

# *The Efficiency of Rocket Fuel*

Chengkai Wang

*The Affiliated International School (YHV) of Shenzhen University, Shenzhen, China*  
*wshero@vip.163.com*

**Abstract.** Selecting a chemical propellant for a launch vehicle isn't just about looking up specific impulse values in a handbook. Vacuum  $I_{sp}$  matters a great deal for upper stages, obviously, but for a reusable booster that has to fly again next week, a bunch of other considerations start to dominate. Things like tank volume, whether the fuel gunks up the cooling channels, and how much time the ground crew spends scrubbing things between flights. This article walks through the main propellant types—liquids, solids, hybrids—with a focus on the trade-offs that actually matter in a high-cadence operation. Liquid methane gets a lot of attention here because it sits in an interesting middle spot: denser than hydrogen but not nearly as dense as kerosene, and it burns clean enough that coking is less of a pain. It's not the absolute best at any one metric, but the overall balance is what makes it attractive for the current generation of reusable rockets.

**Keywords:** Rocket Propellant Efficiency, Specific Impulse, Reusable Launch Vehicles, Green Propellants, Methane Propulsion

## 1. Introduction

The rocket equation has been around for a long time—more than a century, in fact. It's simple to write down:  $\Delta v = I_{sp} \times g_0 \times \ln(m_0/m_f)$ . Everyone who designs a launch vehicle has to live within the limits it sets. For decades, the main focus was on squeezing out every possible second of specific impulse. If that meant dealing with gigantic tanks or hazardous chemicals, well, that was just accepted as the price of doing business.

But the landscape has shifted in the last ten or fifteen years. Rockets are launching on a weekly cadence now, and boosters are routinely flying back to land on drone ships or concrete pads. That changes the calculus. A propellant that delivers excellent vacuum performance but forces you to partially tear down the engine and scrub out carbon after every few flights? That's not really a winning formula anymore. Turnaround time—sometimes just called "hangar time"—has become a first-order economic driver. The mass of all the ground support junk matters too. And helium. Helium is surprisingly expensive, and supply chain hiccups have been known to cause headaches for launch campaigns. This is the kind of thing that doesn't show up in a textbook performance table.

If you look at the operational rockets flying today, there's a pretty clear split. Upper stages almost all use liquid hydrogen with liquid oxygen. First-stage boosters, on the other hand, have historically used RP-1 kerosene. This split is not an accident. Hydrogen gives you extremely high exhaust velocity, which is exactly what you want once you're out of the thick lower atmosphere. But

hydrogen is a pain to handle operationally—the boiling point is 20 K, the tanks are enormous for the mass they contain, and boil-off during countdown is a constant nuisance. Kerosene, by contrast, is dense and relatively easy to work with, but it has a well-documented drawback: it tends to form solid carbon deposits inside the narrow cooling channels of a regeneratively cooled engine. Sutton and Biblarz cover these fundamental trade-offs in their textbook [1].

The push toward reusability—I mean real reusability, where the same booster flies ten or twenty times—has really put the spotlight on that coking problem. If an engine is supposed to accumulate hours of hot-fire time without a major overhaul, carbon buildup becomes a serious obstacle. This paper looks at propellant options from that operational angle, focusing on what happens after the engine shuts down, not just during the burn.

This paper looks at propellant options from that operational angle, focusing on what happens after the engine shuts down, not just during the burn. It should be noted that detailed refurbishment data is proprietary, and the analysis that follows is based on publicly available presentations and inferred industry practices.

## 2. Why $I_{sp}$ alone isn't a great metric for boosters

Specific impulse, usually given in seconds, measures the thrust you get per unit weight flow of propellant. It's a handy number. For a vacuum upper stage, every extra second of  $I_{sp}$  translates almost directly into more payload—or less propellant needed for the mission. The RL10B-2, an expander-cycle engine that burns liquid oxygen and liquid hydrogen, gets about 462 s in vacuum. That's about as good as it gets with chemical propulsion [2]. For a geostationary transfer mission, a hydrogen upper stage can save something on the order of 20 tonnes of propellant compared to a kerosene stage delivering the same payload. That's a real mass saving, and it cascades through the rest of the vehicle design.

But for a first stage, the situation is different. A booster climbing through the dense lower atmosphere has to fight two things: gravity and aerodynamic drag. Gravity losses are lower if you can accelerate hard off the pad. Drag losses are lower if the vehicle cross-section is small. Both of these point you toward dense propellants—fuels that pack a lot of mass into a relatively small volume.

Liquid hydrogen is not dense. At its boiling point, the density is only about 71 kg/m<sup>3</sup>. A typical LOX/RP-1 combo, on the other hand, has a bulk density of roughly 1030 kg/m<sup>3</sup>. That's a factor of fourteen difference. For a fixed propellant mass of 100 tonnes, a hydrogen tank needs something like 1410 m<sup>3</sup> of internal volume, while an RP-1 tank needs only about 97 m<sup>3</sup>. The larger tank means more surface area, more structural mass, and a larger frontal area that increases drag during the transonic and max-Q portions of ascent.

There's a historical example that illustrates this pretty well. During the early design work for the Saturn V's S-IC first stage, the team briefly looked at using hydrogen for the booster. They ran the numbers and quickly dropped the idea when it became clear that the tank diameter would need to exceed 12 m—too big for the tooling and transport setup at the Michoud Assembly Facility. So the F-1 engines were designed around RP-1 instead. The decision was driven by volume, not by  $I_{sp}$ .

A more useful metric for booster propellants is density-specific impulse, which is just the product of average bulk density and specific impulse. For hydrogen, this number is something like 32,000 in SI units (kg·s/m<sup>3</sup>). For kerosene, it's roughly 360,000. Methane lands in between, at around 158,000. It's an intermediate position in this regard—not as bulky as hydrogen, not as dense as kerosene. But as it turns out, that middle position is a pretty good place to be for a reusable booster.

And then there's the coking thing. In a regeneratively cooled engine, the fuel is circulated through passages in the combustion chamber wall before being injected. This keeps the wall from melting and recovers some of the heat. Kerosene, though, is a complex mixture of hydrocarbons—paraffins, naphthenes, aromatics—with an average carbon number around twelve. When it gets hot, above roughly 560 K under typical cooling-channel pressures, the long-chain molecules start to crack. This process leaves behind solid carbon—coke—on the channel walls.

The coke layer acts as thermal insulation. It reduces heat transfer from the wall to the coolant, which causes the wall temperature to rise, which in turn accelerates further decomposition. It's a feedback loop. Naraghi and co-workers looked at using dual-circuit cooling to spread the thermal load and lower peak wall temperatures [3]. That helps, but it doesn't change the underlying chemistry. For an expendable engine that burns for a few minutes and then takes a swim in the ocean, this degradation is acceptable. For an engine that's expected to rack up hours of hot-fire time over dozens of flights, the progressive buildup of carbon becomes a real maintenance headache. Somebody has to go in there and clean out the injector face and the cooling channels. That takes time and money.

The 560 K threshold is a useful reference point, but the actual onset of coking depends on a combination of pressure, residence time, and the specific hydrocarbon distribution within the RP-1 specification. As RP-1 is defined by a performance specification rather than a fixed molecular composition, batch-to-batch variations in coking propensity are observed in practice (Figure 1).

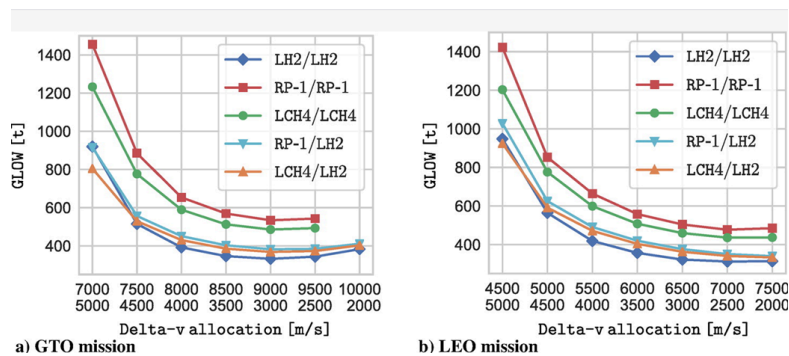


Figure 1. Comparison of bulk density and specific impulse for common propellant combinations

### 3. A look at the main propellant families

Rocket propellants are usually grouped by their physical state and how they're stored. Each category has its own set of advantages and headaches. I'll go through them one by one.

#### 3.1. Liquid Bi-propellants

LOX and liquid hydrogen give the highest specific impulse of any practical chemical combination. The exhaust is light—mostly water vapor with some unreacted hydrogen—so the exhaust velocity is high. This makes LOX/LH<sub>2</sub> the standard choice for upper stages and deep-space departure burns. The Centaur upper stage has been flying in one form or another since the 1960s, and it still sets the benchmark for high-energy missions.

But liquid hydrogen is operationally challenging. Boiling point is twenty Kelvin. Even with good insulation, you lose a few percent per hour to boil-off. The Shuttle's External Tank had to be continuously topped off until the final minutes of the countdown. And the tank itself is huge. All that volume translates into more surface area, more insulation mass, and a heavier structure overall [1].

LOX and RP-1 has been the booster workhorse for decades. Saturn V's F-1 used it. Falcon 9's Merlin uses it. The high density keeps the tanks compact, and the turbopumps don't have to work as hard because they're moving less volume. The main downside, as already discussed, is coking in the cooling channels. For a partially reusable vehicle like Falcon 9, the Merlin engines get inspected and cleaned after some number of flights. The exact refurbishment interval isn't public, but it's definitely part of the routine.

Liquid methane with liquid oxygen—often just called methalox—has moved from the lab to operational hardware pretty recently. Density at the boiling point is about  $420 \text{ kg/m}^3$ , which is roughly six times that of hydrogen. Vacuum  $I_{\text{sp}}$  for methalox is usually in the 370–380 s range, a nice bump over RP-1's 340–360 s. Methane's molecular structure is simple:  $\text{CH}_4$ . One carbon, four hydrogens. There are no long chains to crack, so coking is dramatically reduced. NASA Marshall ran methane-cooled thrust chambers for thousands of seconds of hot-fire time and saw essentially no measurable carbon buildup [4]. That's a huge deal for reusability.

Methane also plays nice with autogenous pressurization. In a traditional system, you need high-pressure helium bottles to keep the tanks pressurized. Helium is pricey, the composite tanks are heavy, and you have to refill them between flights. With autogenous pressurization, you just tap off a bit of gaseous propellant from the engine, run it through a heat exchanger, and use that to maintain tank pressure. Methane's vapor pressure at operating temperatures is high enough to make this work—though at 90 K the saturation pressure is only about 0.1 MPa, so you do need to heat it up quite a bit. The engine provides plenty of waste heat for that.

Hypergolic combos—MMH and NTO, for instance—ignite on contact. No ignition system needed. They're storable at room temperature for long periods. That's why they've been the standard for spacecraft thrusters, orbital maneuvering systems, and landers. Dragon's Dracos use them. Apollo LM used them. But they're toxic and corrosive. Ground crews need full SCAPE suits, and any spill requires extensive decontamination. For a reusable first stage that lands back at the pad, hypergolics would be a logistical nightmare. Nobody wants that.

Incidentally, there's work on "green" hypergolics that aim to cut the toxicity while keeping the storability. More on that later.

### 3.2. Solid propellants

Solid motors are mechanically simple. A typical composite grain is ammonium perchlorate as oxidizer (around 70%), HTPB rubber as binder and fuel (about 15%), and aluminium powder (about 15%) to boost combustion temperature and damp instabilities. No pumps, no valves, nothing moving. Thrust-to-weight can be over 100:1. That's why solids are great for strap-on boosters—they give a big kick early in flight and then get dropped. Shuttle's SRBs each cranked out about 12.5 MN at liftoff.

The exhaust plume, however, is less lovely. It contains hydrogen chloride gas, which turns into hydrochloric acid in moist air, plus fine alumina particles. There's been studies on the environmental side. Denison and others modelled the plume and concluded that chlorine radicals in the stratosphere could cause localized, short-term ozone depletion [5]. The white smoke trail you see from a solid booster is mostly that stuff.

Solids can't be throttled or shut down once lit. That limits where you can use them. Also, the grain can develop cracks over time—a failure mode that has to be managed carefully. The Shuttle SRB field joint issue was a different thing (O-rings), but it shows solids have their own peculiar failure modes.

### 3.3. Hybrid systems

Hybrids try to split the difference. Solid fuel grain, liquid oxidizer. Burn only happens when oxidizer is flowing. This is safer than a solid motor because the fuel and ox are separate and in different phases. A cracked case doesn't blow up. You can also stop and restart, which solids can't do.

The big technical hurdle for hybrids has always been regression rate—how fast the fuel surface burns away. Classic fuels like HTPB burn slow. You end up needing complicated multi-port grains just to get enough surface area for decent thrust. Paraffin-based fuels burn faster because a thin melt layer forms and gets unstable, shearing off into droplets that get swept into the flow. This can bump the regression rate by a factor of three or more compared to HTPB [6]. But if you add polyethylene wax to make the grain stronger, the regression rate drops again—by like 33% at 20% wax loading—because the wax calms down the instabilities that help strip the melt layer [7].

Hybrids also have a shifting O/F ratio during the burn. As the port gets bigger, the mass flux changes, and the mixture ratio drifts. You can compensate with active control, but that adds complexity. And one of the selling points of hybrids was supposed to be simplicity. So yeah. They've flown a few times—SpaceShipTwo's motor is one example—but they haven't taken over the world (Table 1).

Table 1. Performance characteristics and critical limitations of five major propellant types

Propellant Combination	Type	Typical Vacuum Isp (s)	Bulk Density (kg/m <sup>3</sup> )	Critical Limitation
LOX / LH2	Cryogenic	450 - 465	Around 71	Volumetric inefficiency
LOX / RP-1	Cryogenic	340 - 360	Around 1030	Coking / Reusability barrier
LOX / LCH4	Cryogenic	370 - 380	Around 420	Higher vapour pressure
NTO / MMH	Storable Liquid	310 - 340	Around 1200	Toxic hazard mitigation
APCP (Composite)	Solid	260 - 280	Around 1800	Environmental acid deposition

### 4. Booster versus upper stage: different animals

What works for the first two minutes of flight isn't necessarily what works for a six-hour coast and a restart in vacuum.

#### 4.1. The first stage

A first-stage booster has to lift the whole stack and punch through the thickest part of the atmosphere. To minimize gravity losses, you want high initial thrust-to-weight. To minimize drag, you want a compact cross-section. Both favor dense propellants.

A hydrogen first stage would be enormous. Big tanks, heavy structure, high drag. Delta IV Heavy uses hydrogen for all three cores, and it's an expensive vehicle for a reason—the hydrogen tankage drives a lot of that cost and dry mass.

RP-1 gives you density. But then you have the coking headache during reuse. Methane sits in a middle ground that looks better the more you look at it. Its density is okay—not amazing, but okay. For 100 tonnes of propellant, a methalox tank is around 240 m<sup>3</sup>, versus 97 m<sup>3</sup> for RP-1 and 1410 m<sup>3</sup> for hydrogen. And the clean burn means way less cleaning between flights.

I remember seeing a presentation a few years back where an engineer from one of the commercial launch companies mentioned that after a full-duration methalox engine test, you could

practically wipe the chamber with a rag and it'd be clean. With kerosene, you're scraping carbon for hours. That's not a published number, but it sticks in your head.

Look at what's coming out of the BE-4 and Raptor programs. BE-4 has run for thousands of seconds on the test stand. Post-test inspections reportedly show clean chambers with no major coke [8]. Whether the fifty-flight target is actually achievable remains to be proven—there's a difference between test-stand time and real flight cycles with all the thermal and structural loads—but the direction makes sense.

## 4.2. The upper stage

Once you're above the sensible atmosphere, the game changes. Tank volume doesn't matter much for drag anymore. The structural mass penalty of a big hydrogen tank is less important because the stage is already much lighter than the booster. Specific impulse becomes the dominant figure of merit.

Hydrogen still rules up here. Centaur III can send about 4 tonnes to Mars. A kerosene upper stage with the same starting mass would send a lot less because the lower  $I_{sp}$  forces a higher propellant mass fraction for the same  $\Delta V$ . The rocket equation compounds the advantage.

Methane has also been examined for upper-stage applications. While a methalox upper stage incurs a specific impulse penalty of approximately 20% relative to hydrogen, it offers a 5–10% improvement over kerosene and eliminates the operational overhead of deep-cryogenic hydrogen handling. For missions that do not require maximum geocentric escape energy, this trade-off may be favorable. However, published trade studies have not yet established a consensus on the optimal configuration.

## 4.3. A bit more on autogenous pressurization

I talked about this earlier, but it's worth coming back to. Getting rid of the helium system is a bigger deal than it might first seem. Helium is light, yeah, but the composite tanks to store it at high pressure add maybe 100–200 kg to a medium launcher. And you have to refill them between flights. And helium prices have been known to spike. Autogenous pressurization just uses the propellant itself. Methane works well for this. Kerosene would need crazy high temperatures to vaporize enough. Hydrogen has super low density even as gas. Methane is in that sweet spot where it's doable without heroic engineering.

## 5. Some current research directions

The field isn't standing still. A few areas are worth keeping an eye on.

Hydrazine is nasty stuff—carcinogenic, toxic, a headache to handle. But it works, and it's storable at room temp. That's why it's been the standard for spacecraft propulsion forever. Finding a "green" replacement that's safer but still performs has been a goal for ages.

Two candidates are fairly far along. LMP-103S, from Sweden, is based on ammonium dinitramide (ADN). It's flown on a number of smallsat missions. AF-M315E, based on hydroxylammonium nitrate (HAN), has been tested a bunch by the Air Force and NASA. These ionic liquids are denser than hydrazine and have similar or slightly better  $I_{sp}$ . They're way less toxic, though not exactly harmless.

The catch is they burn hotter. Hydrazine decomposition runs around 1300 K. These green props can hit 2000 K or more. That's rough on catalyst beds. The traditional Shell 405 catalyst degrades

fast at those temps. New catalyst formulations are being worked on, and some have shown decent lifetimes, but the technology isn't as mature as hydrazine yet [9].

The old-school way to develop new propellants was to mix stuff and light it on fire. Slow. Expensive. Now people are using machine learning to screen huge libraries of candidate molecules. Train a model on known compounds, let it guess properties of new ones.

Huang and others used genetic algorithms to search for additives that could bump solid propellant  $I_{sp}$  by a few seconds [10]. Cool idea. The bottleneck is experimental validation—you still have to actually make the molecule and fire it in a motor. And predicting long-term storage stability from computational models is really hard. Molecules do weird things over years.

If you're going to Mars and want to come back, you probably want to make your return propellant on the surface. The Sabatier reaction takes  $\text{CO}_2$  and hydrogen and makes methane and water. Mars air is mostly  $\text{CO}_2$ . Water ice is known to exist underground in lots of places. Split the water, get hydrogen, make methane. Boom, you have methalox on Mars. This is one reason methalox is popular for Mars ascent vehicle concepts [11].

3D printing is changing how combustion chambers are made. You can print copper alloy chambers with cooling channels that would be impossible to machine conventionally. NASA's RAMPT project has demonstrated large-scale printed chambers with integral cooling channels. This lets you tailor the thermal design to whatever propellant you're using, maybe squeezing out a bit more performance or extending life.

## 6. Conclusion

So where does all this leave us? For upper stages, liquid hydrogen is still the champion. The high  $I_{sp}$  is worth the tank volume penalty once you're out of the atmosphere.

For reusable first-stage boosters, the wind is blowing toward methane. It's denser than hydrogen, cleaner than kerosene, and works nicely with autogenous pressurization. It's not the best at any single thing—density,  $I_{sp}$ , or handling—but it's good enough across the board that the overall package is compelling.

Solids still have a place for cheap, high-thrust augmentation, especially in strap-ons. Hybrids remain a niche with potential but some tough problems.

There's still plenty we don't know. Long-term fatigue of methane engine parts over hundreds of cycles isn't well documented in the open literature. Scaling up green ionic liquids from lab to tons is a whole other engineering challenge. And ML-based propellant discovery needs a lot more experimental validation before it starts delivering flight-ready formulations.

But the overall trajectory seems clear. The industry is moving from a single-minded focus on specific impulse toward a broader, lifecycle-oriented view of what makes a propellant "efficient." And that's probably a good thing.

## References

- [1] Sutton, G. P., and O. Biblarz. Rocket Propulsion Elements, 9th ed. John Wiley & Sons, 2016.
- [2] Aerojet Rocketdyne. "RL10 Propulsion System Data Sheet." 2018. <https://www.rocket.com/space/liquid-engines/rl10>
- [3] Naraghi, N. H., S. Dunn, and D. Coats. "Effectiveness of dual cooling to lower maximum wall temperature of regeneratively cooled engines." AIAA 2006-4367, 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference, Sacramento, CA, July 2006.
- [4] NASA Marshall Space Flight Center. "Methane engine combustion stability and heat transfer investigations." NASA/TP—2015-218559, 2015.

- [5] Denison, M. R., et al. "Solid rocket exhaust in the stratosphere: Plume diffusion and chemical reactions." AIAA 92-3388, Nashville, TN, July 1992.
- [6] Karabeyoglu, M. A., D. Altman, and B. J. Cantwell. "Combustion of liquefying hybrid propellants: Part 1, general theory." *Journal of Propulsion and Power*, Vol. 21, No. 4, pp. 610–620, 2005.
- [7] Tian, H., Y. Li, and G. Cai. "Experimental investigation of paraffin-based hybrid rocket fuels with polyethylene additives." *Aerospace Science and Technology*, Vol. 68, pp. 564–572, 2017. doi: 10.1016/j.ast.2017.06.012
- [8] Dodd, F. E. "BE-4 engine development update." IAC-20-C4.1.2, 71st International Astronautical Congress, October 2020.
- [9] Brechbill, S. C., and T. W. Hawkins. "Development of an AF-M315E fueled emergency power unit." AIAA 2012-4291, 48th AIAA Joint Propulsion Conference, Atlanta, GA, 2012.
- [10] Huang, Z., V. Yang, and R. A. Yetter. "Data-driven modeling of rocket propellant thermochemistry using machine learning." *Aerospace Science and Technology*, Vol. 115, 106828, 2021.
- [11] Muscatello, A. C., and E. Santiago-Maldonado. "Mars in situ resource utilization technology evaluation." AIAA 2012-360, 50th AIAA Aerospace Sciences Meeting, Nashville, TN, January 2012.