

The physics of Nitrous Oxide

Most amateur groups choose Nitrous Oxide, often referred to as 'nitrous', and sometimes 'nitrogenous oxide', (though *not* 'nitro' which is nitromethane) as the oxidiser for their hybrid rocket engines. Another hybrid rocketry nick-name for Nitrous Oxide is 'Nox', but 'NOx' is actually a broad environmental term for any of the various compounds and derivatives in the family of nitrogen oxides, including nitrogen dioxide, nitric acid, nitrous oxide, nitrates, and nitric oxide.

Nitrous oxide's chemical formula (N_2O) shows a predominance of Nitrogen, which doesn't help at all with burning; it appears at first sight just to be dead-weight that has to be carried aloft, though it cools the nozzle sufficiently that graphite nozzle inserts can be re-used many times.

But in fact the 'inert' nitrogen actually performs a very useful function within hybrid combustion chambers: Nitrogen mass flow is the majority of the nitrous mass flowing down the central hole or fuel port, and so this does most of the eroding of fresh fuel to burn with the oxygen within the nitrous.

So it significantly aids the fuel erosion rate. (see our Introduction to hybrid design paper)

That's why it's much harder to get a decent thrust out of pure oxygen hybrids: the **Specific Impulse** is higher, but the fuel erosion rate at **Stoichiometric** mixture is much lower.

So blasting another inert gas down the port of a liquid oxygen hybrid would help up the fuel liberation (but that's heading back towards nitrous oxide again.)

Still, tweaked for performance, nitrous hybrids will outperform all but the highest performance solids, so the following are some of the points to consider when designing and/or using a nitrous hybrid.

The overtly useful aspects of nitrous are:

1) The simple gas bottles nitrous has to be stored in are a lot cheaper to buy or rent than, say, liquid oxygen or hydrogen peroxide containers, so at the small quantities most amateur groups use, nitrous systems work out cheaper, even though the nitrous itself is quite expensive per litre.

Oxygen has to be chilled below about minus 120 degrees C before it provides a reasonably dense liquid phase: most of the money Aspirespace spends on each H2O hybrid test goes on the containers and cryogens necessary just to keep the liquid oxygen (Lox) this cold for quarter of an hour or so in the test-tank.

Nitrous is readily available from many sources, such as hotrod car shops, whereas a helluva lot of health and safety paperwork has to be done before anyone will sell you Lox or High Test-concentration Peroxide (HTP), or even worse, the utterly toxic Red Fuming Nitric Acid (RFNA).

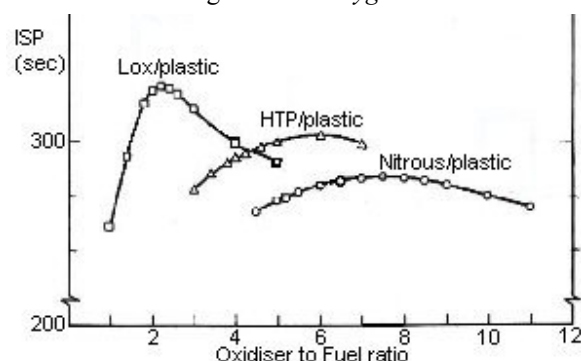
2) Just like peroxide, a large oxidizer to fuel ratio is required when burning nitrous in the combustion chamber (around 7:1 by mass) which results in a requirement for vast quantities of nitrous, and so a large tank onboard which it's difficult to keep from being heavy.

This high ratio isn't all bad news, because as the oxygen within is a low fraction of the total nitrous, you can be quite sloppy with the 7:1 *nitrous* to fuel ratio without altering the actual *oxygen* to fuel ratio within, much.

This means that unlike other oxidizers, a graph of **Specific Impulse** (in a vacuum) plotted against oxidiser-to-fuel ratio *doesn't* have a sharp peak at best (**stoichiometric**) mix that drops off sharply on either side of the peak.

The graph shown for Nitrous/plastic combustion is the flattest compared to the others, decreasing by less than 5% of optimum ISP over a **stoichiometric** range of 5:1 to 10:1 Ox:Fuel ratio.

So you'll still get plenty of thrust even if your mixture ratio of nitrous to fuel is way off 7:1, which is good if your test rig can't give you accurate figures to let you tune up the motor: the first few flights will still be adequate, provided the motor doesn't melt.





3) Like bottled CO₂, nitrous is **subcritical** at room temperature meaning that both a liquid and a gas phase can coexist within a *closed* tank.

I'll elaborate on this in a moment, but the gist of it is that the moderately dense liquid phase of nitrous can therefore be stored in a compact tank on the pad in the British climate.

4) At room temperature, Nitrous is *only just* subcritical by a few degrees.

This is nitrous's most unexpectedly useful property, because this close to the Critical temperature, small drops in tank pressure cause large-scale production of extra vapour. This extra vapour strives to maintain the tank pressure at high value as the tank empties.

A traditional blowdown system, e.g. using an ideal gas such as helium, loses tank pressure at a much higher rate during the burn.

This willingness to vapourise with small pressure drops means that the nitrous will vapourise within a few millimetres downstream of even the crudest injector, typically even a simple single hole.

Injector design is therefore trivial, though noticeable improvements still can be gained from more traditionally complex designs.

5) As an added bonus, the pressure of the nitrous gas phase (termed the 'vapour pressure') is seriously high at room temperature, at around 55 Bar (800 PSI).

The gas phase can therefore be used aquajet-style to squirt the liquid phase into the combustion chamber at very high pressure. Some groups call this 'Vapak' pressurization (Vapour pressurization). This means you can tweak the combustion-chamber to be at almost this high a pressure and the nitrous will still run downstream (in a pressure sense) into the chamber.

The higher the chamber pressure, the higher the Specific Impulse of the motor: our AspireSpace hybrids run at about 35 Bar chamber pressure, which is about as high as you can get whilst still having enough of a pressure drop between tank and chamber to prevent a screaming motor.

6) Nitrous has to be raised to a moderately high temperature before it will decompose and release its oxygen.

This is very good from a safety point of view, but it does mean that a lot of heat has to be pumped into the nitrous from some other source at ignition, or the hybrid simply won't light-up.

Once the plastic fuel is burning though, the temperature in the combustion-chamber is high enough to decompose the rest of the nitrous as it feeds-in from the tank during the burn.

Filling the run tank:

I tend to call the beefy container supplied with the nitrous the **fill-tank**, whereas the lightweight tank inside your rocket-vehicle that it fills I call the **run-tank**. (the term 'fuel tank' is just plain wrong; the fuel is the plastic in the combustion chamber.)

The run tank is filled using a difference in pressure between the fill and run tanks:

The pressure difference (or 'head') caused by gravity when the run tank is connected to a fill tank that is physically higher than it will fill the run-tank with the denser liquid phase, while the lighter gas phase will bubble back up into the fill tank.

Alternatively, the run-tank has a vent-hole in it, which lowers the pressure in the run-tank relative to the fill tank (see below).

The massive pressure difference between the inside of the fill-tank and the outside air will happily carry the nitrous several metres 'uphill' into the run-tank, so the fill-tank can then be physically lower than the run tank; it is typically lying on the ground whilst the run-tank is up in the rocket up on the launcher.

A question often asked is how full can you fill the run tank?

As we'll see later in this article, the expulsion of the liquid nitrous phase out of the tank during the burn *is not* a simple blowdown process, because the nitrous vapour is definitely not an ideal gas.

Nitrous performs much better than this, and in fact our test-firings and simulations show that the graph of tank pressure drop with time (during the firing) does not depend upon the amount of nitrous vapour originally in the top of the tank, so you could fill the tank completely full of liquid.



Technical papers

But if the run-tank is to be completely sealed after filling, but some time before firing, then for safety reasons (see our hybrid safety article on hydraulic overpressure), a small percentage of the tank volume is deliberately left free of liquid to allow for liquid nitrous expansion with the ensuing increase of temperature.

Vents

On many hybrid systems, this ‘ullage’ or ‘head-space’ of gas is achieved by a vent-hole or vent-pipe with an inlet situated slightly below the top of the tank, the outlet is open to the atmosphere outside.

A vent works exactly like the overflow outlet on a bathtub in that the liquid never fills higher than the vent. (provided that you fill it reasonably slowly.)

The outlet of the vent-pipe can be higher than the vent inlet if required, because the massive pressure difference between inside the tank and outside will happily carry the nitrous several metres ‘uphill’.

As soon as the nitrous reaches the level of the vent, you’ll see the plume issuing from the vent thicken and whiten appreciably, and that’s the time to stop filling.

A dark background behind the vent outlet aids this visual check.

If your hybrid design allows, now’s also the time to close the vent hole to stop the loss.


Most commercial nitrous hybrid systems keep the vent open permanently, therefore nitrous is continuously being lost.

Although a small enough vent diameter will keep the tank pressure high for some time, this progressively lowers the tank vapour-pressure over time as discussed below, so such a design has to be launched *immediately* after filling.

Faff around on the pad for too long, and significant thrust is lost.

In this diagram, the fill-tank on the left has to be tilted-up to get liquid phase out, whereas the fill-tank on the right has a ‘dip-tube’ running down inside it so that it can be sat upright.

When you purchase your nitrous, remember to ask whether the fill tank has a dip tube fitted or not.



Subcriticality and Supercriticality:

The apparent simplicity of nitrous hybrids comes at a price.

The nitrous is typically at a temperature where its physics is anything *but* simple, but as in every other branch of rocketry, do thy homework to get thy max performance.

Most substances, below a Critical point (each substance has its own Critical temperature and pressure), can exist as more than one phase simultaneously; they are then termed subcritical.

For example any combination of two of the solid phase, liquid phase, or gas phase, can exist together in a tank in ‘phase equilibrium’, or even all three at the same time at the ‘Triple point.’

Nitrous oxide sitting inside a closed container at room temperature is subcritical: partly liquid, and partly gas which being less dense collects at the top of the container.

Strictly, the term subcritical is taken to mean ‘just subcritical, but near to the Critical point’ but this applies to nitrous as we’ll encounter it.



Technical papers

Nitrous properties

Below is a table of nitrous properties reproduced from Ref.1.

ρ is the symbol for density.

Note how the vapour pressure and vapour density increase with increasing temperature, whilst the liquid density decreases with temperature.

| Temperature degrees. C. | Vapour Pressure Bar Abs | ρ_{liquid} kg/m ³ | ρ_{vapour} kg/m ³ |
|-------------------------------|-------------------------------|---|---|
| —20 | 18.01 | 995.4 | 46.82 |
| —15 | 20.83 | 975.2 | 54.47 |
| —10 | 23.97 | 953.9 | 63.21 |
| —5 | 27.44 | 931.4 | 73.26 |
| 0 | 31.27 | 907.4 | 84.86 |
| 5 | 35.47 | 881.6 | 98.41 |
| 10 | 40.07 | 853.5 | 114.5 |
| 15 | 45.10 | 822.2 | 133.9 |
| 20 | 50.60 | 786.6 | 158.1 |
| 25 | 56.60 | 743.9 | 190.0 |
| 30 | 63.15 | 688.0 | 236.7 |
| 35 | 70.33 | 589.4 | 330.4 |
| T_{crit} 36.42 | 72.51 | 452.0 | 452.0 |

Firstly, a definition:

The word 'vapour' is usually used to refer to a gas when it's below its Critical temperature and pressure, and so is existing alongside some other phase.

It's purely a matter of context: there's no physical difference between a vapour and a gas, they're exactly the same thing.

Living at the bottom of Earth's atmosphere as we do, all of our experience of phase changes, usually of water, occur with a constant pressure of 1 atmosphere around us, which usually swamps the results of our experiments.

If the atmosphere wasn't there, water would behave quite differently from our usual experience.

To start with, water's subcritical below 374 degrees C so there are always at least *two* phases present below this Critical temperature.

One phase may well be much less obvious than the other though, in fact it's only when the temperature has climbed to 100 degrees C that the pressure of water's vapour phase gets as high as the atmosphere around it.

What we call boiling is when bubbles of water vapour can exist without getting squashed flat by the pressure of the Atmosphere.

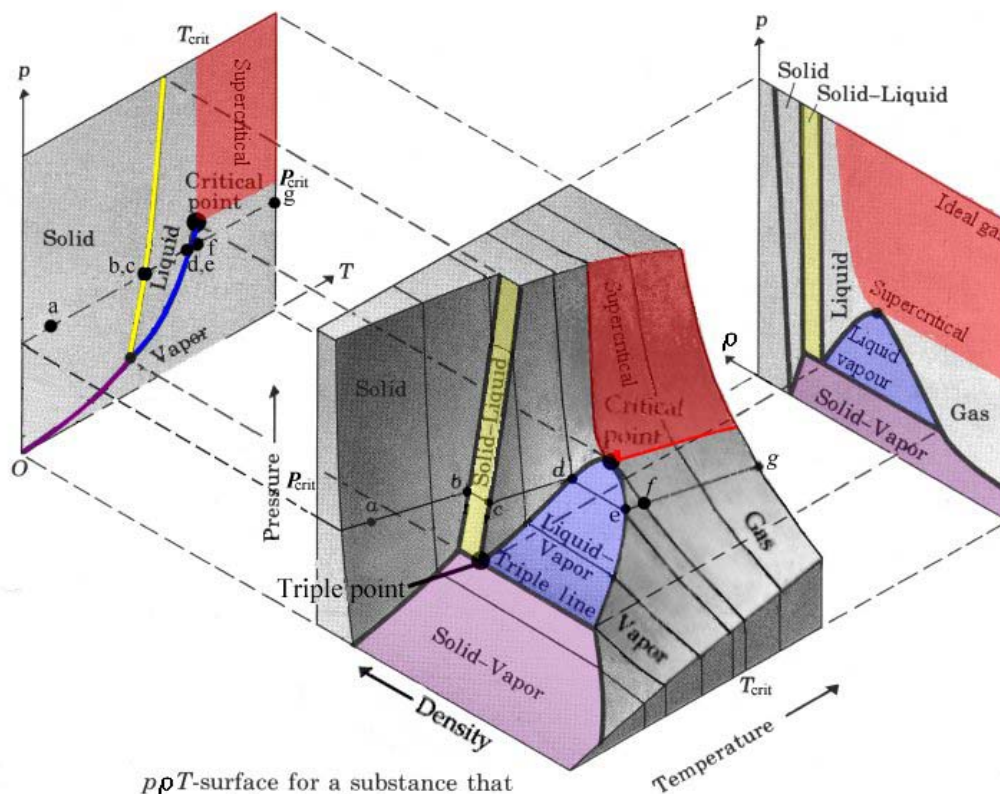
So though we're used to thinking that only liquid exists below 100 degrees, and only gas above 100 degrees, this is actually a high school physics simplification. This is Britain after all; we do get the odd rain-cloud.

Nitrous goes supercritical at plus 36 degrees C, so it's very easy to overheat it into supercriticality:

In the heat of the desert launching campaigns in the United States, the nitrous in several hybrids went supercritical.

Supercritical nitrous requires special injector design, so almost all thrust was lost using the standard injectors.

Here's a 3-D graphical representation (not to scale) known as a phase diagram, of the physical properties of any substance that expands on melting, such as nitrous oxide.



$p\rho T$ -surface for a substance that expands on melting. Projections of the surface on the pT - and $p\rho$ -planes are also shown.

The slopes of this chunk of 'mount nitrous' represent the values that nitrous physically can exist as; pressure being shown as height.

The path a-g on the upper diagram shows the 'isobaric path' (constant pressure contour), i.e. the experience we're used to with water under the constant pressure of 1 atmosphere around us as described above.

b-c and d-e show the sudden changes of phase at constant temperature that we're used to.


(Actually, water is one of the few substances that contracts on melting, so water's phase diagram has 'c' at a higher density than 'b'; it's solid-liquid 'cliff' faces away from us instead of towards us as shown for nitrous, but in all other respects the shape of the 'water mountain' is the same.)

On our planet, nitrous' vapour pressure is well above the pressure of the Atmosphere at the temperatures we'll play with it: Boiling point for nitrous is about 90 degrees C. below zero.

So any air trapped in our nitrous tanks that doesn't immediately get squirted out of the vent hole by nitrous' high room-temperature vapour pressure might as well not be there.

The tank behaves as if it contained only pure nitrous.

In this diagram we zoom-in on the range of pressures and temperatures we'll encounter in rocketry. The density graph shows the view from 'above'.



The liquid-vapour area describes what's happening in your tanks: a subcritical region where both the liquid and vapour phases coexist.

When heated, the liquid phase of nitrous follows the saturated liquid line on the graph whereas the vapour phase follows the saturated vapour line.

The series of parallel lines (parallel to the density axis) that cross lines X and Y are known as 'tie-lines', and it's a *convention* to represent how much mass of each phase there is (as a fraction of the total mass in the tank) by the position along the tie-line.

So by this *convention* (each phase *actually* follows its respective saturation line), the exact path up the coloured section depends upon what fraction of the mass of the substance was in the form of each phase when you started heating it:

For example, path X would be a tank of nitrous mostly filled with liquid, whereas path Y would be a tank of nitrous with mostly vapour in it.

By this *convention*, the liquid saturation line is therefore the path of a tank completely full of vapour.

of liquid that is warming up, whereas the vapour saturation line is the path of a tank completely full of vapour.

(In the upper diagram, lines b-c and d-e are tie-lines.)

Notice that as the temperature increases, the density of the liquid saturation line decreases while the density of the vapour saturation line increases.

This phase diagram is based on real data (see the table above from Ref.1): at the Critical point, the densities do become the same; the two phases merge into one single phase, so paths X and Y both pass through the Critical point.

These are photos taken through a tank window as a substance is heated (left to right) to its Critical point. The line of the liquid surface (the meniscus) disappears: the phases merge completely.



Supercritical nitrous can therefore be regarded as either a super-dense gas, or a low density liquid.

At much higher temperatures, the density of supercritical fluid drops much lower: oxygen or nitrogen at room temperature are well-supercritical, hence we refer to them as 'ideal' gasses at these conditions.

Looking at the density versus temperature diagram, you can also see that the change in density of both phases of nitrous per degree change in temperature is largest (steepest) just before the Critical point.



Technical papers

It turns out that the change in vapour pressure per degree C. is also largest just before the Critical point. For nitrous, even the Scottish climate is still rather close to it's Critical temperature of 36 Degrees C, so sadly, you suffer *big* changes in pressure and density with *small* changes in temperature. A whopping two Bar decrease in vapour-pressure per degree C decrease in temperature is typical in Britain, so if your nitrous gets too chilly, you'll get a lot less pressure in the tank, so obtain a lot less thrust than you expected.

This close to the Critical temperature, the nitrous *vapour* phase is actually moderately dense and can't be ignored; it has a sizable mass inside the run tank. (and inside the combustion-chamber eventually.) Conversely, the liquid phase isn't terribly dense, and is progressively less dense as it is warmed: heat it too much and you won't get as much mass of liquid in the run tank's internal volume. (but chill it too much and you lose a lot of vapour-pressure, select your own preferred temperature.)

From the above table of nitrous properties you can see that at 15 degrees C. (standard U.K. day) the liquid phase is only six times denser than the vapour phase.

Historically, it is the liquid phase that is used in the combustion-chamber. The vapour phase then causes extra thrust after the liquid runs out, but its lower density means that the burning is considerably fuel-rich, so the extra thrust it gives is small.

Changes in liquid/vapour proportion due to temperature alone:

Going back to the lower diagram, look closely at the tie-lines, recalling what they represent, and you'll notice something odd about the paths X and Y.

The ratio of liquid to vapour within a closed run- tank changes with temperature.

This means that the amount of *liquid* nitrous that you *think* is in your run-tank will change over time if you don't take care to keep its temperature constant between the time that you *start* filling and the time that you launch.

So while it may at first seem a good idea to pre-chill the run-tank to get a good fill of dense liquid phase in there, after several minutes of faffing on the pad the nitrous has warmed up and so the situation has changed.

The nitrous is contained inside the fixed volume of the closed bottle that is the run tank, and so it's mass can't change with time:

$$m_{total} = m_{liquid} + m_{vapour} = \text{a constant} \quad (1)$$

So it's forced to self-adjust so that it can physically fit inside the tank as the densities of the two phases change with temperature.

The way it physically alters the volumes of the liquid and vapour phases is that a rise in temperature causes some of the liquid to vaporise into vapour, whilst a drop in temperature causes some of the vapour to condense into liquid.

It's forced to follow a volume formula: $V_{vapour} + V_{liquid} = V_{bottle}$ or, $\frac{m_{liquid}}{\rho_{liquid}} + \frac{m_{vapour}}{\rho_{vapour}} = V_{bottle}$ (2)

where ρ is density, m is mass.

Actually, this 'self-adjustment' phenomenon is very similar to a reversible chemical reaction:

Temperature is defined as the *average* speed of the molecules of the nitrous: some are moving slower than the average, while some are moving faster, possibly fast enough to break away from the liquid phase and become part of the vapour. This is known as evaporation.

Conversely, some of the slower vapour molecules that 'impact' the liquid phase remain as part of the liquid, a process known as condensation.

At equilibrium (the nitrous has reached constant temperature and pressure), the rate of condensation is exactly balanced by the rate of evaporation, so no net change occurs with time.

It's only when the nitrous is no longer in equilibrium that one of the rates exceeds the other, and an overall change occurs.

This all occurs within your closed run tank and so you can't see it happening.

Worse still, the total mass of nitrous in the bottle remains the same of course, so weighing scales won't pick up any changes in the proportion of liquid to vapour.



The following resolves this problem:

Fill calculation (after closing the vent valve):

Assuming that you filled the run-tank slowly then you know what mass of nitrous went into the tank, because the volume of tank above the vent-hole should have been vapour alone, and the volume of tank below the vent-hole should have been liquid alone.

So for example if the head space was 15% of the tank volume, then just at the end of filling:

$$\frac{m_{vapour}}{\rho_{vapour}} = 0.15(V_{bottle}) \quad (3) \quad \text{and} \quad \frac{m_{liquid}}{\rho_{liquid}} = (1 - 0.15)(V_{bottle}) \quad (4)$$

If you don't have weighing scales, these two combine to give:

$$m_{total} = m_{vapour} + m_{liquid} = \rho_{vapour} (0.15)V_{bottle} + \rho_{liquid} (1 - 0.15)V_{bottle} \quad (5)$$

The densities of the saturated liquid and saturated vapour phases can be read off of a lookup table such as given above from Ref.1

A run-tank pressure-gauge is invaluable here, perhaps read using binoculars, to discern what temperature caused this run-tank vapour-pressure reading; it may not have reached ambient temperature yet.

The changes in the proportions of the two phases after some time when the temperature has changed (noted by a change in the vapour pressure reading) can then be calculated by rearranging equation (2) and combining with equations (1) and (5) :

$$m_{liquid} = \frac{\left(V_{bottle} - \frac{m_{total}}{\rho_{vapour}} \right)}{\left(\frac{1}{\rho_{liquid}} - \frac{1}{\rho_{vapour}} \right)} \quad (6) \quad m_{vapour} = m_{total} - m_{liquid} \quad (7)$$

where the densities are those at the new temperature, and m_{total} and V_{bottle} have of course remained constant.

The classic mistake is to forget that pressure gauges measure relative to the atmosphere outside their casing; so one must add one atmosphere (1.013 Bar) to the gauge pressure reading to get the Absolute pressure reading required for the lookup table above.

It's good practice to always label your pressure data 'Bar gauge' or 'Bar abs'.

Some electronic pressure-sensors measure absolute: check their data-sheets.

Proportion changes due to outflow:

Because pressure, temperature, and density are connected, if we cause changes in *pressure* within our run-tank, either during filling, or when we empty its contents into the combustion-chamber, temperature changes will then occur.

And as we've seen, temperature changes cause the ratio of liquid mass to vapour mass in the run tank to change.

Several examples of this occur during hybrid operation:

Firstly, the vent-hole relies on the fact that the vapour-pressure inside the run tank is higher than the atmosphere outside, and so an outflow is established.

The vent should either be of tiny diameter, or be a pipe with a restriction of tiny diameter somewhere along it. (0.3mm diameter is typical.)

A large diameter vent is undesirable because it provides little resistance to the flow pouring out of it, so the drop in pressure between tank and outside occurs more within the tank than within the vent hole.

(electronics engineers will recall the principle of a Potential Divider).



Technical papers

The nitrous responds to this low tank pressure by vapourising its liquid away large-scale. Moreover, the flowrate of nitrous leaving via the vent-pipe is much higher, so it'll all disappear after a short time.

Also, a vent produces gas thrust like any rocket, so you want this flowrate to be small if it's venting sideways.

Similarly, when the **run valve** opens, (the valve between run tank and combustion chamber) the gas phase forces the liquid out of the tank in the manner of an aquajet, because the combustion-chamber connected below the tank is also at lower pressure (unless you've made the nozzle throat too small.) As the tank empties, the liquid level obviously drops, so the volume available to the vapour phase above the liquid increases, so the vapour expands. And like any gas, it's pressure drops as it expands.

Whatever caused the vapour-phase's pressure to drop, be it venting or emptying, the pressure is now lower than it ought to be (it ought to be at it's vapour-pressure) and this drop in pressure is 'sensed' by the liquid phase below it.

Some of the liquid phase will then vaporise in an attempt to create more vapour to raise the tank pressure back up to vapour pressure: the lower the pressure imbalance, the higher the vaporisation rate.

Now the process of vaporising liquid into vapour requires energy (called the latent heat of vaporisation), and this energy has to come from somewhere.

The heat energy required is drained from the nearest available source, which in this case is the remaining liquid nitrous itself, which therefore gets cooled by an amount determined by its Specific Heat Capacity. (Heat capacity per kilo of liquid per degree temperature change.)

Oddly enough, my experiments and simulations show that the metal wall of a nitrous tank doesn't give up heat that quickly into the liquid even though you'd expect it to: the metal may be a conductor, but the liquid isn't.

So the metal of the tank can be ignored as a heat-source for pressure changes, *provided* that they occur in a short time, say the 10 seconds or less that are typical of a hybrid firing.

From an engineering point of view, the various thermal layering effects (known as 'stratification') that occur within the nitrous, wherein the liquid and vapour closest to the boundary between liquid and vapour ought to be the coldest because that's where the vaporisation occurs, can be completely ignored in computer modelling.

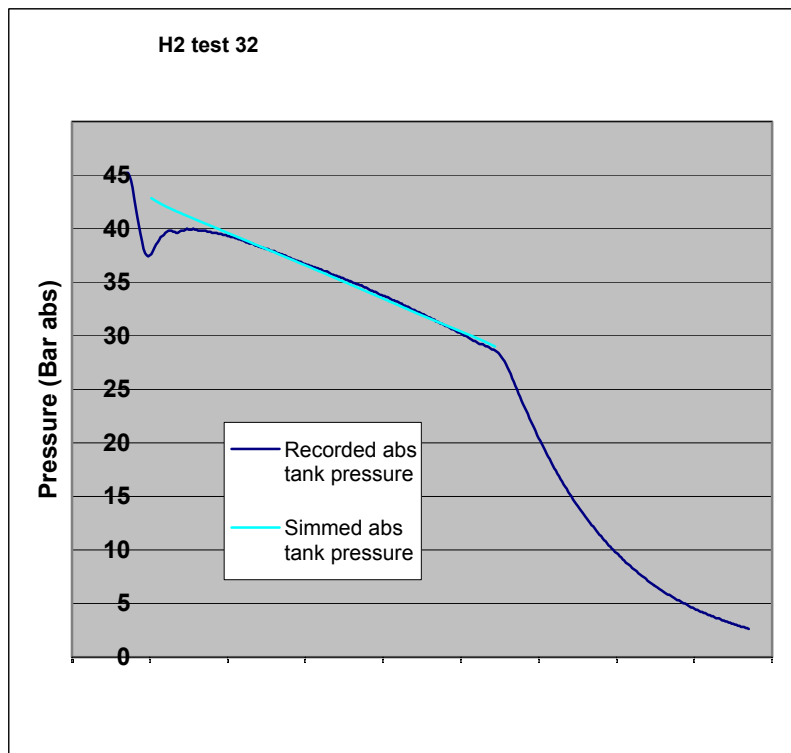
Perhaps this is because the colder nitrous will be denser, so will try to sink to the bottom of the tank and so the liquid gets evenly mixed. Also, changes in pressure affect the whole of the nitrous at once.

Experiments show (Ref. 4 and confirmed by our own) that other effects cancel stratification out, and so the liquid and vapour can be simply modelled as 'blocks' at uniform temperature.

Read our paper 'Modelling the tank emptying' for a mathematical model and software to model the nitrous leaving the tank.

This cooling of the remaining liquid (and therefore any future gas to be vaporised from it's surface as the emptying progresses) means that the vapour-pressure (the tank pressure) will slowly drop over the burn time.

In this graph, burnout was taken as the point when the liquid phase ran out (the graph suddenly steepens)



The lower the pressure drops below vapour-pressure, the more vapour is required to raise the pressure back up, and the more chilled the liquid-phase becomes as it provides this vapour. This is why leaks in any pipe-joints carrying the liquid phase of nitrous oxide show up as regions covered in ice; the nitrous sucks heat out of the atmosphere as it leaks out to atmospheric pressure and vaporises, freezing the water-vapour in the air around the leak.

It'll freeze your eyes, face, or hands too if they're near a leak: **wear goggles and gloves when you work with nitrous.**

So if you crank open the vent (to the atmosphere outside) to huge diameter in an attempt to perform a quick fill, you'll lower the tank pressure *way* below vapour pressure, and so the nitrous will vaporise big-time, chilling itself seriously cold in the process as it drains heat from itself.

If the leak is plugged, for example by shutting a valve on the vent-line, or by shutting the **run valve** mid-burn, liquid will continue to vaporise inside until the vapour-pressure is restored. (albeit the lower vapour-pressure you get at a colder nitrous temperature.) Then as heat from outside *slowly* trickles back into the liquid through the tank walls (this takes a long time, so the tank *does* count as a heat source), the vapour-pressure will slowly rise again until the liquid is back at ambient temperature, then no more heat can flow in. This can take a good 15 minutes even for small run-tanks though.

If the nitrous was originally very chilled (from too fast a fill) an awful lot of it will vaporise during this time, so what started out as a run tank nearly full of liquid might well now be mostly vapour.

The rate of decrease of tank pressure with time (the slope of the graph above) depends on how quickly you empty the tank:

Experiments at Surrey Satellites Technology Ltd have shown that if the nitrous is emptied at a tiny flowrate, less than 10 grams per second or so, then the tank pressure remains constant because the small inflow of heat through the tank wall is just enough.

The vaporisation of the liquid phase into gas is known to resemble conventional 'boiling'. Analogous to the phenomena of supercooling, the boiling of water at atmospheric pressure sometimes doesn't occur at the boiling temperature of 100 deg C.; sometimes the temperature continues to rise higher until some tiny dent or scratch in the container wall (called a 'vapour nucleation site') forms a bubble that breaks loose and sets the wholesale boiling off.



Technical papers

Chemists often drop ‘boiling stones’ (small porous bubble-producing ‘pebbles’) into beakers to ensure that boiling occurs at the temperature expected.

Experiments from Ref. 4 show that mechanical agitation will also trigger boiling in fluid that ought to be boiling but as yet is not.

Once any tiny amount of local boiling kicks in, the resulting bubbles agitate the liquid, greatly increasing the boiling rate, and this feedback mechanism then cascades to produce serious boiling. We see this in our hybrid tests too: the graph shown above is typical, and shows a kink at the start. It seems that when we open our hybrid **run valve**, the initial drop in the liquid level catches the nitrous ‘unawares’, and so there is very little vaporisation, and the ensuing graph of pressure drop has the characteristic steepness of an expanding ideal gas.

Once the hybrid fires up though, the ensuing motor vibrations shake the rocket, and hence the tank. This shaking appears to trigger large-scale boiling, and the tank pressure graph then rises sharply, before descending at the gentler slope that you’d expect from a vaporising subcritical fluid.

On the pad: a quick recap of what all this esoteric physics means to us on the launchpad:

- 1) If you plug the vent-hole after filling to preserve a higher tank pressure and so better performance, your tank better have a head-space or your innocent-looking run-tank may hydraulically overpressure (go boom, see our hybrid safety article) a few minutes after filling.
- 2) It is the liquid phase that we use in the combustion-chamber, so we want to preserve as much of this as possible. Though the gas phase will cause extra thrust after the liquid runs out, its thrust is very small.
- 3) Use a small vent-hole so that the run-tank fills slowly, or a lot of the liquid you put in there will have vaporised by the time you fire it, if it hasn’t all leaked away out the vent.
- 4) If you fill the tank too quickly by cranking open the vent-hole, you’ll over-chill the nitrous, so if you fire it straight away, you’ve got very little tank pressure which will reduce combustion-chamber pressure and so kill most of the thrust.
- 5) If you quick-fill and then wait several minutes before firing, then (assuming you’ve plugged the vent) there will be much less liquid in there than there was 5 minutes ago; it’ll have vaporised into vapour in the tank.
- 6) It may seem cool (sic) to pre-chill the tank to densify the liquid phase to get a lot in there, but you’ll get all the problems due to over-chilling mentioned in 3) 4) and 5).
- 7) If it’s cold outside, warm the run-tank (remotely!).
- 8) If it’s too hot outside, chill the run-tank to keep the liquid density reasonable, or even to prevent the nitrous going supercritical.



Glossary:

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

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

Run-valve: the valve that lets the nitrous flow from the run-tank into the combustion chamber.

Specific Impulse (ISP): the miles-to-the-gallon of a rocket propellant combination as it were, equal to the thrust generated per Newton of fuel used per second (per mass flowrate of fuel times 1gravity).
Units are seconds.

Stoichiometric: the fuel to oxidizer ratio that yields best performance (usually near that for best **Specific Impulse**)

References:

Ref. 1: Engineering Sciences Data Unit (ESDU) sheet 91022,
Thermophysical properties of nitrous oxide.
Available in hardcopy from some U.K. University libraries, or accessible over the Web to students with an ATHENS password.

Ref. 2: University Physics 6th edition
Sears Zemansky and Young
Addison Wesley world student series ISBN 0-201-07199-1

Ref. 3: Space Propulsion Analysis and Design
by Ronald .W. Humble, Gary .N. Henry and Wiley J. Larson
McGraw Hill Space Technology Series ISBN 0-07-031320-2

Ref. 4: 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;
many thanks.