope the solid grain with metal particles, which turns it back towards being a solid motor.)
- The solid grain is mechanically robust, therefore a malfunctioning hybrid that hits the ground won't explode. It did bury itself a long way down though!

Although safer, hybrids can still misbehave (see our 'hybrid safety' paper).

Modelling the hybrid regression rate

The **regression rate** that the solid propellant is eroding at a point x metres down the **port** of the hybrid is typically a function of the **mass flux 'G'** of gas that the grain is suffering at that point along the port. G is simply the total of the fuel mass flux plus the oxidiser mass flux:

$$G_{total}(x) = G_{fuel}(x) + G_{ox}$$

I'll use the term liquid propellant plus solid propellant instead of fuel plus oxidiser, as this covers both forward and reverse hybrids, but take this to mean gasses that were once liquid propellant that entered the top of the port via the injector, plus gasses that were progressively eroded and vaporised from the solid propellant port.

Thus, in a hybrid, the fuel vaporized depends on the **mass flux** of liquid propellant flowing down the port, and so maximizing $G_{liquid_propellant}$ by increasing its **mass flow rate** (by opening up the injector), and also by decreasing the port diameter, will maximise thrust.

There is a limit, however, to the maximum mass flux attainable. If too high then the flame will literally be blown out.

The maximum liquid propellant flowrate occurs on startup because the chamber is still at low pressure, and so the pressure difference between tank and chamber is at maximum. So the more energetic the ignition system, the higher flux the motor can take.



After considerable development, Aspirespace now have an ignition system that will handle extremely high startup fluxes. The trick is to site the ignition system just downstream of the injector so that it decomposes the nitrous to release hot oxygen that then meets the fuel grain.

The primary transport mechanism

A typical equation for the **regression rate** of a plastic fuel grain (i.e. no secondary transport mechanism) burning with liquid oxidiser is:

$$\dot{r}(x) = a G^n(x) x^m$$

where a , n , and m are constants, and x is the distance down the port.

The **mass flow rate** of fuel vaporized is then:

$$\dot{m}_{fuel}(x) = \rho_{fuel} \dot{r}(x)$$

per square meter of port surface area 'S' at the x of interest, and this is integrated along the length of the port to get the total mass flowrate of fuel heading for the nozzle:

$$\dot{m}_{fuel} = \int_0^x \rho_{fuel} \dot{r}(x) S(x) dx$$

The above equation for $\dot{r}(x)$ suggests that the diameter of the port will increase (regress) much more rapidly at the downstream end (larger x) than the upstream end. However, the boundary layer (whose thickness partially blankets the fuel-grain from the heat of combustion) grows larger with x as well, so the diameter of the port grows only slightly larger along its length.

Furthermore, the burning is self-levelling: any 'humps' in the fuel-grain increase the local erosion rate around themselves to smooth themselves out.

For these reasons, much more useful expressions for comparison with test firing results can be based on the average values along the port at some time point of interest during the burn:

$$\dot{r}_{average} = a G_{average}^n L^m$$

where x has now been replaced by L , the total port length, and:

$$G_{average} = \frac{\dot{m}_{total}}{A_{average}}$$

i.e. based on an average measurement of the port cross-sectional area.

The values of a , n and m are now (slightly) different; typical values for this equation for a Lox-plastic hybrid are: (from AMROC) $n = 0.76$, $m = -0.24$, (values of a are always jealously guarded.)

The mass flow rate of fuel vaporized by the end of the port is then: $\dot{m}_{fuel} = \rho_{fuel} \dot{r}_{average} S$

where S is the current port surface area $\pi D L$ (assuming a circular port, otherwise use the **hydraulic diameter**).



Note that $\dot{m}_{total} = \dot{m}_{fuel} + \dot{m}_{oxidiser}$ is the *total* mass flow rate at the end of the port and hence flowing towards the nozzle. Therefore, the solid erosion is a cascading effect: the more fuel gets eroded, the larger the total mass flow rate, hence the more fuel that gets eroded. It's a feedback mechanism, but fortunately it's a converging one, and well-damped. Needless to say though, an implicit, iterating software loop is therefore required to get the final value of \dot{m}_{total}

Excel spreadsheets will handle iterative calculation loops well, though you've got to tick the iteration function to on. (On the top menu select: file [or tools on older versions], options, formula, calculation options, and tick the box 'enable iterative calculations'.)

From reference 2, a simple, robust, and rapidly converging algorithm you can use for this is:

1. Start by assuming that $\dot{m}_{total} = \dot{m}_{oxidiser}$ only (i.e. $\dot{m}_{fuel} = 0$)
2. Determine $\dot{r}_{average}$ and hence \dot{m}_{fuel}
3. Add $\dot{m}_{fuel} + \dot{m}_{oxidiser}$ to get a new \dot{m}_{total}
4. Has the new \dot{m}_{total} settled down (converged) to equal the previous step's \dot{m}_{total} to the required accuracy? If it hasn't, use this new \dot{m}_{total} and loop back to 2. If it has, the calculation is complete.

Excel handles the final check automatically: you simply set the convergence limits in the same sub-menu, and then make spreadsheet cell 1 equal to spreadsheet cell 3 to close the loop.

The exponent m in the above regression equations tends to be a very weak power of length, much less than 1.0, so for the moderate port length changes involved in tuning a hybrid for performance it can be ignored.

Furthermore, when modifying an existing design, one will use the same ignition system, so will expect to see similar fluxes down the port.

In this case, you can get away with an engineer's simplification of the regression equation:

$\dot{r}_{average} = a G_{liquid}^n$ where the length term has been lumped-in to the catch-all coefficient a .

This is much easier to use as it's based only on the liquid flowing in from the injector, so no iterative calculation is required. This equation simply won't work if you double the size of the hybrid for example, and may well not be transferable to someone else's design.

Appendix 1 details how to convert regression rate data between systems of units.

The secondary transport mechanism (see above)

This is modelled as detailed in Reference 4. If your hybrid uses a polymer fuel then this mechanism won't occur, so you can skip this section.

As the 'wind' down the port increases (**mass flux** increasing) it's found that the wavelength between the waves on the melt layer decreases. So the number of wave crests available in the port to provide a spray of droplets increases, and that's doubtless one reason why the regression rate increases with mass flux.

From Ref. 4 a general expression for the **mass flow rate** of fuel liberated by the secondary transport mechanism alone (sometimes called the 'entrained' flow) per square meter of port surface area is given as:



$$\dot{m}_{entrained} = \rho_{liquid} \dot{r}_{entrained} \text{ varies with: } \frac{G_{total}^{2a} h^b}{\mu_{liquid}^c \sigma_{liquid}^d \rho_{port_gasses}^a}$$

where h is the thickness of the melt layer, μ is the viscosity of the melt layer liquid and σ is its surface tension: these last two are assumed constant for a particular fuel type. (a, b, c, d are constants.)

An expression is given to compare the entrained regression rate to the total regression rate:

$$\dot{r}_{entrained} = C \frac{G_{total}^{2a}}{\dot{r}_{total}^b} \text{ where } C \text{ is effectively a constant for each solid fuel.}$$

Initial experimental data for the constants are: a is around 1.5, b is about 2, d is less than unity; possibly a lot less, and c is greater than d , expressing the view that a low viscosity has a much more significant effect than low surface tension.

Interestingly, experiment shows that if the port mass flux is too low, there will be no secondary transport mechanism, there's a lower flux limit. In practice, for all but viscous polymers, typical hybrid port mass fluxes are much higher than this minimum.

From the above expressions, the regression rate depends on the thickness of the melt layer (typically of order 0.3 mm).

You'd typically make your fuel opaque to infrared radiation from the flame zone, typically by doping it with black dye or 'carbon black' (soot), so that only the top layer of the fuel gets heated, otherwise infrared radiation could penetrate deep into the fuel and it might start to creep or slump as it heats (which generally happens in practice!)

If the fuel is opaque, then one would also expect the melt-layer to be opaque. An expression for the thickness of an opaque melt layer is:

$$h = \left(\frac{\kappa_{liquid} \rho_{solid}}{\dot{r}_{total} \rho_{liquid}} \right) \ln \left[1 + \frac{C_{liquid}(T_v - T_m)}{L_m + C_{solid}(T_m - T_a)} \right]$$

where T_a, T_m, T_v are the initial, melting, and vaporization temperatures of the fuel respectively (Kelvin), C is the specific heat capacity (KJ/kg), L_m is the Latent heat of melting (KJ/kg), κ is the thermal diffusivity (m²/s), and ρ is density (kg/m³).


In this equation, the thickness h is inversely proportional to \dot{r}_{total} , the total (primary plus secondary transport mechanism) regression rate in mm/sec. Ln is the natural logarithm function. Annoyingly, this does make the equation rather recursive; you need experimental regression data to estimate h .

Post-combustion chamber

Even with an optimum port length, it's usually the case that not all of the vaporised solid and liquid have been able to mix and burn before exiting the end of the **port**; the centre of the port tends to be fuel-lean whilst the flow at the wall is fuel rich.

Simply adding an empty space downstream of the port just before the nozzle is all that is required to recoup nearly all of this loss.

As the flow exits the port, it goes over the edge where the grain ends causing a trapped toroidal vortex to occur, which mixes most of the flow before it encounters the nozzle.



This post-combustion chamber is typically one chamber diameter in length: the extra chamber mass required is more than recouped by the increase in overall performance.

Propellant choices

Despite the fact that the first hybrid was doubtless tested centuries ago, only recently have they caught the world's attention, simply because hybrid's sedate **thrust-curves** have no military applications. A fact we're quietly delighted about.

Thus the actual number of possible hybrid propellant combinations that have been tested so far is a tiny fraction of all possible hybrid combinations (anything that'll burn). We'll consider the liquid propellants first, but the choice of solid propellants is more complex, so I'll get to that later.

Liquid propellant selection

The simplest liquid propellant choice, and certainly with adequate performance, are the just-subcritical propellants such as nitrous oxide or near-critical oxygen for forward hybrids. Or ethane, propane, LPG, ethylene etc for reverse hybrids. (see our 'physics of nitrous oxide' article for the definition of just-subcritical.) They all have the advantage of not requiring a separate pressurant or pump, as they self pressurise.

To achieve a decent thrust from a rocket motor, its combustion chamber has to operate at a high internal pressure.

Newton's laws dictate that a mass of liquid or gas will only flow from a higher pressure region to a lower pressure region so the pipework and tank upstream of the motor has to be pressurised to a higher pressure than the combustion chamber in order to get the propellant in.

Nitrous oxide and ethane in particular both give moderate density, good **specific impulse**, and high pressure at room temperature.

These just-subcritical propellants will vaporize readily with small drops in pressure, allowing the use of the most trivial injectors. (see our 'physics of nitrous oxide' article.)

The downside is that for accurate testing and tuning, you need to develop a mathematical model of the near-critical 'blowdown' tank-emptying process to get a handle on the variation of oxidiser mass flowrate through the injector with time. The coding of this was a serious chore; it's detailed in our paper 'modelling the nitrous run tank emptying'.

Next are the 'traditional' liquids that are *not* near their critical point: liquid oxygen or hydrogen peroxide for forward hybrids, or alcohols, gasolines, etc for reverse hybrids. These *do* require a separate pressurant, such as helium, nitrogen, or CO₂, or alternatively some type of pump, to raise them above the pressure of the combustion chamber so that they will flow into it.



Furthermore, they require a more complex injector in order to produce streams of sufficiently small droplets: the cross-sectional area of the streams leaving each orifice of the injector must be small enough that surface tension forces come into effect, and break the stream into a hail of tiny droplets that quickly get vaporised by the hot gasses in the combustion chamber.

For smooth motor running, it's important to ensure that only gas enters the solid-grain **port**. If big droplets of liquid get into the port, they've got enough momentum to penetrate the boundary layer and splash against the solid propellant, whereupon the droplet explodes with the solid in its immediate vicinity, causing a minor crater that destroys the symmetry of the burning surface, and also produces a pressure pulse inside the chamber.

A continuous hail of liquid droplets will cause a continual barrage of tiny explosions that severely rattle the motor, and also cause a multitude of random pressure spikes that just might add together at some point during the burn to dangerously overpressure the combustion-chamber.

Lox hybrids are particularly prone to this, and the only remedy is to pre-heat the oxygen. So technically Aspirespace's H2O Lox/polyethylene hybrid is actually a Gox/polyethylene hybrid with cryogenic storage of the oxygen to increase its density.

When using these non- near-critical liquids, you typically need an empty space in the combustion chamber between the injector and the start of the solid grain. This 'vaporisation chamber' gives the incoming liquid propellant a few more milliseconds to break up properly before pouring down the **port**.

Some hybrids, such as our H2O, provide a little burning fuel in the vaporization chamber to heat up the incoming liquid, in order to vaporize it better. This is then called a 'pre-combustion chamber'. The just-subcritical propellants above don't need either.

On the plus side, while one could develop the most horrendously complex tank-emptying mathematical models for these more traditional liquids, you can get away with simple models, because you can directly measure the liquid flow rate using flow-meters between the tank and injector.

Solid propellant selection

The hybrid is primarily designed to the **thrust-curve** that you want, though try to be flexible. It's much easier to design hybrid solid grains to produce the long, gentle burns suitable for horizontal-takeoff winged launch vehicles than it is for the short, high-thrust 'grunty' burns needed for vertical takeoff launch vehicles or boosters.

The reason is simply that the thrust at any point in the burn depends primarily upon \dot{m}_{total} , which itself depends upon how rapidly the solid grain is regressing, so the higher the thrust required, the more rapidly the **port** diameter has to increase.

This poses a geometrical problem:

$$A = \pi r^2$$

so from the chain rule of differentiation, as the port regresses, this area increases with time at an over-linear rate:

$$\frac{dA}{dt} = \frac{dA}{dr} \frac{dr}{dt} = 2\pi r \dot{r}$$



Unfortunately, assuming a constant liquid propellant mass flowrate into the port, this means that the total port **mass flux** is decreasing at this over-linear rate because (from the mass flow rate formula) it is inversely proportional ($=1/A$) to area A . (see glossary) and so the thrust will drop, and so will the solid propellant mass flux, which means that for a forward hybrid, the burn is getting progressively leaner (too little fuel) with time. So even more thrust is lost.

Fortunately, as the port widens, its total surface area $S (= 2 \pi r L$ which is the area available to farm fuel mass flow from) increases, but it's only increasing linearly with the regression rate:

$$\frac{dS}{dt} = \frac{dS}{dr} \frac{dr}{dt} = 2\pi L \dot{r} \quad (\text{where } L \text{ is the total port length})$$

whereas as we've just seen, the mass flux is dropping over-linearly, and so is dominating.

So contrary to popular opinion, the best hybrid propellant combination is the one that has the lowest regression rate that you can use for your application, because the thrust and fuel-to-oxidiser ratio will then wander the least with time.

A low regression rate tends to require a long port to provide enough surface area to farm sufficient solid propellant, which creates a long, skinny combustion chamber.

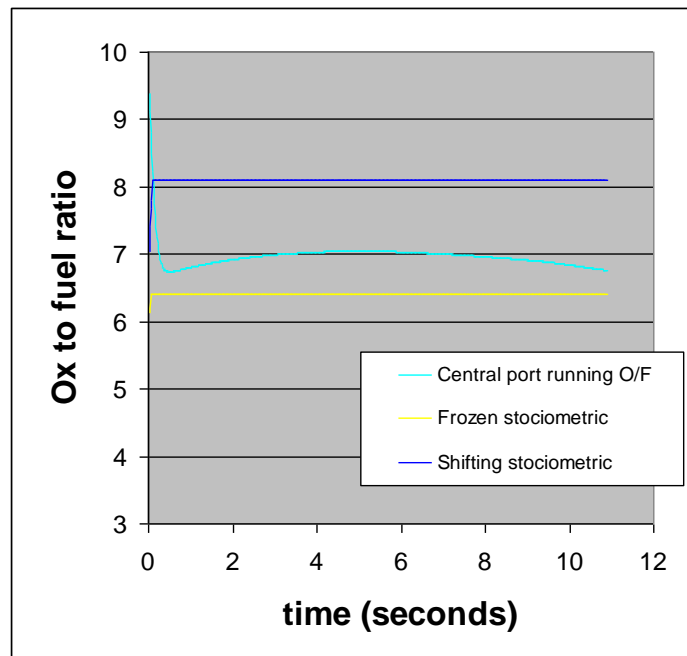
Designers of vertical-takeoff rocket vehicles shy away from skinny chambers as they're a long way from the ideal lightest pressure vessel which is a sphere. But aircraft designers needn't worry: a long, skinny hybrid chamber provides a stiff longitudinal beam.

The above \dot{r} versus \dot{r} squared geometrical imbalance means that some propellant combinations that work at one physical scale are not good at others.

For example, at the size of hybrid required for 'HPR' rocketry levels 1-3 (see the UKRA website, www.ukra.org.uk), simple single-port plastic solid grains are fine.

In fact, the plastic's regression rate is near-optimum for this engine size:
By tuning the injector, we've managed to tailor the drop in tank pressure with the nitrous flowrate out of the tank to balance-out the drop in high-density polyethylene fuel mass flux as the port grows in diameter.

So the burn remains pretty-well stoichiometric for our Aspirespace ADV2b hybrid as simmed here:





Polymer fuel grains

Typical plastic fuels are:

- Synthetic Rubber (though this can cause a sooty exhaust).
- Perspex (as used in our classroom demonstrator H0 because it is transparent).
- Nylon.
- High-Density PolyEthylene (which we use, as it has a slightly higher density and **Specific impulse**).
- Hydroxyl-Terminated PolyButadiene (which isn't a terribly good choice, but solid-motor manufacturers have a lot of it lying around.)

Other suitable polymers are wood, rolled paper, compost, and almost any foodstuff.

Attempts to use polymers at larger scales have generally run into trouble, especially when using oxygen with plastic as the **stoichiometric** ratio for that particular combination is low, and hence a higher proportion of fuel is required. The hybrid fuel grain becomes excessively long and skinny, and would need to be folded or coiled.


One solution is either to cluster a batch of smaller hybrid combustion chambers in parallel, or, doing effectively the same thing, cluster a batch of **ports** in parallel in one, large, solid grain.

From a mass point of view, the latter looks more attractive because only the one combustion chamber casing is required, and it's incorrectly assumed that this is the only way to allow the use of a single large nozzle.

But in practice, multi-port hybrids are a design pain:

As the ports regress larger, eventually they'll meet and join-up, and as they join, large slivers of solid fuel can break free and block the nozzle. A blocked nozzle is dangerous for any type of rocket (see our 'hybrid safety' article).

So either you leave a lot of unburnt plastic remaining after engine burnout (i.e. space the ports further apart so that they don't join up) or you have to use a lot of inert structural support (such as phenolic) to wall the ports off from one another.



Schematic of a Fuel-Web Support.
As the fuel burns, it becomes too thin to support itself, so insert inert structural supports.

Either way, this leaves a lot of dead weight in the chamber after engine burnout, whereas in contrast our Rickrock (ADV2c) single-port hybrid burns the plastic right out to the liner on the chamber wall: no plastic remains.


The HTPB/rubber fuel grain used in Rutan's X-prize-winning Spaceship One was a quad port similar to the picture above. Spaceship Two may use six or eight ports.

Several alternative ways have been developed to allow the use of polymers at larger scales:

In a vortex hybrid, as mastered by the Orbitec Corporation, the liquid propellant is injected into the port with a large tangential component.

So although the *average* mass flux down the port is of typical value, centrifugal effects throw the gasses out to the walls: the *local* mass flux *at the solid fuel surface* is much higher.

This gives the impression that the regression rate of the plastic *based on the average port mass flux* has somehow magically increased. In any case the regression is indeed much higher, and is aided by the further centrifugal effect that the denser unburnt gasses force their way into the flame zone under high gee more readily, because the burnt gasses have expanded and are less dense.




In the pancake hybrid, a variant of which has been successfully developed by Surrey Satellite Technology Ltd, the solid fuel grain is a flat disk with a small hole in the centre, that sits on the floor of a squat combustion chamber.

In the SSTL design, the liquid propellant was injected through a ring running around the circumference of the disc. It regressed the top surface of the plastic disc, then exited down the central 'plug-hole' to the nozzle underneath.

Though the regression rate was essentially unchanged, the geometry of the resulting motor was a lot more manageable: short and squat.

Also, the port cross-sectional area depends on the distance from the top of the fuel-grain to the ceiling of the combustion chamber: as the grain regresses towards the floor, this distance increases linearly with time.

So the port cross-sectional area grows *linearly* with time, *not* with \dot{r} squared as in a single circular port, so the burn remains more on-design as time passes.



In solid rocket motor design, cunning grains have been devised that keep the port surface area constant with time as this keeps their thrust constant with time. You could try this with hybrids, but you have to juggle the extra ball that you've also got to keep the port cross-sectional area constant with time as well. It's doubtless easier to devise a grain that keeps the two areas increasing at the same rate.


Cryogenic and wax fuel grains

If one wants to use a single port at larger scale, geometry dictates that you need a solid grain that regresses faster.

In the 1990's tests were performed at Edwards (Reference 5) of solid grains formed by freezing pentane and ethylene using liquid nitrogen. These 'cryogenic' hybrid grains were found to regress much faster: three times faster than plastic for pentane and acetone, and ten times faster for ethylene.

Ref. 5 describes how the mechanism of the faster regression rates was discovered: the secondary transport mechanism that I described earlier.

By the turn of the Millenium, researchers were freezing all sorts of hydrocarbons, such as Kerosene, methane, etc.



Regression rates are plotted versus the average oxidiser mass fluxes for various materials tested by AFRL. Curve fit expression for pentane data is $\dot{r} = 0.123(G_{ox})^{0.63}$ (where \dot{r} is in mm/sec and G_0 is in kg/m²-sec). For comparison purposes the burning law of a plastic propellant, HTPB, is included in the plot.

This secondary transport mechanism relies on low viscosity in the melt-layer on the surface of the solid grain (and to a lesser extent lowered surface tension as well).

Presumably, the lower the viscosity, the faster the regression rate. There is a simple relationship between the viscosity and the hydrocarbon's complexity (molecular weight): the simpler the hydrocarbon, the lower the viscosity, taking hydrogen at one end of the scale, and plastic at the other. Tests by Orbitec of frozen (!) hydrogen (H₂) and frozen methane (CH₄) reveal staggeringly fast regression rates, around 30 times faster than plastic.

Orbitec tell me that the above data is from tests in conventional hybrids, and *not* their vortex hybrids with the centrifugal effects discussed above.

These rates are so enormously fast that there isn't a hope of keeping a single-port grain stochiometric for any length of time if the port is changing geometry so rapidly, unless the hybrid was truly titanic in size.


Orbitec certainly couldn't keep the port stochiometric on their small lab-scale test motors, they had to inject a secondary stream of liquid propellant into the post-combustion chamber to trim the burn. This 'afterburner' concept is a neat idea: it's turning the hybrid back towards a liquid, which will up the performance a little, because less of the total propellant will absorb heat (of fusion) as it turns from solid to liquid, but it avoids the chance of a liquid engine hard start. The hybrid grain is then functioning rather like a pre-burner.

Still, this 'tribrid' adds complexity; Orbitec had to burn the port way off stochiometric, perhaps to keep the temperature of the flow into the post-combustion chamber low enough to prevent melting the secondary injector, though more likely to *reduce* the port regression rate!

For the larger motors we amateur groups are devising, we want a grain that'll regress somewhat faster than plastic, but not much faster. On the hydrocarbon graph, we're heading back towards the plastics, but trying to avoid actual long-chain polymerisation. The resulting choices are wax, soap, bitumen etc, the largest of which yet tested in the UK was when a few gallons of liquid oxygen got accidentally dumped on Cranfield University's carpark in the 1950's, and regarded the tarmac as fuel...

The Space Propulsion Group, (which budded off from Stanford University) in collaboration with NASA Ames, have been testing and flying small to medium motors (up to 190 mm diameter) using pretty-much standard candle wax, which happens to be paraffin wax, and therefore can be regarded as effectively solid kerosene.

They recommend using the more mechanically strong, higher melting temperature wax used for making Hurricane candles, though they've added a few secret additives for extra strength, and opacity to infra-red.



Unfortunately, though many groups have tried to replicate their results, nobody can. Aspirespace have experimented with wax hybrids (see our website), and all our tests but one just generated a gooey mess. The majority of the wax melted instantly and poured out the rocket's nozzle like water spewing from a firehose: the regression rate was way too high to be useful. Only a 50/50 mix of wax and low-density low-melt temperature polyethylene performed as expected, but not much better than straight polyethylene.

The California Rocket Society tested a form of wax in a hybrid in 1938, but it wasn't successful; probably for the same gooey reasons.

The Space Propulsion Group would appear to be using an extra secret ingredients in their wax: we'd all love to know what they're using.

Aspire added a little carbon-black to our wax to make it opaque to infrared to block the infrared radiation in the chamber from cooking the wax right through, but this gave a surprisingly soft wax that ran out the nozzle like water.

With the obvious benefits of high regression, non-toxicity, and very easy mouldability, why haven't wax hybrids 'taken off' in a big way?

Informed rumour says that they suffer from a problem: one of the positive safety aspects of hybrid combustion is that the burning rate is insensitive to cracks or voids in the solid propellant grain. However, it appears that wax hybrids (and probably cryogenic hybrids as well) *do* react to voids or cracks in the grain: any voids or cracks promote a circulation that greatly increases the regression rate locally, so that the void or crack grows rapidly in size. This will eventually expose the bare combustion chamber insulation and so could lead to a burn-through of the chamber wall.

Personally, I feel that this problem is fixable. What is needed is a composite grain that mechanically prevents cracks or voids from growing: perhaps one could use wax sandwiched between sheets or nets of plastic or paper. This would also mechanically strengthen the grain, which is another possible issue with large wax grains.



Glossary:

Fill-tank: the commercial container supplied with the nitrous.

Hydraulic diameter: for a port of non-circular cross section, we need a metric equivalent to the diameter of a circular port in order to calculate mass flux. This equivalent is known as the hydraulic diameter and is defined as: $D_H = \frac{4A}{P}$ where A is the cross sectional area and P is the wetted perimeter of the cross-section.

Injector: a device similar to, sometimes identical to, a shower-head. It is designed produce streams of droplets that are small enough to vapourise quickly before leaving the exit of the heated combustion chamber. The more complex injectors needed for *non* near-critical propellants splash multiple exit streams into each other to break them up; these are known as impinging-stream injectors.

Mass flow rate \dot{m} : is the speed of propellant flow down a pipe or **port**, in kilograms per second. The Mass continuity equation can be used to calculate this as $\dot{m} = \rho AV$ where ρ is the gas density in kilograms per cubic metre, A is the cross-sectional area of the pipe the propellant is flowing down in square meters, and V is the propellant flow velocity in metres per second.

Mass flux G : (strictly \dot{G} but oddly seldom written this way) is the 'draught' or 'density flow rate' of propellant passing over the surface of the **solid grain**. This is equal to the **mass flow rate** (see above) divided by the cross-sectional area of the pipe/**port** it is being forced to flow through, i.e.

$G = \frac{\rho AV}{A} = \rho V$ where ρ is the gas density in kilograms per cubic metre, and v is the flow velocity in metres per second.

The higher the density and the higher the velocity (so the higher the mass flux), the more molecules of propellant will encounter the boundary layer on the internal surface of the pipe per second.

Port: from solid rocketry parlance, this is the name of the hole down the inside of the **solid grain**.

Regression rate \dot{r} : is the velocity (typically millimeters per second) at which the surface of the solid propellant is eroding 'inwards' as it vaporizes due to the heat of combustion.

Run-tank: the lightweight tank inside your rocket-vehicle that is filled from the fill tank. (in a forward hybrid, the term 'fuel tank' is just plain wrong as the fuel is the plastic in the combustion chamber.)

Solid grain: from solid rocketry parlance, the block, rod, or tube, of solid propellant within the combustion chamber.

Specific impulse (ISP): the thrust produced per unit weight of propellant burnt per second. Equivalent to the 'miles per gallon' metric of cars.

Stoichiometric: the mixture ratio of fuel to oxidiser giving best rocket performance. This is usually also the mixture ratio that ideally would totally consume both fuel and oxidiser, though in practice the mix for best thrust tends to be a little fuel-rich.

Thrust-curve: in rocketry parlance, the graph of thrust (vertical axis) versus time (horizontal axis). From calculus, the area bounded by this curve and both axes is the integral that leads to how fast your vehicle will be traveling at motor burnout. This area is known as the Total impulse (in Newton seconds).



Technical papers

Viscosity: Viscosity, just as in car engine oil grades, is basically the 'syrupiness' of the fluid. Light machine oil having low viscosity, and thick tar having high viscosity. Viscosity is technically defined as resistance to shearing (rate) of alternate layers of fluid.

Conversion factors:

Mass flux: 1.0 gram/centimetre² = 10.0 kilogram/metre²

Thrust: 1.0 pounds force = 4.448222 Newtons



References:

Ref. 1: 'Thermophysical properties of nitrous oxide', Engineering Sciences Data Unit (ESDU) sheet 91022

Ref. 2: 'Space Propulsion Analysis and Design', Ronald .W. Humble, Gary .N. Henry and Wiley J. Larson, McGraw Hill Space Technology Series ISBN 0-07-031320-2

Ref. 3: Dr Bruce P. Dunn, University of British Columbia and Dunn Engineering. Several articles on self pressurised peroxide rockets and experiments on propane tanks, as well as email communications with the author on the subject of numerical modelling of the tank emptying process; with many thanks.

Ref. 4: 'Development and testing of paraffin based hybrid rocket fuels', AIAA paper 2001-4503, M.A. Karabeyoglu, B.J. Cantwell, and D. Altman

Ref. 5: 'Lab scale test and evaluation of Cryogenic Solid Hybrid Rocket Fuels', AIAA paper 95-2948, Patrick G Carrick and C. William Larson, Phillips Laboratory, Propulsion directorate Edwards Airforce Base, CA

Ref. 6: Aerocon hybrid factors paper 1997, and design manual 1995

Ref. 7: Space Propulsion Group website

Ref. 8: Wikipedia, 'hydraulic diameter'



Appendix 1: converting regression data

Experimental **regression rate** data can be given in many units, particularly as the Americans haven't joined the civilised world and gone metric.

To convert regression rate data from one set of units to another, take account of the power that the mass flux is raised to:

For example:

$0.104 \times G^{0.352}$ mm/s regression rate per g/cm²s mass flux

equals:

$(0.104 / 1000) \times (1000/10000)^{0.352} \times G^{0.352}$ m/s regression rate per kg/m²s mass flux