
How fire can be tamed

Graham R.L. Cowan

105, 1144 Division St., Cobourg, Ontario K9A 4J9, Canada

E-mail: gcowan@eagle.ca

Abstract: Combustion dies at the interface between breathable air and macroscopic pieces of certain involatile fuels. If fed only them, in a compressed oxygen chamber, it makes an almost sunlike flame that cannot run wild. If upon dilution to manageable coolness the ash drops to the chamber bottom, and from there can be removed without the diluent, true harnessbrokenness is possible. Excess oxygen can be the diluent without thereby being wasted. It can rid itself of the diluted flame's heat, and spare many trees from becoming newsprint bearing motor fuel mishap reports, by working in a thermodynamic cycle. Some ashes, especially boria, both precipitate well from the diluted flame and travel well. By visiting faraway solar or fission power stations, and returning to the chamber as regenerated fuel, they can make combustion both docile, and subsidiary to docile primary energies.

Keywords: alternative fuels; decarbonisation; energy carriers; energy storage; fireproof fuels; fuel safety; hydrogen economy; hydrogen public acceptance; hypoxia; ignition resistant fuels; nuclear production of motor fuels; solar production of motor fuels; tame combustion; tame fire; zero emission vehicles.

Reference to this paper should be made as follows: Cowan, G.R.L. (2008) 'How fire can be tamed', *Int. J. Nuclear Hydrogen Production and Applications*, Vol. 1, No. 3, pp.235–248.

Biographical note: Graham Cowan obtained a programming diploma from the Honeywell Institute in 1984. He is a researcher in clean vehicle propulsion. His interests include solar and nuclear hydrogen production for use as motor fuel, and direct nuclear propulsion of larger vehicles, including air-cooling of their reactors without irradiation of the air, through the use of low-vapour-pressure oxides as heat transfer media. He has published popular articles and a technical conference paper on these same oxides' potential to be manageable ashes in combustion power systems.

1 Introduction

Where fuel combustion seems most certain to continue and hugely increase, tame or not, is in car motors. Taming it there can be considered to have six parts:

- 1 Liven up the oxygen (LUTO): equip cars with extractors of oxygen from air. (These normally deny a little oxygen to air-breathing bystanders so as to provide it, purified, to something that needs it so. On occasion they must cut that thing entirely off in the bystanders' favour. Accordingly, and to save space, they will be called 'oxygen deniers' or 'O-deniers' from here on.)

- 2 Burners, learn to burn eclectically (BLeBE): choose a fuel whose ash can readily be let out of a chamber full of hot compressed oxygen without letting much of the oxygen come along. The fuel's combustion must propagate in compressed oxygen but die in air at breathable temperature and pressure.
- 3 Star-pyrolyse ash to fuel and oxygen (SPAFO): develop nuclear or solar power stations to import the ash and convert it back to oxygen and fuel, the latter in small uniform pellets shaped for feedability into pressure chambers.
- 4 Strongly confine oxygen in oxide (SCOO): make the inner surfaces of the oxygen chamber of materials that hot, compressed, nearly pure oxygen cannot attack. Include yielding surfaces so that when the fuel of part two burns within it, the oxygen can expand and do work.
- 5 Let engines eat pellets (LEEP): equip cars with bins for the fuel pellets of part 3 and ways for the internal combustion engines of part 4 to take them from the bins.
- 6 Return ash to sender (RATS): Equip cars with bins in which ash can stay until refuelling time, so that it then can be offloaded in exchange for fuel, and begin its journey to the power stations (part 3).

A linear heap of one-carat (0.0002-kg) diamonds probably would satisfy the dying-in-air condition of part 2. Strong heating at one end would cause the diamonds there to burn, but the fire would not travel along the heap. Raising the pressure and, as part one implies, making the oxygen less dilute could change this. But diamond ash is gaseous, and this makes part two's other requirement, that departing ash not take extra oxygen with it, hard to meet.

Also, one variety of diamond ash is carbon monoxide, a gas that kills stealthily when inhaled. Diamond fires, and carbon-oxidising fires in general, are not tame.

The earth's atmosphere normally contains enough of the other diamond ash – carbon dioxide – that the power stations of part three could import all the ash they would need merely by letting it come to them in the wind. Part six would be unnecessary: all the cars, all the power stations, and all of us would be in a single ash bin, the same one that now contains just us and the cars. However, a car motor that burns diamonds in oxygen has another difficulty: at present-day levels of efficiency, about half its power would go to the oxygen denier.

A much smaller share would be taken if the heat engine could get the same heat using much less oxygen. This turns out to be a real option, as potential fuels that are more suitable than carbon in other ways typically also work their oxygen much harder.

2 Oxygen denier readiness

Although today's oxygen deniers leave room for improvement, they are good enough to go on with. A typical one is said to produce 34.5 kg/h of oxygen and mass 1590 kg.

Air sieve materials that let argon pass have been developed, but the portable oxygen deniers now in commercial service use materials that sift out oxygen and argon alike. Their normal product is a mixture with an oxygen mole fraction of 0.945, argon in the same proportion to oxygen as in air, balance nitrogen (Santos et al., 2006).

The theoretical work done by a single-stage compressor in feeding air through them has been assessed in terms of megajoules per cubic metre of product at standard temperature and pressure (STP). Plots of that work versus the production rate show it increasing slowly for a while, then more quickly. For a commercial device using three different sieve materials – Oxysiv 5, Oxysiv 7, MS S 624 – the knees in their plots appear, for the first and last, respectively at 1.60 MJ/(STP m³) and 1.20 MJ/(STP m³). For Oxysiv 7, the greatest work requirement is 1.02 MJ/(STP m³) ('TE' curves, Santos et al., 2006, Figure 4).

If Oxysiv 7 is used and a car motor turns a compressor that indeed has just one stage, but is only 60% efficient, the car motor must provide 1.70 MJ/(STP m³) of shaft work. That is 0.38 kilowatt-hours per kilogram of oxygen.

2.1 *The fate of the argon . . .*

At its simplest, a tame fire chamber would only let gas in, not out. It would admit no gas but oxygen, and as much oxygen as came in would be incorporated into non-gaseous ash and exit in that form.

Revised to use an oxygen denier like those available today, it would admit gases, principally argon, that fire cannot condense. Argon would build up and the pressure would increase. This would snuff the fire in one of a number of ways: increasing the chamber pressure until no more oxygen could be injected, or until the chamber burst, for instance. The fuel will not burn in air.

2.2 *. . . And the resulting waste of oxygen*

At the cost of increasing the oxygen denier's energy demand by about 10%,¹ this problem can be solved. The extra energy would be needed to bring in that much more oxygen. Along with argon, it will be allowed to flow back out of a small gas exit port. Argon's accumulation will thus be limited, and will not prevent continuous combustion.

The oxygen denier that takes 0.38 shaft kWh per produced kgO₂ will therefore need 0.42 shaft kWh per kilogram of oxygen that the fire actually uses.

NIST data imply that a fire whose minus-delta-'G' at 298.15 K is 1 kWh must, if zirconium-fuelled, consume 0.111 kg of oxygen. Most other fuels take more: carbon takes 0.292 kgO₂/kWh, iron, if being oxidised to the (II, III) oxide magnetite, takes 0.226 kgO₂/kWh (Linstrom and Mallard, 2005).

Multiplying these oxygen requirements by 0.42 kWh/kgO₂ yields the fractions of the fire's energy the O-denier would take: for zirconium 0.0466, for carbon 0.123, and for iron 0.0952. But since the O-denier runs on mechanical work, these fractions need to be divided by a typical heat engine efficiency, say 0.3, and for zirconium, carbon, and iron that gives 0.155, 0.409, and 0.317.

2.3 *Capping the oxygen denier's fractional take . . .*

If there is a type of mechanical linkage that can divide the power from a turning shaft in a fixed proportion between two loads, it could ensure that an O-denier that is expected to need 0.317 of the power of the motor it feeds can never take more than 0.32.

The energy cost of extracting oxygen from air rises as the oxygen fraction diminishes, so with such a power-fraction-limiting linkage, an oxygen denier might fall behind the motor's oxygen demand if its intake air began to contain less than 20% oxygen.

2.4 . . . means a motor can starve for oxygen . . .

The motor would begin to lose oxygen pressure, and therefore power. The O-denier, able to take no more than the fixed fraction of the now reduced power, would still not be able to keep up. Soon it and the motor would stop.

2.5 . . . in the midst of plenty

This turns out to be just what they ought to do. People in a freezing house may someday seek warmth in the garage by using there, as a furnace, an air-breathing car that oxidises some carbon-free fuel. If the car has an oxygen denier that cannot, as it deoxygenates the garage's air, take more power to compensate, it will soon live up to its name by cutting off the motor.

It will do this even though, with a less stingy allocation of the motor's output, they both might still have been running hours later, long after hypoxia, unheralded by any rise in carbon dioxide levels, had stealthily and permanently shut down the people.

Portable oxygen deniers' energy needs can in theory still decline by an order of magnitude. If reductions continue incrementally, designers who pair them with heat engines will be obliged incrementally to ratchet down, also, the fraction of the engines' output they can take, keeping it barely over the minimum needed in fresh air.

Calling them oxygen deniers may make this critical denial function harder to overlook.

2.6 *The inert hordes dodged a bullet*

The necessity of thinking about hypoxia arises not just with tame fires but with the encompassing class of carbon-free fuel-air reactions.

After hydrogen cars were impressively demonstrated in the mid to late 1970s, many hoped or believed they were within 5–10 years of catching on. Had this been true, perhaps by the late 1980s all cars would have been safe for prolonged running indoors, but that safety is not intrinsic to hydrogen-air flames nor to fuel cell cathodes.

2.7 *Internal iron combustion power*

If a truck has a 1590-kg oxygen denier that provides 34.5 kg of oxygen per hour, and of that, the fire in its heat engine combines 31.4 kg per hour with iron to form magnetite, it will go:

$$\begin{aligned} \text{Net power} &= (31.4 \text{ kg O}_2 \text{ per hour}) * \\ &\quad (\text{heat engine efficiency/mass of oxygen required per unit energy yielded by fire} \\ &\quad - (0.42 \text{ kWh per kg O}_2)) \end{aligned}$$

although perhaps not very quickly. Using 0.3 for heat engine efficiency, the above yields for carbon 19 kW (26 hp), for iron 28 kW (38 hp), and for zirconium 72 kW (96 hp).

Dividing the fire's oxygen consumption rate by the net power yields oxygen consumption per net shaft kWh: carbon takes 1.65 kgO₂/kWh, iron takes 1.12 kgO₂/kWh, and zirconium takes 0.44 kgO₂/kWh.

3 Burners, learn to burn eclectically (BLeBE)

The fuel chosen should burn readily enough in purified compressed oxygen, but be non-ignitable in breathable air. If pulverisation would make it ignitable in air whereas in palpably large pieces it is not, then in a fire-taming effort one would not pulverise it. The side effects of making it inhalable and hard to clean up when spilled, would also be avoided. Similarly, one would not have it generate another fuel, one that burns readily or explosively in air, by taking oxygen from that other fuel's ash.

Aluminium is the best-known fireproof fuel, but there are others. They include beryllium – toxic and very expensive – and boron and silicon. They include heavier elements, such as the above-mentioned zirconium, that clothe themselves in oxide films that strongly impede ignition. If a net shaft kWh from a zirconium-burning vehicle motor requires 0.44 kg oxygen, which over time averages roughly half in the ash reservoir, half in the atmosphere or in transit back to it, then that kWh also requires 1.23 kg zirconium, which is always on board.

With boron, aluminium, and silicon the always-on-board masses are, respectively, 0.276, 0.483 and 0.485 kg, and the combining oxygen masses are 0.613, 0.430 and 0.552 kg, per net shaft kWh. This again is on the basis of 0.3 heat engine efficiency and an oxygen denier that requires 0.42 kWh per combining kilogram of oxygen. Unless for some heavy element there is a compelling reason to accept the extra mass, tame combustion should be fed one of these three.

The net-power calculation of the previous section predicts 51 kW for boron, 73 kW for aluminium, and 57 kW for silicon. The truck mentioned there could keep up with traffic if it were burning any of these.

3.1 Ability of ash to precipitate from diluteness in hot oxygen

An ash that fails to do this, and is produced in a flame much hotter than a solid pressure envelope can bear, will not lend itself to that flame's dilution down to bearable temperature with excess oxygen, for letting the ash out will tend to entail also letting out, and wasting, the oxygen.

If, however, cooling due to dilution causes the ash to condense and fall out, a quantity of oxygen can repeatedly do three things – dilute ash, drop it, and rid itself of the heat the ash gave it by acting as heat engine working fluid – before in its turn being incorporated into ash and falling to the vessel's bottom.

Among combustibles that will not burn in air, few present any difficulty for the ash-fall part of this scheme; maybe only osmium, since its tetroxide's normal boiling point is 403 K. If there are osmium-burning-car enthusiasts, their dream will not be advanced by the taming of fire.

Of elemental boron, aluminium, and silicon, none is screened out, for their ashes – boria, corundum (aka alumina), and silica – all have normal boiling points well above 2000 K. By the time their flames in oxygen have diluted themselves below this temperature, they have become smoke plumes, upwellings of hot oxygen in which involatile ash particles fall slowly.

3.2 *Selective removability of precipitated corundum*

If the purpose the smoke is to serve is expanding and pushing a piston in a cylinder, this typically needs to be repeatable with the same piston and cylinder about 100 million times: thousands of times per minute for thousands of hours. This will not happen with corundum-bearing oxygen, as corundum is abrasive. In an environment of hot dusty oxygen that is just beginning to do its work of expansion in an efficient automotive heat engine, at the temperature of the surface upon which the oxygen presses, corundum is harder than any other material that would be stable.

That surface could well be made of corundum. Strong confinement of oxygen in oxide (SCOO) calls for a strong oxide. Corundum is a strong oxide. But if they are not to scratch it, ash particles in the oxygen must be made of something softer.

Aluminium-burning rocket motors embody another solution to this problem – accept a short lifetime. But there may be a way an aluminium burner can live a long life, despite expelling corundum dust throughout it: let aluminium fill its combustion chamber, as liquid, and burn oxygen in an inverse flame in mid-metal.

Liquid aluminium has a history of being fairly successfully managed, including travelling hundreds of kilometres on public highways as 5-m³ crucibles-full. At departure from a foundry such a load is at 1150–1200 K. Ideally it has not yet cooled below 1050 K, more than 100 K above its freezing temperature, when it arrives at a casting site.

It seems reasonable that compressed oxygen injected from non-oxidisable nozzles deeply submerged in pressurised liquid aluminium would be consumed near the nozzles, none of it reaching the upper surface, and that the combustion products would initially rise through the liquid as bubbles, but quickly lose heat to it and collapse. Farther up, the flame would be one of hot liquid metal carrying corundum particles up through cooler liquid metal.

Compared to a corundum-bearing updraft of oxygen, this other convecting fluid offers the advantage that a material exists – aluminium diboride – that promises to resist both it (through being saturated with it) and its corundum particles (through possibly superior hardness). Although combustion would be internal to the liquid metal, it would be external to aluminium diboride heat exchanger tubes placed in the updraft's way, so the system would be an external combustion engine. The working fluid in the tubes could not be oxygen, but it would not have to be.

At the chamber bottom, where the liquid aluminium would be relatively cool, there could be a drain with a screw impeller in it, both also made of aluminium diboride, in order to let the precipitated corundum particles and the liquid metal entrained by them exit as a sludge or paste. The drain would give onto a chamber where aliquots of the sludge would fall onto a draining surface. As much liquid metal as might readily drain out of them would do so, and be returned to the upper chamber. The aliquots would then be cooled and put in an ash ingot bin.

This would amount to somewhat unselective ash removal because some liquid aluminium between particles would get frozen into the ingots. Per net shaft kWh, more than the previously noted 0.483 kg of aluminium would move from fuel bin to ash bin. By helping to hold the ash ingots together, the unburned metal would aid in returning ash to sender (RATS).

3.3 Selective removability of precipitated boria and silica

Boron melts above 2000 K and silicon at 1687 K. Neither, when burned with pure oxygen, can usefully be present in excess to dilute the flame, for neither shares aluminium's ability to remain fluid when cooled to temperatures well below the maximum a heat engine's solid parts can long endure. Because of this, they cannot coolly surround a flame that starts out much hotter than that maximum and through turbulent mixing form a diluted flame whose temperature is just right. Of the flame's two reagents, only the oxygen can do this. The result is a smoke of boria in oxygen or silica in oxygen.

Either smoke's expansion would be less rough on a heat engine than that of corundum-laden oxygen. Silica, whose crystals melt at 1983 K, is, like every non-corundum ash that might be so suspended, less abrasive than corundum. If it originally was vapour and has become a suspended condensate through being diluted and cooled by the oxygen around it, boria exists as droplets, which supercool. When they land on a surface that is hotter than about 0.8 of boria crystals' 723-K melting point, they coat it. Water droplet impact in engines powered by wet steam can be destructive, but boria droplets are very different. A film formed by earlier-arriving boria droplets can so cover a surface that latecomers cannot abrade it at all.

On oxygen-boron combustion chamber walls warm enough for accreting boria droplets to merge into a film, the film will flow. Along with direct droplet precipitation, this flow can create a chamber-bottom boria lake. A dynamic equilibrium can be established where boria comes to the lake's surface and, at the same rate, removes itself, plus as much oxygen as was soluble in the surface layers, through a lake-bottom drain. This removal is selective because gaseous oxygen cannot get down there.

If in the drain the boria has a temperature of 873 K, somewhat cooler than liquid aluminium, it is runny enough to drain at a useful rate. Unlike a paste of corundum and liquid aluminium, it is entirely liquid and unabrasive. It needs no screw impeller to help it along.

Its dynamic viscosity and density, represented below by ' μ ' and ' σ ', are 480 pascal seconds and 1608 kg/m³ (Smith and McBroom, 1999). A 0.25-m-deep lake of it that is held at a combustor bottom by its standard terrestrial weight of 9.80665 N/kg will have a top-to-bottom hydrostatic pressure difference, represented below by ' Δp ', of 3942 Pa.

A laminar flow calculation that neglects end effects:

$$\text{Mass throughput} = \pi \sigma D^4 * (\Delta p) / (128 \mu L)$$

predicts that boria will exit through a circular hole in the bottom at a rate of 0.0189 kg/s if the hole's length L and diameter D are both 0.0388 m, and its walls are at the same 873-K temperature as the liquid.

In preparation for return of ash to sender – RATS – the boria extruding below the bottom of the drain could periodically be cut off. The now detached gob could be cooled by air flow, and then, with a little more air flow, blown into its bin.

Raising the temperature by 200 K could speed the flow up about 18-fold. A spun lake, held to its bed by centrifugal force rather than planetary gravity, could drain much faster still. But the 0.0189-kg/s oozing, with an average flow speed of 0.0100 m/s, turns out to be adequate for draining the small volumes per unit time that a car-scale combustor would produce. According to the previously noted per-net-shaft-kWh boron and oxygen

masses, which sum to 0.889 kg, the net driveshaft power that could be co-produced is 76 kW.

For silica that mass sum is 1.037 kg, highest of the three fuel-ash pairs under consideration. Since corundum is more abrasive than silica, conceivably an engine powered by expansion of hot compressed silica-bearing oxygen could be internally surfaced with corundum, and not be eviscerated by the silica particles. But wherever they settled, they would tend to stay. They would not be able to merge and flow. If they did merge, it would amount to sintering. Perhaps they would form tough artificial quartzite or vitreous silica scale on heat engine surfaces. Along with the extra mass, this prospect makes internal silicon combustion relatively uninteresting.


4 Central station deoxidation

Billions are determined to become motorists. One way or another, motor fuel production rates will rise by many terawatts over the next few decades. They can be increased cleanly by constructing nuclear or solar power stations that take in alumina or boria ingots and remove the oxygen. Such stations could be rated in terms of aluminium terawatts, TW(Al), or boron terawatts, TW(B).

If each station were to produce 0.1 TW(Al), it would make about three times as much as all of today's smelters. Since this is roughly a two-orders-of-magnitude increase in individual plant scale, by the inverse-square-root-of-scale rule (Marchetti, 2006), product unit cost would decline by one order of magnitude, from around US\$0.30/kWh(Al) down to about US\$0.03/kWh(Al). In aid of this cost reduction, inert anodes might be put into service (Welch, 1999), so that where now a significant fraction of the energy stored in aluminium comes from oxidising carbon anodes, scaled-up plants would need only electricity.

A 0.1-TW(B) power station might need only heat, provided heat of a certain intensity was available. Boria follows the general rule for very stable oxides that mere direct heating, however intense, will not cause macroscopic separation into fuel and oxygen. Star pyrolysis of ash to fuel and oxygen (SPAFO) yields fuel and oxygen atoms. They remain mixed, and will not refrain from reattaching to each other if one cools the vapour down to a temperature where, if they would so refrain, separation might conveniently occur.

However, heat from a source significantly hotter than 2500 K can usefully act on another oxide that is less stable, and – being an ore of iron – very much cheaper and more abundant: magnetite (Ehrensberger et al., 1997; Mohai et al., 2007):




For $x=0$, NIST data imply this process has enthalpy change +372.3 kJ/mol, plus another 39.2 kJ/mol that the oxygen would give back in being cooled from 2500 to 298.15 K. Being liquid, the ferrous oxide tends to separate from the oxygen, so they can be cooled without recombining.

Losing the 39.2 kJ would be reasonable for a solar power station that focussed a large image of the sun down onto a high-altitude outdoor stream of magnetite, for then the half-mole of oxygen could go directly into the upper air. If such a station annually turned 32.6 billion kg of magnetite into 1.9 billion kg of oxygen and 30.7 billion kg of ferrous oxide, its annual average output could be expressed as 1 GW(FeO).




Where summer is much sunnier than winter, ferrous oxide production rates in winter, spring, summer, and fall might average respectively zero, 1, 2, and 1 GW(FeO). By summer's end, 7.7 billion kg of ferrous oxide, a gigawatt-season's worth, could accumulate, perhaps as an outdoor conical heap 300 m across the base. If a steady year-round ferrous oxide gigawatt were taken, the iron by winter's end would be in a slightly larger magnetite pile. Other kinds of gigawatt-season energy reservoir – two billion lead-acid car batteries, a cubic km of water raised 800 m – are larger or more costly or both.

Boria, like corundum, can dissolve in molten salts. Converted to 300–400 MW of electricity, the GW(FeO) could electrolyse it there year-round, but it may also be able to react directly with ferrous oxide and nitrogen, in a way that corundum cannot:




This condenses nitrogen without producing any other gas, but is exothermic and spontaneous at room temperature and pressure.


Raising the nitrogen pressure to a few tens of MPa should make it go at 500 K. Some compound of boria and ferrous oxide may be so stable as to sidetrack it, or it may still be too slow. It should be tried. If boron nitride and magnetite are seen to be producible quickly enough, and they can be separated, three more spontaneous reactions will lever boron up to freedom:




The above would transform the problem of deoxidising boron into one of converting magnetite and ferrous halides back into free halogens, ferrous oxide, and iron. This highly exothermic reaction would free the halogens:



Iron could be produced by consuming more ferrous oxide.



Finally, the magnetite could become oxygen and ferrous oxide,



and the ferrous oxide could freeze, and if $x=0$, 29 moles of ferrous oxide would have been oxidised to 29/3 moles of magnetite and then recovered. Magnetite's above-noted 372.3-kJ/mol enthalpy of partial deoxidation, times 29/3, gives the delta H of this step:

3.98 MJ/mol if oxygen is let go while still hot, 3.60 MJ/mol if its heat is recovered. The net result would be the dissociation of one mole of boria, $\Delta H +1.25336$ MJ. Of 3.98 MJ, that is 31.5%, about the same as if electricity had been made, but boria can be wirelessly gathered in, and boron wirelessly distributed.

The magnetite heating could occur in an annular curtain of falling particles bathed in rising high-pressure helium. Each particle would begin its fall as a magnetite particle. As it descended it would give up oxygen until it was a ferrous oxide particle. The oxygen would be swept upwards by the helium.

The heat could be radiant heat from a central fountain of chunks of an actinide element or a mixture of actinides. Fission in the chunks would melt and evaporate them so that they would merge as they approached the apices of their free flights. Because their vapour would always be much denser than the surrounding helium, each chunk's expanded remains, although flying upwards as the chunk had been, would still be decelerating and soon would fall back. The structure would be a dense vapour fountain.

Continuing fission in the vapour slug a chunk had become, would raise its temperature. The resulting expansion would reduce its opacity to neutrons, so the process would limit itself. As it fell back, later-arriving chunks, still condensed, could pass through it.

Surrounding this fountain and reflecting neutrons to it, would be a region filled with cool flowing helium. The high transparency of helium to thermal radiation would allow it to fill the space between the fountain and the sheath of iron oxide particles while remaining cooler than either. Its flow would continuously remove the dense vapour fountain's outer surface, the remains of the longest-serving chunks, and carry it downwards.

Compared to operation at normal atmospheric pressure, the high pressure would serve two purposes. One is to enable the cool helium to promote chain fission by converting fast neutrons that escape the fountain into cool thermal ones before they have gone far, so that they have a good chance of diffusing back into it.


If nuclei of natural uranium are the only ones in the vapour, these returning cool neutrons are more likely to cause fission than in the case where they are captured while still fast, because they are more likely to avoid capture by nuclei of the 99.28% non-fissionable majority, and instead be captured by the fissionable minority. If the surrounding helium has to cool too many neutrons, and its temperature increases over time, so does theirs, and their selectivity diminishes. Therefore, in raising its own temperature, this reactor would reduce its moderator's effectiveness, and in so doing, would reduce its own power. This self-regulating tendency is one it would share with existing nuclear power reactors.

The second purpose pressure would serve is to raise the temperature at which the actinide would boil, and thereafter reduce the rate at which it would expand, so that it could stay near the centre longer, reach a higher temperature, and shine more strongly on the iron oxide.

The cascade of actinide vapour and the sheath of helium, neutron-heated nearest the vapour but still cool farther out, would together descend through a ring of nozzles that would douse them with fluoride particles. This would condense the actinide vapour, making the heat earlier stored in it unusable for magnetite pyrolysis, so ideally, for efficiency's sake, each piece of vapour would shine away much more energy during its time of high luminosity than it would then spend in this manner.

The fluoride particles might be 16 mole parts sodium fluoride, four parts potassium fluoride, and five parts magnesium fluoride, a composition sometimes referred to as

NaF-16KF-20MgF₂. It melts at 1077 K (Misra and Whittenberger, 1987). The amount thrown into the descending flame would be adjusted so as to end up above this temperature, in the hope that the particles of actinide soot would end up inside its droplets, and there give up some fission fragments to it (Lemort, 1997):



Gravity and a small pressure difference, maybe 1 kPa, together would pull the droplets, any unincorporated actinide soot grains, and the helium around them all down into a lower chamber. There they would impinge straight down upon the centre of a broad pool of liquid magnesium near its 923-K freezing point. Falling into this, the fluoride droplets would freeze into beads. These beads and the actinide soot, some of it inside them and isolated from the liquid, would fall through it and come to rest on a layer of older beads.

The beads and soot would contain fast-decaying fission fragments. An influx of pieces of solid magnesium would be maintained, therefore, to keep the liquid magnesium near its freezing point. They would sink through it alongside the beads and lie with them on the bead layer, and be melted by decay heat there. As much magnesium as was being added as solid would be removed from the surface as liquid and dropped through helium, which would take away the heat, probably to a heat engine.

The oldest, least heat-producing beads and soot grains would be at the bottom of the bead layer. Under them would be a deep foundation of solid NaF-16KF-20MgF₂. It would extend far enough down to provide a long-term backup heat sink. After solid magnesium pieces have stopped being thrown in, and all present have melted, the liquid magnesium warms to 1077 K and the beads immersed in it melt. They form a liquid NaF-16KF-20MgF₂ layer between the solid fluoride and the magnesium. The actinide grains formerly trapped in them fall onto the solid fluoride. Their continuing but diminishing heat production drives the liquid-solid interface downwards. Occasionally it reaches pockets of magnesium and causes them to drain upwards. Eventually the inward leakage of the outside world's coolness overwhelms the fission fragments and the melt front becomes a freeze front. It rises above the grains, trapping them.

Usually the beads would not be allowed to melt and a dredge could selectively dredge the deepest-down, least heat-producing ones off the solid NaF-16KF-20MgF₂ surface. The liquid magnesium could be evaporated off them, and then they could be pulverised under liquid xenon. NIST tables say natural air-derived xenon at 175 K is dense enough under its own 173.25-kPa vapour pressure to float NaF-16KF-20MgF₂ bead fragments but not actinide soot grains (Lemmon et al., 2005).

An alternative fluoride composition, NaF-13MgF₂-22CaF₂, melts at 1027 K (Misra and Whittenberger, 1987). That is 50 K lower than NaF-16KF-20MgF₂, but still higher than magnesium's freezing point, so much the same actinide-soot-gathering and bead-forming behaviour can be expected of it. However, its floatability in liquid natural xenon, even at the latter's triple point, is uncertain. A better bet for floating it would be the large amounts of slightly denser xenon that can be extracted from existing commercial reactors' spent fuel.

The actinide soot would be taken out from under the liquid xenon and the fluoride bead fragments skimmed off the top. The soot would be pressed into chunks. Flung again up the fountain's axis, they would again become bright aliquots of vapour. The fluoride would be used to cool the fountain's effluent again.

Helium in the gas lying over the liquid magnesium would be cooled, filtered, and returned to its place surrounding the fountain and enabling it to shine. Some fission fragments' inability to pass through the helium filters, and others' binding by the fluoride, are what make requiring the fuel to boil itself reasonable: those fission fragments not incorporated when it condenses are promptly trapped in a trap good enough to be their final resting place.

5 Summary

Fuelling combustion only with pellets of aluminium or boron, small but not so small as to be difficult to see and handle individually, can tame it. Compared to burners of oil in air, internal boron combustion cars and external aluminium combustion ones would both significantly advance safety, as spilled aluminium or boron pellets would not be able to burn in the dilute, low-pressure oxygen we live on.

Non-combustion primary energies can supply aluminium and boron by extracting them from their oxides. These would conveniently be available as tame combustors' well-consolidated ash ingots, as the same use of purified oxygen that overcomes these fuels' ignition resistance also makes it easy to retain their ashes.

Cars would need machinery to make ingots and bins to hold them. They would need oxygen deniers. These O-deniers would reduce the engines' net efficiency by taking some of their output, and they and the ash bins would make cars larger and heavier, other things being equal, than oil burners.

5.1 Implications

The cars would advance environmental stewardship: ashes would not be emitted to the atmosphere. Chemically unaltered air would carry waste heat away in a stream that would probably be directed aft, and the oxygen denier would emit another relatively slow, relatively narrow stream of slightly oxygen-depleted air, probably upwards.

The cars would advance ease of use. Letting the ashes of hydrocarbon motor fuels simply blow away can be easier than dealing with ash ingots, but certain ways of dealing with the fuels themselves, ways that would otherwise be convenient, would cause frequent severe accidents. Like the severe but infrequent harm of post-crash fuel-fed fires, these large risks can be reduced to zero, and the potential conveniences made actual, by a switch to innocuous fuel that will not burn.

People who carry flasks of reserve fuel in their cars may appreciate the following illustration. The low-fuel light has been on for some time. The car falters. Reaching into the back seat without looking, one raises the reserve flask's lid and scoops out a handful of aluminium or boron pellets. The car falters again, and in haste one fumbles some of the handful onto the cabin floor. The rest are successfully dropped into the in-cabin refuelling port. They roll down its throat into the main fuel bin in the front bumper. No further hesitation occurs; additional handfuls cause the low-fuel light to go out. The dropped pellets can be let lie until one is at one's destination, then gathered up and put back in the reserve flask.

Compared to those of liquid hydrocarbon-burning cars, zero-local-emission cars' stores of propulsion energy have been significantly less convenient to replenish and much smaller. The bargain they have offered early adopters has not been cachet for a price, but

cachet for two prices: pay more and tolerate more inconvenience. Early adopters of tame combustion cars would also have to pay more, but convenience would become one of two benefits rather than one of two costs.

Having two refuelling ports on a car, one inside the cabin and one outside, is one such potential convenience. Another is the option of buying the car and its whole lifetime supply of about 10 m³ of fuel at the same time. Neither geometry nor prudence would forbid storing the fuel in one's 200 m³ basement apartment, and refuelling at home.

The ash ingots would not be so compact. Home refuellers could accept the inconvenience of packaging them and shipping them off to power stations, but fuel retailers would soon see an opportunity in doing this for them. Since the aluminium or boron pellets would be uniform, durable, and non-ignitable in air, a fuel bin wall ventilated with holes too small for them to fit through would let a retailer blow them in with an air hose. The ash ingots might need two hoses, one to vacuum and one to blow. This and the fuel transfer could be done simultaneously. The whole three-hose process could be quick.

5.2 Prediction

When developed, tame combustion cars will promptly take over the high end of the car market and then push hydrocarbon-air combustion progressively down and out.

References

- Ehrensberger, K., Palumbo, R., Larson, C. and Steinfeld, A. (1997) 'Production of carbon from CO₂ with iron oxides and high-temperature solar energy', *Industrial & Engineering Chemistry Research*, Vol. 36, pp.645–648.
- Lemmon, E.W., McLinden, M.O. and Friend, D.G. (2005) 'Thermophysical properties of fluid systems' in P.J. Linstrom and W.G. Mallard (Eds), *NIST Chemistry WebBook, NIST Standard Reference Database Number 69*, National Institute of Standards and Technology, Gaithersburg MD, 20899, June. Available at: <http://webbook.nist.gov>.
- Lemort, F. (1997) 'Étude de la Séparation Actinides-Lanthanides des Déchets Nucléaires par un Procédé Pyrochimique Nouveau', Commissariat à l'Énergie Atomique France Rapport CEA-R-5760, pp.173–174. Available at: <http://www-ist.cea.fr/publiccea/exl-doc/00000036953.pdf>.
- Linstrom, P.J. and Mallard, W.G. (Eds) (2005) *NIST Chemistry WebBook, NIST Standard Reference Database Number 69*, National Institute of Standards and Technology, Gaithersburg MD, 20899, June. Available at: <http://webbook.nist.gov>.
- Marchetti, C. (2006) 'Long-term global vision of nuclear-produced hydrogen', *Int. J. Nuclear Hydrogen Production and Applications*, Vol. 1, pp.13–19. Available at: <http://www.inderscience.com/storage/f126210871115349.pdf>.
- Misra, A.K. and Whittenberger, J.D. (1987) 'Fluoride salts and container materials for thermal energy storage applications in the temperature range 973 to 1400 K', NASA Technical Memorandum 89913, AIAA-87-9226. Available at: http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19870014593_1987014593.pdf.
- Mohai, I., Gál, L., Szépvölgyi, J., Gubicza, J. and Farkas, Z. (2007) 'Synthesis of nanosized zinc ferrites from liquid precursors in RF thermal plasma reactor', *Journal of the European Ceramic Society*, Vol. 27, pp.941–945.
- Santos, J.C., Portugal, A.F., Magalhães, F.D. and Mendes, A. (2006) 'Optimization of Medical PSA Units for Oxygen Production', *Industrial & Engineering Chemistry Research*, Vol. 45, pp.1085–1096.

Smith, R.A. and McBroom, R.B. (1999) 'Physical Properties of Vitreous Boric Oxide', article 'Boron Oxides, Boric Acid and Borates', *Kirk-Othmer Encyclopedia of Chemical Technology*, Table 4, 4th edition.

Welch, B.J. (1999) 'Aluminum production paths in the new millennium', *JOM*, Vol. 51, No. 5, pp.24–28. Available at: <http://www.tms.org/pubs/journals/JOM/9905/Welch-9905.html>.

Notes

¹ One of the things determining how many kilograms of oxygen an O-denier must supply per kilogram a tame fire uses is the minimum allowed oxidiser fraction in the tame fire chamber, represented below by t . Air's 0.21 is by definition significantly less than t .

Supposing $t=0.6$, assuming nitrogen behaves as an inert gas, and setting the oxygen denier product's oxygen, argon, and nitrogen mole fractions x_{O_2} , x_{Ar} , and x_{N_2} respectively to 0.945, 0.042, and 0.013:

$$x_{O_2} / (x_{O_2} - t * (x_{Ar} + x_{N_2})) / (1 - t),$$

yields a value near 1.096.

Nitrogen in the chamber does not behave exactly as an inert gas, however. Were it to show complete non-inertness, every molecule combining with two oxygen molecules to form two nitrogen dioxide molecules, it would make its volume and diluting effect zero by hiding itself in the oxidiser. x_{N_2} would disappear from the ratio expression:

$$x_{O_2} / (x_{O_2} - t * x_{Ar} / (1 - t))$$

For the same t and mole fractions, that is around 1.071.

A catalytic converter might be needed in the gas exit channel to deal with nitrogen dioxide. It would have about 75 times less gas to scrub than its counterpart in an oil/air-burning car, so a bed of alkali, although consumable, might serve well enough instead.