Some Thoughts On End-Burning Amateur Rockets

Steve Peterson

Version 1.0 - 21 Dec 2023
Version 1.1 – 2 April 2024: expanded section 1.3 Dual Thrust

Introduction

In the field of rocketry there is a maxim that "in the absence of gravity, fly as slowly as possible while in the absence of drag, fly as fast as possible."

Given that HPR rockets fly in the presence of drag, it is regrettable that nearly all commercial and amateur HPR motors burn radially (BATES, cruciform, finocyl, etc.) making them high-thrust with short burn times. This tends to create flights characterized by high velocity and therefore high drag, thus directly contravening the advice given above. Such flights require large amounts of impulse in order to achieve high altitudes, but simply adding impulse in order to gain altitude runs into difficulties caused by quickly diminishing returns. And when impulse is restricted, as it is for altitude records, then the only option is to increase thrust by shortening burn times which soon runs into the same difficulty.

Consequently, a number of years ago the author concluded that it was time to seriously examine an approach which actually attempts to follow the maxim and this led directly to in-depth study of end-burning grains.

End-burning grains have probably been considered by practically everyone who has looked into the problem of reducing the drag created by high-thrust designs. What often seems to happen, however, is that, for the propellants available to amateurs, the low thrust produced by end-burning grains causes experimenters to immediately conclude that a faster burning propellant is required and the investigation ends.

Given the difficulty of developing a faster burning propellant, the author decided to investigate just how much thrust was actually required to overcome drag at low velocities, whether that thrust could in fact be produced by an end-burning grain made with existing sugar propellants, and how grain geometry might be used if it were necessary to increase the thrust for critical segments of a rocket’s flight.

This investigation, conducted off and on over more than a decade, has consumed probably hundreds of hours of simulation and online research, as well as countless hours at the workbench attempting to fabricate insulating liners and other components for long-burning, low-thrust motors. Numerous rat-holes have been thoroughly explored. This work is nowhere near complete, but it is perhaps appropriate to pause the work and take stock of what has been learned to date. That is the purpose of this writeup.

The first few pages attempt to show, in an abstract way, just what is to be gained by low-thrust, long-burning approaches to propulsion, without consideration of any practical issues (like casing insulation) that might stand in the way. This section depends entirely on simulation (with whatever problems may be inherent in such an approach*) to illustrate a number of principles that may not be apparent to enthusiasts who have not ventured beyond, say, BATES-type motors. In particular, it is shown that for many flights which might be considered by amateur motor-makers, burn rate is not the issue many consider it to be.

The next section then looks at actual rockets, including two end-burners, to put concrete numbers to some of the principles of the first section.

This is then followed by a section describing the design of end-burning grains which will achieve thrust profiles that produce much more optimal use of propellant than is possible with radially-burning grains.

Finally, some consideration is given to the practical issues that must be overcome if long-burning motors are to be successful, including L/D ratios, casing insulation, and gravity turns, along with indications of progress the author has made to address them.

---

*The simulations were run with the author’s own simulator, developed/improved over the past two decades. As part of preparing this writeup, a number of the rockets were also simulated using RASAero II which the author considers to be the “gold standard” of simulators, especially for supersonic flights. The difference in altitude between the two simulators varied, but was never more than 5%, the author’s simulator giving more conservative numbers in every case. Given that the goal of this writeup is to present principles rather than exact numbers, this correlation is more than adequate.

Note: it is not valid to compare efficiency numbers given for simulated rockets with those of actual rockets described in section 2: the actual rockets used AP-based propellants and the simulations used sugar-based propellants which have much lower specific impulse (meaning that they weigh more for a given total impulse).

1.0 Why End-Burners?

1.1 Thrust, Drag, Velocity, and Altitude

Let us start by looking at the effect of drag on the velocity profile for a rocket flown with different levels of thrust.

Below is a graph of velocities for five simulated flights of the same 2” diameter rocket but with thrusts of 640N, 640N, 256N, 85N, 26N, and 14N. In all cases the total impulse was 1280N-s, so the burn times were lengthened to achieve the lower thrusts. Mass was optimized to produce the greatest altitude for a given thrust. The red curve is with Cd = 0 (i.e., no drag), the rest are with a dynamically calculated Cd in order to illustrate the effects of drag on velocity when various thrust levels are used. The smaller curves have been offset along the x-axis for greater clarity.

It is not news, of course, that drag dramatically affects a rocket’s flight (red curve vs the other curves). In the graph, the peak (burnout) velocity is obviously greatly diminished from flight 1 to flight 2 (same 640N thrust with and without drag), but more importantly, the velocity at every point of the curve is decreased leading to a very reduced coasting profile (notice how the straight red curve becomes concave when drag is introduced). That is “bad news” for a flight that depends on coasting for much of its altitude gain.

However, the real point of this chart is how the *shapes* of the curves change as thrust levels are reduced. Notice that the area under the curve *increases* each time the thrust is reduced. The total area under a velocity curve, is, of course, the altitude, so it is clear from visual inspection that, other things equal (specifically total impulse), as thrust decreases, altitude increases (altitudes for flights 2-5 were 12699’, 13508’, 18005’, and 31696’--see the table below the chart). Notice, however, that flight 6 only managed 25922’ showing that the idea has limits.

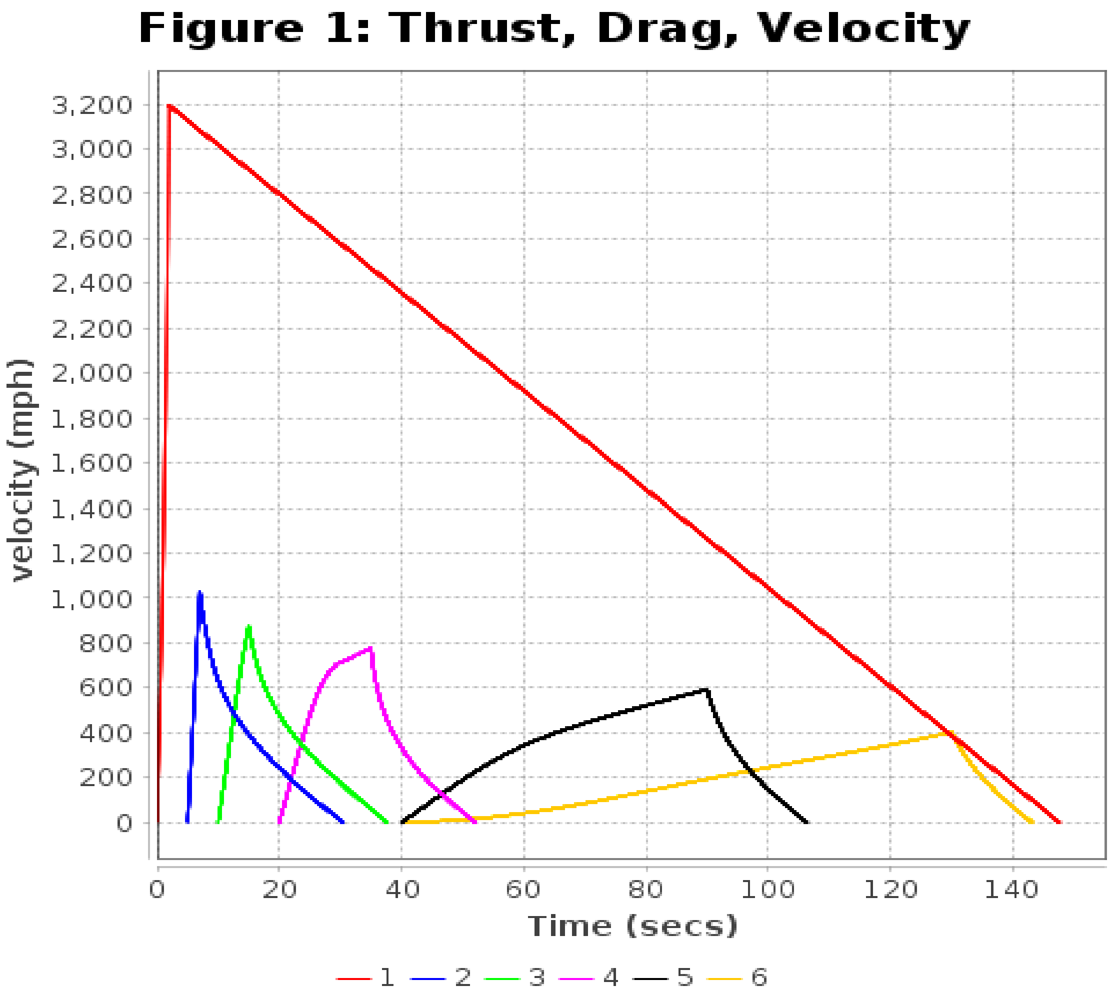

Rocket #	Mass kg	Thrust N	Altitude ft	Efficiency ft/N-s
1	.5	640	----	(no drag, just for comparison)
2	.5	640	12699’	9.9
3	2.6	256	13508’	10.6
4	.7	85	18005’	19.1
5	.5	26	31696’	24.8
6	.5	14	25922’	20.3

Notice that decreasing the thrust from 640N to 256N (flights 2 and 3) did little to affect the altitude. If the benefits of low thrust are to be realized, the thrust must be *really* low, at least by HPR standards.

One can introduce the concept of “efficiency”, which, like efficiency for cars, is a measure of how far the vehicle goes per unit of “propellant”. In the U.S., a car’s efficiency is measured in miles/gallon. Similarly, a rocket’s efficiency can be measured in feet (or meters) of altitude per N-sec of total impulse.

In this case, notice that the efficiency (last column) goes up as thrust is reduced, although, as noted, the idea can be taken too far (rocket 6).

It is also worth noting that the optimal masses for the low-thrust rockets were significantly lower than for the higher-thrust ones. This is due to coasting being far less important than the boost phase in low-thrust flight regimes, and the lower the thrust, the more this is true. (In fact, the 0.5kg mass of rocket 5 is not optimal, but that’s the lowest it could be and still be realistic—even then, it will take some skill and attention to detail to build something that light.) Finally, notice how little thrust is being used to achieve the greatest altitude: 26N of rocket 5 vs the 640N of rocket 2; this contrast will be seen again later when actual rockets are compared.

Given the predilection of many HPR enthusiasts to simply throw more impulse into a rocket in order to achieve greater altitude, the following demonstrates the result (next page):

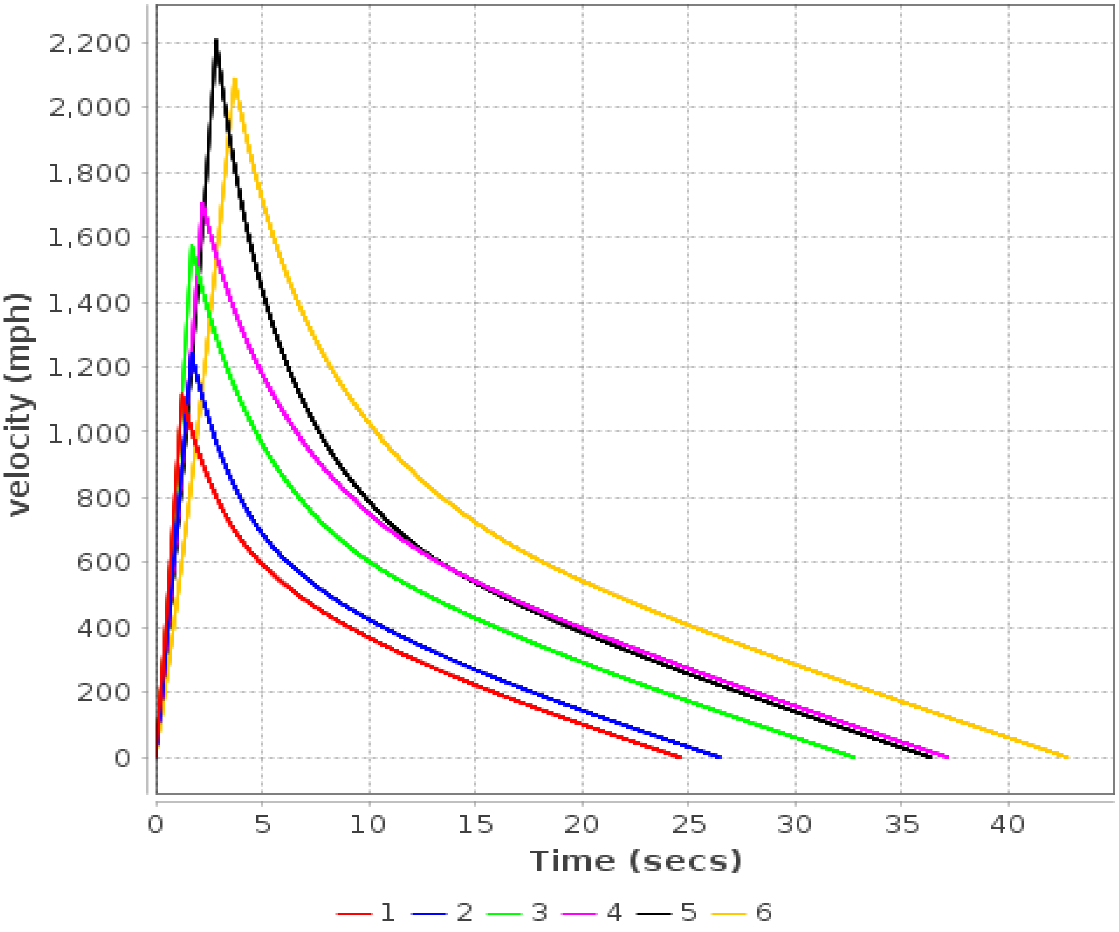

Figure 2: High Thrust, Increasing Impulse

Rocket #	Dry Mass kg	Imp. N-s	Thrust N	Altitude ft	Efficiency ft/N-s
0	1.9	1280	1024	12476’	9.7
1	3.2	2560	1505	14572’	5.7
2	4.9	5120	3011	22368’	4.4
3	8.7	10240	4654	28193’	2.8
4	10.8	20480	7185	30126’	1.5
5	23.8	40960	11070	38075’	0.9

Here, rockets have been given impulses from 1280 (J-class) to 40960 N-s (O-class). Burn time was kept short (2-5 seconds) and the dry masses were adjusted to give roughly maximum altitude. However, as impulse was increased, the diameter of the rocket had to be increased to keep the grain’s L/D ratio at least somewhat realistic. This, of course, increased the drag which impacted the altitude, but that’s exactly what happens as propellant quantities are increased—eventually diameter has to be increased and this necessarily hurts altitude.

The last column, efficiency, tells the story: trying to gain altitude by adding propellant to a fast-burning, high thrust rocket which depends on a long coast to achieve its altitude, is just not a very efficient way to increase the altitude. Drag dramatically reduces the burnout velocity and then quickly reduces the velocity further during the early part of the coast phase. The greater the thrust, the greater the impact of drag, and this shows up both in the curves and in the efficiency numbers. This is not to say that additional altitude is not gained as impulse is added—from the table it can be clearly seen that it is. Rather, it is that the returns are diminishing so that it takes more and more impulse to get the next increment of altitude: in this example, each flight roughly doubles the impulse of the previous one, while altitude increases at a much slower rate.

One might also note the relatively high thrust that is required to achieve a high altitude when compared with the low-thrust rockets in Fig. 1 and that the impulse of 1280 N-s for the rockets in Fig. 1 is about 1/16 that required by the high thrust regime (Rocket 4 in Fig. 2) to achieve approximately the same altitude (31696’ vs 30126’).

Having seen the benefits of flying slower by using less thrust (but for a longer time), we next look at the result of adding impulse to a low-thrust rocket, similar to what was shown in Fig. 2 for a high-thrust series of rockets:

Figure 3: Low Thrust, Increasing Impulse

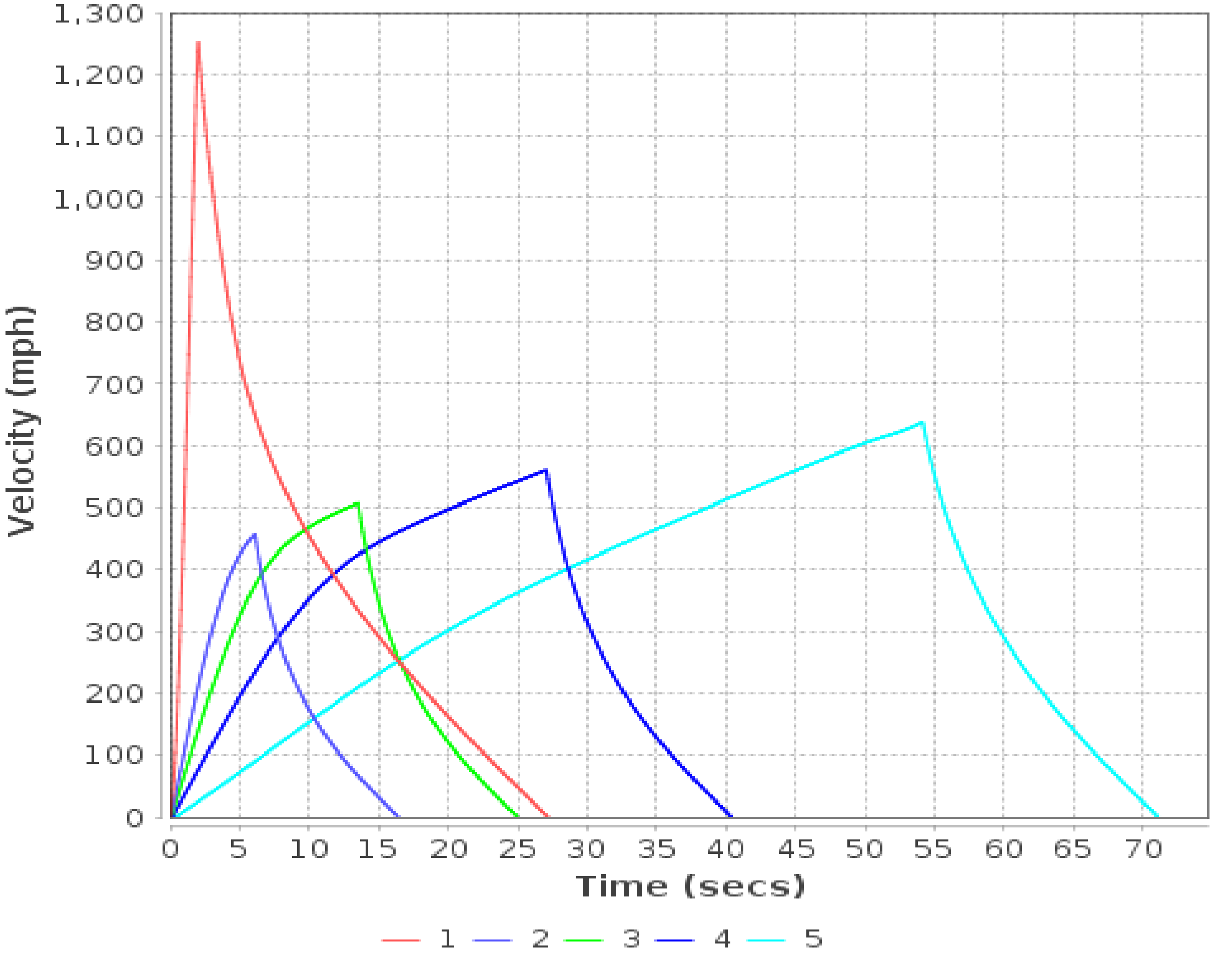

Rocket #	Imp. N-s	Thrust N	Altitude ft	Efficiency ft/N-s	
1	2523	1498	15201’	6.0	(For comparison)
2	209	35	4837’	23.1	
3	470	36	9695’	20.6	
4	956	37	18485’	19.3	
5	1951	39	34380’	17.6	

The main point illustrated here is that propellant can be added to an end-burner while maintaining very high efficiency as compared with a high-thrust motor (compare Fig. 2). Also, the low thrust (35-39N) should be noted, especially in comparison with the altitudes that can be reached. We can use the economic concept of “marginal cost / marginal return” as a figure of merit:

(impulse 2 – impulse 1) /(altitude 2 – altitude 1) where a lower result is better.

Using values from the table for Fig 2. and the table above, we can see that:

High thrust: (20480 – 10240) / (30126 – 28193) = 10240/1933 = 5.3 N-s / ft
Low thrust: (1951 – 956) / (34380 – 18485) = 995 / 15895 = 0.06 N-s / ft

That is, the cost for one more foot of altitude is 5.3 N-s in the high thrust case, while it is only 0.06 N-s in the low thrust case. This is quite a difference (more than 88x).

Another way of looking at it is this: the dollar cost (roughly) for a 2500N-s HPR motor is about $300 whereas the cost for one having about 950N-s is only a little over $100. If the 950N-s motor were an end-burner, it would achieve the same altitude as the 2500N-s motor, meaning that about $200 of the expensive motor was spent just to push the air around rather than achieving altitude.

1.2 Effects of Mass on Low-thrust flights

Having seen the effects of drag on altitude, let us briefly look at the other great determiner of altitude: mass, in particular, the “dry” (or structural) mass of the rocket.

A rocket’s flight can be viewed as a complex interplay of drag and gravitational forces—depending on the rocket, drag forces may dominate throughout the flight or gravitational forces may be dominant, or, if the rocket is somewhat optimized, it may be that one force dominates at the beginning of the flight, while later on, the other dominates. This means that simulation must be used to find the best tradeoff between a rocket’s dry mass and its velocity for a given objective.

Most readers will know that for rockets which depend on long coast times to achieve altitude it is sometimes necessary to increase the rocket’s mass to achieve maximum altitude.

However, the game changes when using a long-burn/low thrust approach to achieving altitude.

In this case, the thrusting portion of the flight achieves most of the altitude and the coasting portion contributes very little. It’s a continuum, of course: one could use a “medium” thrust with a “medium” long burn time. But whatever the design, about the best that can be said is that for a low-thrust/long burn scenario, it is almost always the case that a lighter rocket will fly higher than a heavier one, other things equal. *How much* higher or lower for a given change in mass is not intuitively obvious, however. Two examples will illustrate the problem:

Example 1: a G-class motor, in a 2” diameter rocket, having a dry mass of 341 grams will reach a simulated altitude of 4800 feet. Adding 34 grams (about 10%) causes a reduction of 105 feet, or 2%. However, if the dry mass were instead 640 grams and 64 grams were added (about 10%), the altitude drops from 3352 feet to 3001 feet, a difference of 351’ or 10.5%.

Example 2: an L-class motor, in a 3” diameter rocket, having a dry mass of 887 grams reaches a simulated altitude of about 95,589 feet. Adding 88 grams (about 10%) causes the altitude to drop to 88,635 feet, a difference of 6954 feet or about 7.3%. If the dry mass were 1887 grams the altitude would be about 43,506 feet; adding 10% mass, the altitude drops to 34,824 feet amounting to about a 20% reduction in altitude. In contrast with Example 1, increasing the mass of the lighter rocket by 10% decreased the altitude by 7.3% (not the 2% in Example 1), and increasing the mass of the heavier rocket by 10% decreased the altitude by 20% (vs. 10.5% in Example 1).

Sometimes a slight increase in mass will make almost no difference, sometimes it will make a much larger one. In other words, simulate, simulate, simulate. See section 4.2 for additional examples.

The foregoing comments imply that low-thrust/long-burn is not the solution to every mission one might contemplate—sending a bowling ball to orbit, as an extreme example, would not be a good match. But when the dry mass allows for a low-thrust regime, and the desired altitude is high enough to be worth the effort, then a low-thrust/long burn approach deserves serious consideration.

1.3 Dual Thrust

In examining Figure 3 one notices that most of the velocity profiles are very triangular and that rocket 5’s velocity is quite low at the beginning, possibly indicating that its speed off the pad might be too low to be stable. These observations suggest that a relatively high initial boost, followed by a lower, sustained thrust would be advantageous, both for launch speed and altitude. This is in fact true.

In the chart below, the red curve is the profile of a low thrust rocket with no initial boost showing the usual slow ramp-up of velocity. The blue curve is for the same rocket but with a short initial thrust of about 11x the sustained thrust. Although the total impulse is roughly the same for each rocket, the resulting altitude of the boosted one is almost twice that of the unboosted version (see the table below the chart).

Figure 4: Dual Thrust

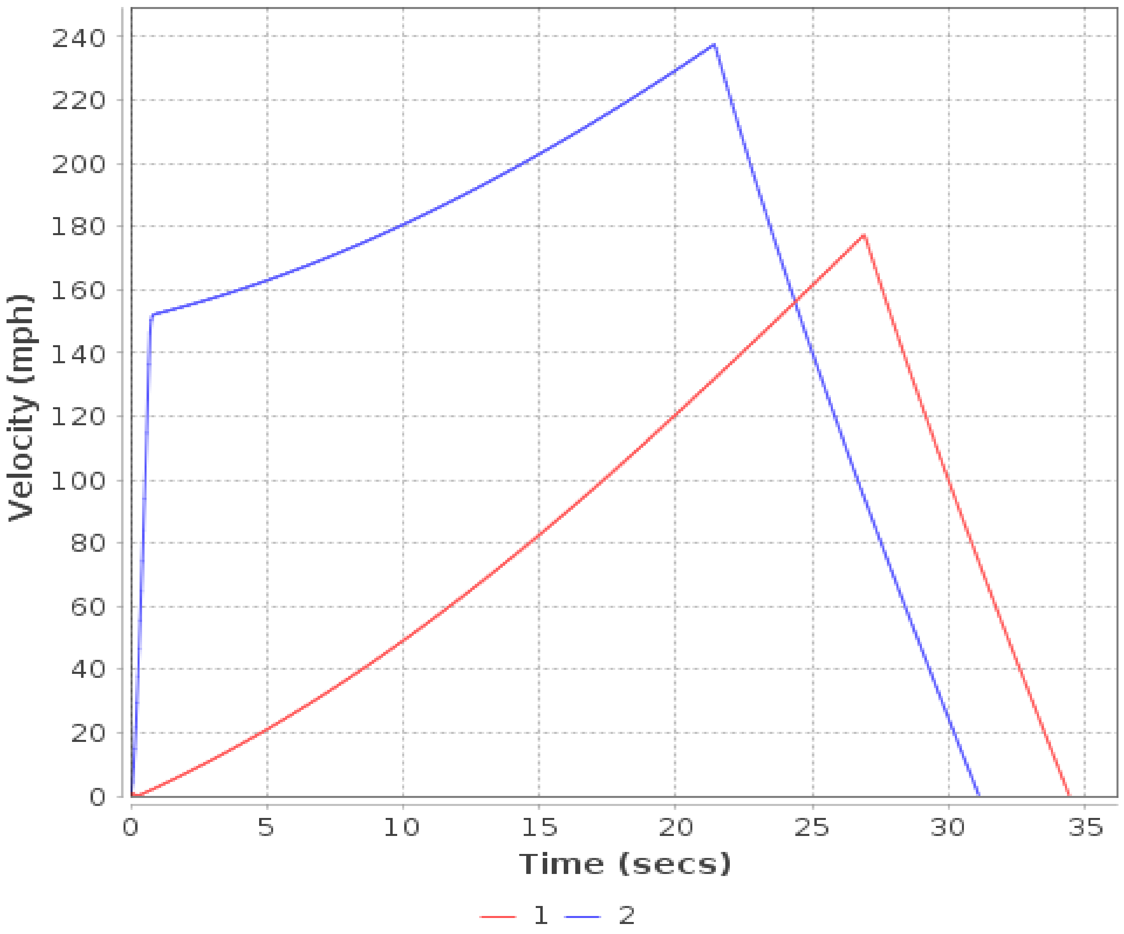

Rocket #	Impulse N-s	Thrust N	Altitude ft
1	612	24	3996’
2	621	253/24	7377’

One might wonder just how it could possibly be the case that one could nearly double the altitude with essentially the same impulse.

The answer is that a bit of the propellant is taken from the end-burning portion and reshaped to provide high thrust for a very short burn time. One might think "that small amount of propellant can't possibly result in enough of a velocity increase to do anything useful."

Behind the scenes of Figure 4, the grain geometry changes from an 8" long end-burning portion for the first curve, to a 6.5" end-burning portion plus a 2" long, 1" core portion for the second. The cored portion has a web thickness of only about 0.33 inches resulting in a burn time of a little more than 0.5 seconds--not much impulse at all and the two grains weigh about the same. The reader will notice in Figure 4 that the burn time has shortened from about 27 seconds to about 21 seconds due to the shortened e-b portion.

Getting something useful from that smidgeon of the grain results from a combination of factors. The goal is not to fly to any great altitude with that cored portion, nor to fly very fast, just to fly faster than without the cored portion. In other words just get the rocket moving at a decent clip and let the long burn time of the e-b portion do the rest. This is helped by the low dry mass of a well-designed end-burning rocket, meaning that "a = F/m" can be greater than it might be for a typical HPR rocket and, compared with the initial velocity of a pure end-burner, we don't have to fly that fast to get a fairly large improvement. Looking at Figure 4, the initial velocity of the squared-up curve is only around 150 mph; if this were a BATES motor designed to reach 7000', that velocity would have to be much greater, which would take a lot more propellant. Notice also that the velocity of the pure end-burner is only 20 mph after 5 seconds (because it's a contrived example), so it isn't all that hard to fly faster. (The more realistic example is found later in section 3).

So while the initial velocity is large compared with the pure end-burner, it is small compared with a BATES motor. As a result not much impulse is required. Said differently, the purpose of the cored portion is not to gain much altitude, but merely to impart initial velocity,

The effect on the velocity curve is that it becomes more rectangular which gives much more area under the curve during the early portion of the burn. In the graph, the right, end portion of the curve is shortened due to the shortened e-b portion of the grain, so a bit of area (i.e., altitude) under that portion of the curve is lost there, but the large increase at the beginning of the burn more than offsets that loss.

Looking ahead to section 3.0, the design process consists in part of juggling these two parts of the grain--trading off propellant devoted to the core (and the resulting velocity) vs. the portion devoted to maintaining a velocity that maximizes altitude (assuming that's the goal).

1.4 Decreasing Thrust to Keep Velocity Low

Finally, let us look at a case where a burn rate *lower* than nominal might be useful. To those accustomed to the high thrusts of BATES grains, and who believe that end-burners produce too little thrust to be useful, it may come as a bit of a shock to be told that sometimes even an end-burner’s “small” amount of thrust can be too much.

As the previous chart shows, and, in fact all the previous profiles of long-burn motors show, there is a tendency for the velocity to rise as the flight proceeds. This is to be expected since the rocket’s mass is decreasing as is the atmospheric density, leading to greater acceleration for a given thrust, which in turn means a greater velocity.

For rockets having sufficient thrust for a long enough period, this “ski jump” profile can turn out to be quite wasteful because the higher velocity is causing the drag to increase. Consider the following chart which shows a J-class (1233 N-s) rocket, 2” in diameter. The red curve results from a thrust profile of 242 N initial thrust followed by a sustained thrust of 22N. Altitude gained is 28910 feet.

However, a bit more altitude can be squeezed out by changing the sustained thrust to a *decreasing* series of thrusts of 22N, 16N, and 13N, which, in this simulation were created by changing the burn rate from nominal to 72% of nominal and 60% of nominal (blue curve). Because the thrust was decreased by decreasing the burn rate, the burn time was increased (for the same amount of propellant). This leads to a *series* of ski-jump shaped curves, but ones which have a smaller amplitude than the original; if the right thrusts are chosen, this decreasing series of thrusts results in greater altitude than would be achieved by a simple dual-thrust regime.

The advantage of this approach is not overwhelmingly large for the example in this chart. For larger rockets, however, the advantage increases significantly and it becomes worthwhile to use a number of thrust levels. Eventually, of course, an altitude will be reached where the atmospheric density is so low that it becomes advantageous to increase the final thrust to the maximum because drag at that point will be practically zero.

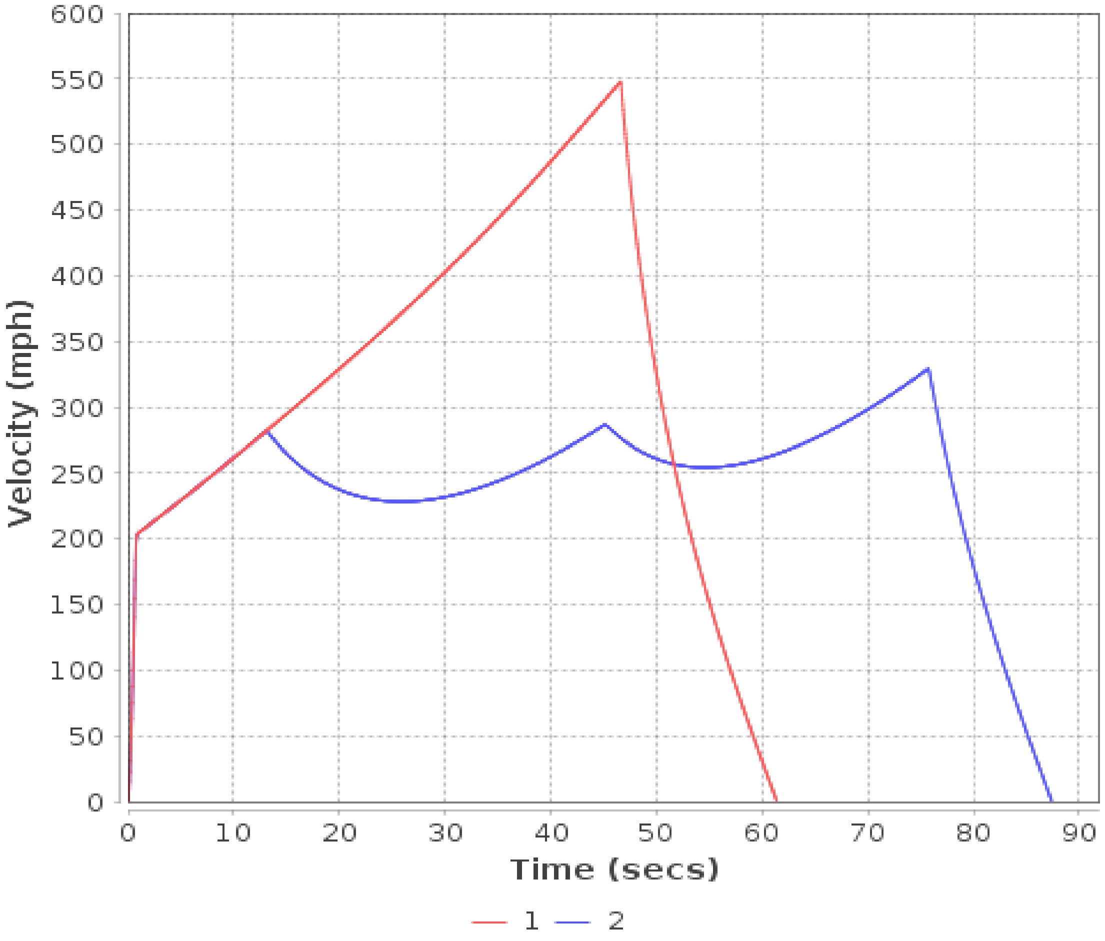

Figure 5: Decreasing Thrust

Rocket #	Imp. N-s	Thrust N	Altitude ft
1	1233	247/22	28910’
2	1242	247/22/16/13	30979’

1.5 Low Thrust in Smaller Rockets

At this point, after examining the results of rockets of J-class and larger motors, one might wonder whether there is a lower limit to rockets that would benefit from a low-thrust, long burn regime. The answer is that, although a rigorous study has not been carried out, simulation seems to indicate that a G-class motor will certainly benefit, but an F-class one may not:

Rocket	Imp. N-s	Thrust N	Altitude ft
G Bates	159	208	3666’
G End-burn	158	23	5082’
F Bates	79	119	3021’
F End-burn	80	10	2903’

In producing this table some optimization was done for each rocket, but it certainly wasn’t exhaustive.

1.6 Summary

This concludes our detailed examination of low-thrust regimes in the abstract. We have seen that:

1) Much greater altitudes can be reached with a given impulse by using less thrust but over a longer period of time as compared with high-thrust regimes.

2) Impulse can be added to a low-thrust rocket with less degradation of efficiency than is the case with high-thrust rockets.

3) Relatively high altitudes can be reached with very little thrust as long as the rockets are lightweight and aerodynamic.

4) A low-thrust approach can be beneficial for rockets with motors as small as G-class.

5) Because the simulations were done with a sugar propellant, it appears that sugar propellants are more than adequate for use in end-burning motors as long as the rockets are aerodynamic and lightweight.

With that, it is now time to leave the world of simulation behind and consider the flights of some actual rockets to see what they might tell us about thrust and efficiency.

2.0 Actual Rockets

At this point, it will be helpful to define an additional metric: *mass fraction*, which is defined as propellant mass / total mass. It is a measure of how efficiently the rocket’s non-propellant (structural) mass is being used. For example, a rocket with a mass fraction of 70% is using far less structural mass in proportion to the propellant mass than is a rocket with the same total mass but with only a 10% mass fraction. This can be a helpful guide when designing/building a rocket for a low-thrust regime since low mass is essential for maximizing altitude per unit of impulse.

The metrics of mass fraction and efficiency, together with what has been presented concerning thrust, can now be applied to some real rockets, information about which was gleaned from various online sources:

Key characteristics of some actual rockets

Rocket ID	GLOW KG	Grain	P-Mass G	T-Imp N-S	BO-Mass KG	MF %	BT Sec	Mx Th N	Mx V MPH	Alt Feet	Eff F/N-S
A	.29	EB	116	211	.174	40	15.0	45/14		15257	72
B	1.105	B	313	636	.792	28	2.9	379		17541	28
C	1.518	B	616	1267	.902	41	1.2	1200	2118	24000	19
D	5.428	B	624	1260	4.8	11	2.8	590		5050	4
Arcas	33.4	EB	18636	42035	14.7	56	29.0	1460	2795	268000	6.4
F	145.5	FCYL	69500	143000	76.0	48	8.0	18000	2181	123000	0.9

GLOW: Gross Lift Off Weight (KG)
P-Mass: Propellant Mass (g)
T-Imp: Total Impulse (N-S)
BO-Mass: Burnout Mass (KG)
MF: Mass Fraction (%)
BT: Burn Time (sec)
Mx Th: Max Thrust (N)
Mx V: Max Velocity (MPH)
Eff: Efficiency (Feet/N-S)

EB: End Burning
B: Bates
FCYL: Finocyl

The first three rockets listed are highly optimized and meant for setting HPR altitude records. All use commercial motors, so were limited to what was available off-the-shelf.

Rocket A is a very light end-burner which achieved an altitude much greater than most amateur motor-makers are likely to reach. The efficiency (72 ft/N-s) is very high, mostly because of the long burn time (15 seconds). The motor had an initial thrust of 45N; the sustained thrust of 14N is astonishingly low when compared with the other rockets in the table, Rocket B in particular. It also shows just how light an HPR rocket can be thereby demonstrating that a high mass fraction can be achieved in a small rocket.

Rocket B used a BATES grain motor with a short burn time of 2.9 seconds. It’s efficiency suffers because of this and it can be seen that its thrust is very high compared with the low-thrust rockets simulated in section 1.1. By HPR standards it is a very light rocket with a decent mass fraction of 28%.

Rocket C is an attempt to reach a high altitude by using a very high thrust (1200N) for a very short period of time (1.2 seconds) and coasting to the desired altitude. It’s efficiency isn’t as bad as it might be primarily because the rocket is very light and it has a good mass fraction of 41%, but it used about 6x the impulse to go about 57% higher than Rocket A.

Rocket D is a typical HPR rocket meant for sport flying. It is fairly heavy (compared with Rockets A-C) with a mass fraction of 11% and its efficiency is only 4ft/N-s. Despite its mass (which would dictate higher thrust than with rockets A-C), its thrust is still quite high. Its total impulse is nearly identical to Rocket C’s, is about 2x Rocket B’s, and about 6x Rocket A’s, yet its altitude is far lower. Thus, its greater mass and diameter (3” vs Rocket A’s 1.2” and Rockets B and C’s 1.6”) show that diameter and mass penalize performance substantially.

Rocket “Arcas” is a sounding rocket produced for the U.S. military from 1959-1963. It’s an end-burner with a steel casing designed to carry payloads of 8.5-15lbs; the heavy casing and payload cause its efficiency to suffer. However, its altitude of more than 250,000’ shows just what can be achieved with an end-burner, even when burdened with such a heavy casing and payload. Because of the end-burning design, its diameter is only 4.5” so its drag losses were still fairly low despite its max velocity of more than Mach 3. Also contributing to its low drag loss is its “low” thrust (with a long burn time of 29 seconds). This caused its velocity to rise fairly slowly so that its top speed was reached at a fairly high altitude (where the air density is quite low). It achieved a 56% mass fraction, which would have been even higher without the payload. This rocket is a testament to what can be achieved with a single-stage end-burner.

Rocket F is an amateur rocket designed to reach slightly more than 100,000’ which in fact it did. Its design contrasts starkly with the Arcas rocket: its diameter was 8” (vs 4.5” for Arcas) leading to very high drag losses, burn time is only 8 seconds vs. the 29 seconds of Arcas, burnout mass is 4x, propellant mass is more than 3.5x, and thrust is 12x, while efficiency is only 0.9ft/N-s. Despite its success at meeting its goal, this is another example of a short-burn/long coast strategy being far from optimal.

In Section 1.2 dealing with dual thrust it was noted that end-burners could have a problem getting off the launch pad quickly enough. A rule-of-thumb in the HPR community is that a rocket's thrust needs to be about 6x its weight to launch successfully most of the time in relatively calm conditions. The two end-burners of this section illustrate two different solutions to this problem.

Arcas had a constant and relatively low thrust/weight ratio (1460N / 33.4 * 9.8 = 1460 / 327 = 4.5x) and used a piston launcher to increase the initial boost. Judging from comments in the development report, an even greater boost would have been desirable, but the method worked well enough that over 1400 flights were successful.

Rocket A used a high initial thrust (45 N) compared with a fairly low takeoff weight (not mass) of 9.8m/s * .39Kg = 3.8Kgf; almost 12x the takeoff weight. This is sometimes called a “2 phase” thrust profile and is the method that will be discussed shortly in the section on end-burning grain design.

The purpose of this section has been to use some actual rockets, mostly from the HPR community, to illustrate some of the principles developed in Section 1.

The most important of these is that in real life, just as it was in simulation, low-thrust, long-burn motors used in lightweight, aerodynamic rockets outperform high-thrust rockets having considerably more total impulse. Rocket B had three times the impulse of Rocket A, but flew only about 15% higher. Arcas had one-third the impulse of Rocket F, but flew more than twice as high even while carrying an 8-15 lb. payload.

Rocket D (when compared with Rockets A-C) shows the penalty caused by drag and mass, the clear implication being that if one is contemplating using an end-burner to achieve significant altitude, attention paid to minimizing mass and drag will pay large dividends.

Rockets A-C show that HPR rockets can be very light and have fairly high mass fractions thereby demonstrating that it just isn’t the case that HPR rockets are necessarily too heavy for end-burning motors.

3.0 End-burning Grain Design

If one examines the formulas for chamber pressure and thrust in a rocket motor, one sees that (simplifying slightly) for a given throat diameter, chamber pressure is the product of burn area and burn rate while thrust is a function of chamber pressure. As mentioned earlier, it seems that many enthusiasts look at the low thrust produced by end-burning grains and conclude that a faster burning propellant is required to achieve the thrust they assume is required—that is, they attack the burn *rate* part of the equation. Since no fast-burning propellant has yet made an appearance in the amateur community, it seems reasonable to attack the other half of the equation: burn *area*.

We saw, in the first section, “Why End-burners?”, that a high initial thrust was beneficial, and, in some cases, essential to getting a rocket off the pad. It was also beneficial to “square up” the velocity curve which resulted in a large increase in altitude. Let us therefore begin by addressing how to provide this high initial thrust by altering the grain’s burn area.

As it happens, the method for producing this initial thrust has been around probably for centuries and has been, and is still used in end-burning model rocket motors: one simply forms a small, or perhaps not so small, “core” at the aft-end of the grain which serves to create an initial burning surface that is larger (or perhaps much larger) than the end of the grain would normally provide. This larger burning area increases the chamber pressure and therefore the thrust. The core quickly consumes itself (it does not propagate up the grain) and the thrust then drops to whatever the end-burning portion of the grain can provide.

In model rocket motors this core is quite small—perhaps 1/8” in diameter and maybe 1/4” deep--but there is no reason it can’t be made larger as long as the strength of the casing is not exceeded by the increased chamber pressure resulting from the increased burn area.

To see how a core can be used (in combination with grain diameter) to achieve a variety of altitudes, let us start with a lightweight 2” diameter rocket with good aerodynamic characteristics. Its motor will contain a 1.5” diameter end-burning grain and its nozzle will have a 0.160” diameter throat. The motor would produce about 18N of sustained thrust.

If a 3” long grain were provided with a 1” long, 1” diameter core, the initial thrust would be about 55N and when the core was on the verge of burning out, the thrust would be almost 80N. Of course, this thrust doesn’t last long—less than a second, in this case—but it is sufficient to get a lightweight rocket off the pad and flying almost 200 miles/hour. At that point, 18N is more than enough thrust to not only maintain that speed, but to increase its velocity to over 300 mph, resulting in an altitude of about 7000’.

If that is not sufficiently high, one can lengthen the grain to, say, 6” and almost double the altitude. Still using the 1”x1” core, a 10” long grain will increase the altitude to about 18,000’. Even a 16” long grain is usable and will get to about 25000’, although the speed off the pad would only be about 30 mph. At this point one could simply lengthen the core to 2”, which would increase the initial chamber pressure from about 375 psi to about 1000 psi creating a liftoff speed of about 45 mph and resulting in an altitude of about 27000’. Here’s a summary:

End-burn Length	Core Length	Altitude feet
3.0”	1”	6839
6.0”	1”	11960
10.0”	1”	18147
16.0”	1”	25247
16.0”	2”	27322

Of course, one cannot continue in this fashion forever, because eventually the rocket becomes so heavy with additional propellant that the thrust produced by the end-burning portion of the grain won’t overcome the gravitational force (plus, of course the relatively small drag force). This causes the high initial velocity created by the core portion of the grain to drop until enough of the grain has burned to reduce these forces. The profile is shown in Fig 6:

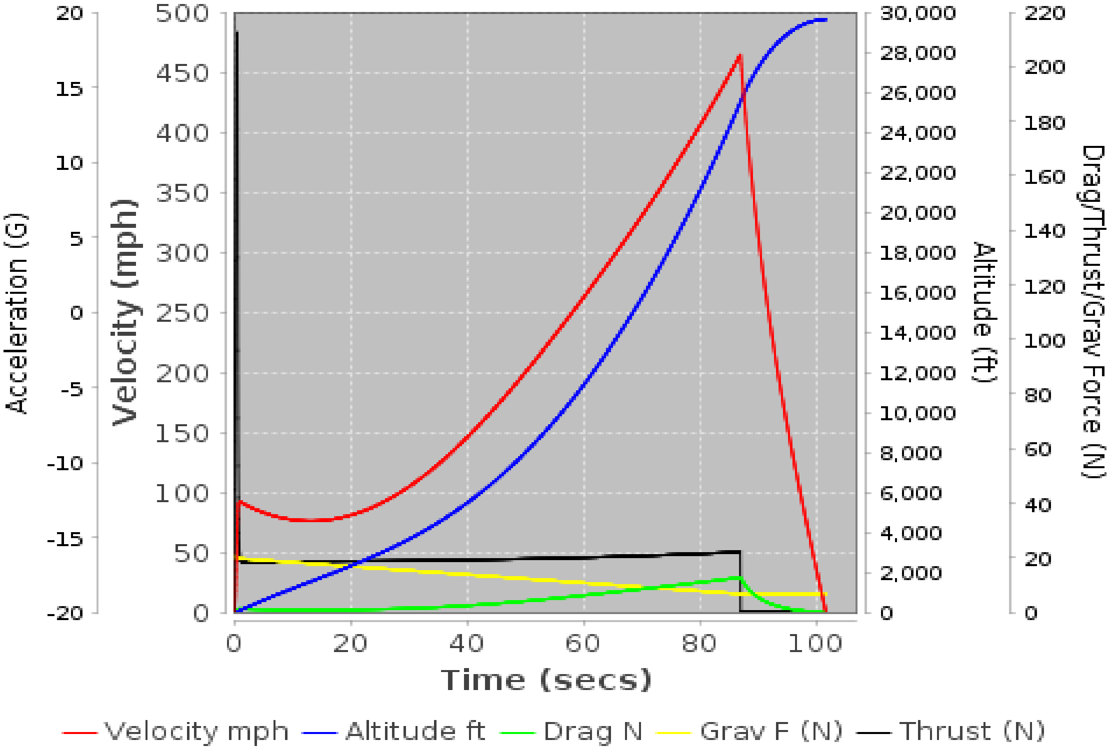

Figure 6: Flight Profile

In this graph, the red curve is velocity, which, as noted, starts out high, but then drops after the core burns out. However, the rocket is still flying, and as propellant burns, the gravitational force (yellow curve) diminishes enough that the thrust (black curve) can overcome it plus the miniscule drag force (green) and start increasing the velocity. However, the altitude is almost certainly going to suffer—for example, a grain length of 28” could result in a lower altitude than a grain length of 16”.

It is at this point that the other main variable in grain design becomes relevant: grain diameter. By increasing the grain diameter, the sustained level of thrust will be increased which solves (or at least mitigates) the problem we just noticed. It will also affect the initial thrust, both increasing it and lengthening its burn time which will be advantageous as well. However, increasing the grain’s diameter requires increasing the rocket’s overall diameter which will increase the drag and the structural (dry) mass, so the increased thrust does not come for free, but by adjusting the core (and thereby the initial velocity) we can gain more than we lose.

If the example rocket’s diameter is increased to 3”, the grain diameter increased to 2.5”, the throat diameter to 0.25” and the grain length to 30”, (but the 2” by 1” core stays the same), we get a profile like this:

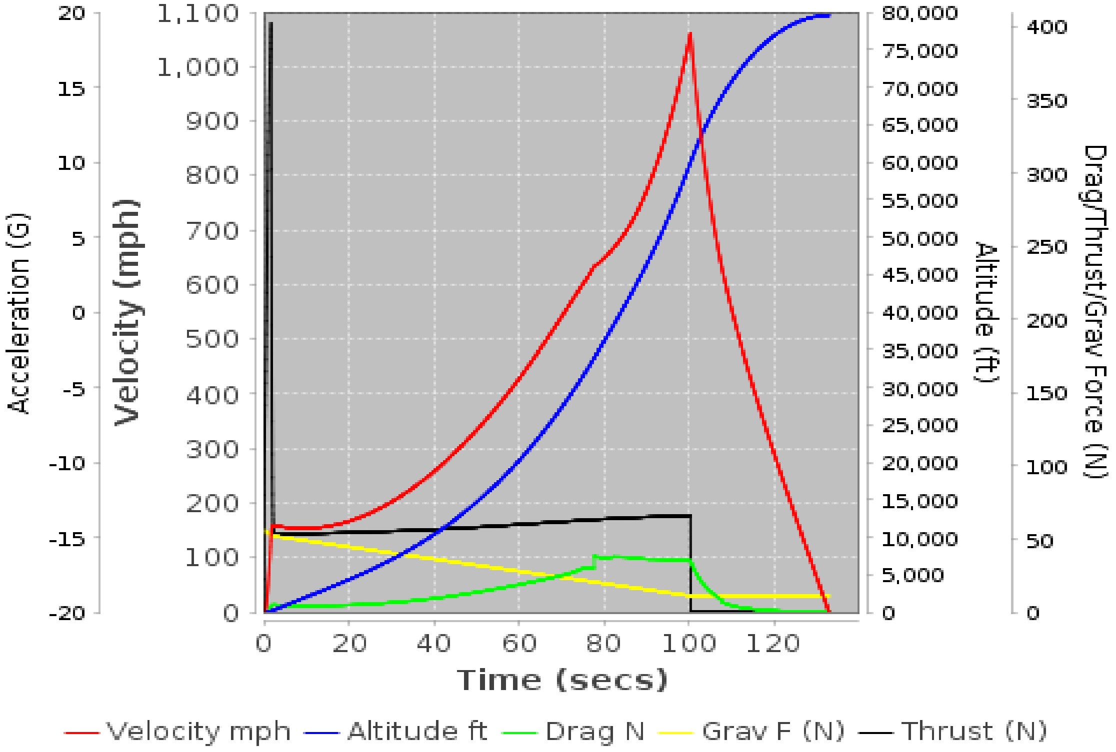

Figure 7: Increased Diameter

This graph bears studying in some detail. In comparison with the 2” rocket, the altitude has vastly increased which shows that we did indeed gain more than we lost by increasing the diameter. The sustained thrust is higher (55-70N—still quite low by HPR standards), and because the rocket is larger and there is a lot more propellant, the gravitational force is also higher. The drag force has increased, as we feared, but there is something fairly interesting about that drag (green) curve: it rises quite gradually and does not exceed the gravitational force until well into the flight where the gravitational force has been cut in half due to propellant burn-off. This is in stark contrast to what happens in a high-thrust/short burn/long coast regime where the gravitational and drag forces are both at their peak shortly after liftoff. An end-burner’s thrust doesn’t have to be large because it does not have to overcome both high drag and large gravitational forces simultaneously.

Another reason a *high-flying* end-burner doesn’t require high thrust (by HPR standards) is revealed by comparing the altitude curve (blue) with the drag curve (green). Notice that at the 60-second mark, the rocket is already at about 20,000 feet by the time the velocity (red) is approaching 450 mph. At this altitude the air density is only 53% of the sea level density which cuts the drag by half. By the time the velocity gets close to transonic speeds, the altitude is around 40,000’ where the density is 25% of the sea level density, cutting the drag to ¼ of what it would be at sea level. Thus, by delaying the high-velocity part of the flight until significant altitude is reached, drag is tremendously reduced. This, again, greatly contrasts with a high-thrust/short burn regime in which the high velocities are reached in the first few thousand feet where the air density is at its greatest.

There is one final advantage an end-burner can have: for a given motor, thrust will increase with altitude. This is true for all motors, but with high-thrust motors that typically run at high (600-1000 psi) chamber pressures and burn out within a few seconds of launch, this effect is negligible. However, demonstrating yet again that the game is different with low-thrust/long-burn motors, the effect can be significant when the chamber pressure is low (say, 150 psi or less) and the motor continues to burn well into a high-altitude flight (say, above 10,000’). Here’s a table showing the effect for a motor having an expansion ratio of 5, a throat of 0.400”, and a chamber pressure of 75 psi:

Altitude (feet)	Thrust (N)	Factor
Sea Level	30.7	1.0
10,000	43.6	1.4
20,000	53.9	1.7
30,000	59.6	1.9
90,000	71.1	2.3

A rocket with this motor had a simulated altitude of about 30,000’ using a constant sea-level thrust, but went to over 80,000’ when altitude’s effect on thrust was considered. Not all end-burners will exhibit such a large increase—as noted above, the effect is highly dependent on chamber pressure. For example, with an expansion ratio of 5, a chamber pressure of 150 psi results in a maximum thrust increase of about 1.4 times that at sea level, and a chamber pressure of 300 results in an increase of only 1.17 times sea level thrust.

Getting back to the main topic of this section, grain design, the process of incrementally adjusting rocket (grain) diameter and core dimensions can, in principle, be continued to achieve whatever altitude one might wish to reach. (One must also consider throat diameter so as not to exceed casing strength.) There are practical considerations, of course, and that is the subject of the next section. But first, a quick summary….

The main point of this section is that geometry can be used to increase the burn area thereby increasing the chamber pressure and therefore the thrust. A fairly simple geometry, that of the core, has been used to illustrate the idea and to show just how much altitude can be achieved with a low sustained-thrust end-burner.

However, other geometrical approaches are possible. One could, for instance, simply put a BATES-style grain segment at the aft-end of an end-burner to achieve even greater initial thrust. Or, increase terminal thrust by also putting a core at the forward end of the grain.

In any case, the idea is to geometrically configure the grain to provide a flight profile that is optimal (for the intended flight) or, at least, *more* optimal than that of the short-burn/high-thrust regime. The strategy is to use low thrust to fly slowly (by HPR standards) through the thickest part of the atmosphere, letting propellant burn-off decrease the gravitational force and the decreasing air density to decrease the drag force until the low thrust naturally increases the velocity.

By manipulating burn *area* to achieve adequate thrust, the difficult issue of increasing burn *rate* through chemistry is thereby avoided.

4.0 Practical Issues with End-burners

To this point, the practical issues of actually achieving the low-thrust, long burns that have been simulated (some might say *cavalierly* simulated) have been ignored in order to streamline the presentation. However, if rockets of this type are actually to be flown, a number of these issues cannot be ignored.

Initially, one might think that there’s not much to do: put a long grain of propellant in a steel casing and send it skyward. Perhaps a bit of casing insulation might be required, but beyond that, what else needs to be considered?

Fortunately, the list does not appear to be long: length/diameter ratio of the rocket, casing insulation, and the flight issues of wind and gravity turns. These will be discussed in turn.

4.1 Length/Diameter Ratio

If the problem with radially burning grains is the constraint of rocket diameter, then the corresponding constraint for end-burners is the length/diameter ratio for the combustion chamber. One cannot simply keep adding more length to the grain forever: at some point the length increases to the point where either the launch forces or aerodynamic pressure will cause the casing to bend. If this bending is mild, then deviation from vertical flight can result; if severe, it can cause breakup.

A very rough rule-of-thumb is that 20:1 is about the limit for most amateur rockets--ratios above that are possible, but require careful structural analysis. Such considerations are beyond the scope of this writeup, but if an L/D ratio of more than 20 is being considered, material selection and sizing will become very important.

However, simulation seems to indicate that before the L/D ratio becomes a limitation, the mass causes the velocity to drop requiring an increase in diameter which in turn results in a reasonable L/D ratio. In other words, the situation appears to be self-correcting. Still, the problem is worth keeping in mind when long combustion chambers are being considered.

4.2 Insulation

Insulating the motor casing is by far the most critical problem facing the construction of a long-burning end-burner. The Arcas developers spent a lot of time and money on it and didn't succeed in finding a material that would allow a fiberglass casing to last for 28-29 seconds. As Adrian Adamson (founder of Featherweight Altimeters) commented in an online post where he discussed breaking the J-class altitude record:

> “If a commercial motor manufacturer ever makes an end-burning 38mm full J, I’m pretty sure the J altitude record would go over 30,000 feet. But It would need one heck of a liner.”

The problem, in a nutshell, is this: there are effective insulators, and there are lightweight insulators, but there are hardly any lightweight, effective insulators. Add to this the problems of forming the insulation into a tube, of cost, of availability to amateurs, and of resistance to combustion chamber temperatures and erosion, and one can start to see the difficulties.

Here is a (very incomplete) table of various kinds of insulation:

Name	Density (g/cc)
Arcas insulation (phenolic/asbestos)	1.5
Commercial paper/phenolic	1.385
EPDM	1.21
Pure phenolic	1.22
Paper (copier paper)	.898
Paper (builder's protective paper)	.670
Bamboo	.5 - .85
Amorim P45 (cork/phenolic)	.33
Fiberfrax	.205
Cork	.161 (varies)
Cryogel	.1385
Balsa	.0277

Work has been done sporadically in this area by the author beginning around 2009 and while progress has been made, no fully satisfactory solution has been found.

Numerous attempts were made to get sodium silicate (waterglass) to work, but the material (when dry) is brittle and composites made with it disintegrated rapidly during combustion.

EPDM, effective in BATES motors, is less effective during long burns and is extremely dense.

Balsa disintegrates quickly.

Bamboo has not been tried. No idea of flame resistance, etc.

Cryogel appears to be unobtainable except as random-sized flakes and chips. Even if it could be obtained, it would have to be preformed into a tube and it’s not clear just how fragile such a thing would be.

Cork sheets aimed at the consumer market fail miserably because the adhesive (which binds the granules together to make the sheet) poorly resists heat; similarly, epoxy and other common adhesives simply melt so composites made with them are useless as insulators. With one exception, no plastic material survives combustion temperatures in reasonable thicknesses.

That exception is phenolic. It has the useful property that it chars rather than melts and is ablative in the combustion chamber. There is one form that is shipped as a dry powder which is then mixed with water. While it is urea-formaldehyde (UF) rather than phenolic, it appears to behave the same. It's obtainable by individuals in one-gallon pails. However, phenolic/UF is quite dense and the UF resin, prior to curing, is quite thick and sticky (about like peanut butter) making it difficult to work with.

The best liner so far has been a convolutely wound tube made from builder's paper (used to protect floors during construction). It is held together with thinned carpenters glue which is not at all resistant to heat. However, this construction seems not to erode excessively during 20-25 second burns. In a 0.25" thickness it appears (with very limited testing) to adequately protect a plastic casing. It is cheap and easily formed into a tube, but is much denser than desired.

A material that is potentially quite light (about 1/3 the density of the paper liner) is a cork/phenolic mixture. There is a commercial supplier (Amorim) of such material but it comes in sheets and and may not be formable into a small diameter (1.5") tube; it isn't clear that it can be obtained by individuals.

Effort has therefore been directed toward forming a UF/cork mixture into a tube by casting or rolling the uncured mixture. The material withstands a propane torch aimed directly at it from a distance of about 2” for a minute which makes it very promising. Very lightweight tubes have been created, but are difficult to form, sometimes have voids, or are generally porous. Density has been similar to, or even less than the Amorim material. Work continues.

The difficulties with insulation may incline some experimenters to just opt for a steel casing, believing that steel would resist the temperatures for the durations involved. There are two problems with this approach: 1) steel is resistant to heat, but it is not impervious to it. Even Arcas, which used a 0.04” thick alloy (4130) steel casing required a 0.150” thick phenolic-asbestos lining. And that was for a 29-second burn. (The need for very carefully designed insulation was proven when early versions of the motor burned through the casing just in front of the nozzle.) 2) steel is very dense, making for a very heavy casing (18 lbs for the 61” long, 4.5” diameter Arcas casing). For this reason, among others, the Arcas team explored the use of fiberglass casings but weren’t able to develop an effective insulation within the constraints of their development program.

As far as weight goes, a steel 2” diameter tube with a 0.040” wall thickness weighs about 26g per lineal inch whereas a carbon fiber casing plus a lightweight liner could weigh as little as 12-15g/inch. However, as noted above, even steel chambers require insulation, so that must be added to the 26g/inch, though perhaps sugar propellants burning at low chamber pressures would allow for less insulation than was required by the Arcas motor. This difference between steel and carbon fiber isn’t large with a short, small diameter chamber, but for a longer/larger one the difference becomes significant.

For example, a 2” rocket with a 6” carbon fiber motor casing simulates to about 4800’; replace the CF casing with 0.040” steel and the altitude drops to about 4268’, about 11% less. With a 17” CF casing the rocket simulates to 19107’ but with a steel one, the altitude drops to 9973’ or about 48% less.

With a larger diameter casing, the situation worsens because of the “minimum gauge” problem, i.e., in a 3” diameter, you can’t obtain a tube with a 0.040” wall thickness, but have to step up to 0.083”.

Also, a steel casing isn’t allowed at Tripoli and NAR launches, so flying such motors would become problematic for many people.

With all this in mind, the author has therefore chosen to focus on insulation, at least until its true limits are established.

4.3 Flight Issues

It has been noted above that unless care is taken, low-thrust flights can sometimes fly rather slowly, either at launch or shortly thereafter. When a rocket is moving slowly it is more affected by wind, so an obvious precaution to take when contemplating end-burner design is to ensure that velocity is adequate for the anticipated wind conditions.

A much more serious problem, however, is the (unintended) gravity turn.

A gravity turn is simply a turn from a vertical trajectory to a more horizontal one with the assistance of gravity. This is intentionally done by rockets heading to orbit because it is more efficient to allow gravity to assist the maneuver rather than doing it purely with guidance propulsion. The turn happens when the rocket is intentionally pitched over slightly from vertical; its thrust vector is then no longer aligned with gravity's vector and the resulting vector is therefore offset from the vertical. This causes the rocket to increasingly pitch over until orbit is reached.

In a rocket intended to fly vertically all the way to apogee, a gravity turn can take place unintentionally if the rocket weathercocks or pitches over for some other reason, such as off-axis thrust. Even though the turn has been started unintentionally, the result is the same: as long as the motor is providing thrust, the thrust vector combines with the gravitational vector and the rotation continues.

This isn't as much of a problem with a typical short-burn HPR motor: in that case, the thrust vector is quite large compared with the gravitational vector (thus overwhelming it) and because of its short duration, only combines with the gravitational vector for a short period. As a result, if the rocket pitches over, no real harm is done, beyond sending the rocket some distance farther afield than might be desired.

But with a motor burning for more than a few seconds, and with less thrust than the typical HPR motor, the turn can become significant and potentially dangerous. Fortunately, those hobbyists flying multi-stage rockets to extreme altitudes have a similar problem; some have therefore begun working on Vertical Trajectory Control Systems designed to keep the rocket flying vertically and several implementations have flown successfully. Some preliminary work in this area has been done by the author, but nothing has been flown yet.

5.0 Concluding Remarks

High-thrust/short-burn motors are well known for resulting in flights characterized by high drag. Simulations done as part of this investigation show that a low-thrust/long burn regime will reduce this drag significantly, with the result that high altitudes can theoretically be achieved with much less propellant and without resorting to staging.

Because propellant chemistry is “hard” and grain geometry is “easy,” grain geometry was used to demonstrate that sugar propellants could, with the right geometry, be used to achieve very high altititudes.

In particular, it was shown that very low thrusts (by HPR standards) are sufficient if the burn time is long enough and the rockets are lightweight and aerodynamic.

Several examples of very lightweight and highly aerodynamic HPR rockets were cited to show that it entirely possible to build rockets that are sufficiently lightweight to use with low-thrust motors.

Since the simulations were done with a typical sugar propellant, this leads to the conclusion that existing sugar propellants are likely more than adequate to achieve very high altitudes. In other words, lack of a “suitable” propellant need not prevent people from serious experimentation with end-burning motors.

Motor casing insulation, however, *does* appear to be a major challenge, at least if lightweight insulation is required. Once that issue is overcome, the flight issue of vertical trajectory control may be a significant problem when attempting high-altitude flights.

The author’s work is currently focused on casing insulation, building very lightweight rockets, and vertical trajectory control.