The Nuclear Option: The feasibility of transport to an interstellar target within a single human lifetime

Jim Cavera, AIM-USA

Abstract

Given the constraint of a single human lifetime, the only possible means of propelling a ship between neighboring star systems is through the use of nuclear propulsion technologies. The first part of this presentation gives an overview of the various strategies available to us (fission, fusion and antimatter) and will attempt to explain the physical principles of each to a general audience.

The second part of this presentation offers a more detailed analysis of a particular propulsion system: a dense plasma focus fusion device capable of driving a manned starship to a fraction of the speed of light. This portion of the presentation will focus, not only on the capabilities, but also on the limitations of such a system. Finally, we briefly discuss some "game-changing" ideas that may hasten the development of these types of interstellar spacecraft.

Introduction

> *"Space," it says, "is big. Really big. You just won't believe how vastly, hugely, mindbogglingly big it is. I mean, you may think it's a long way down the road to the chemist's, but that's just peanuts to space, listen..."*
>
> - Douglas Adams, *The Hitchhikers' Guide to the Galaxy*

Just for the sake of argument, let us propose a trip to the nearest star-system. Let us further place the constraint that we want the transit time to be within one human lifetime; say fifty years. Though, as Douglas Adams so kindly pointed out, space is vast; we do have a few choices of nearby targets. The closest stars to us are (ordered by distance)[1]:

Star	Distance (ly)	Class	Planets?
Proxima Centauri	4.2	M	Unlikely
Alpha Centauri (A and B)	4.4	G, K	Very unlikely
Bernard's Star	5.9	M	Possibly
Wolf 530	7.8	M	Possibly
Lalande 21185	8.3	M	Possibly
Sirius (A and B)	8.6	A, A	Very unlikely
Luyten 726-8 (A and B)	8.7	M, M	Very unlikely
Ross 154	9.7	M	Possibly

Of these choices, most of them have certain issues that make them unappealing targets. Though quite close, the Alpha Centauri system is a binary and is unlikely to have planets. Proxima Centauri may have planets but the likelihood is low as surveys have ruled out anything the size of a "super-earth" or larger. Additionally, it is a "flare star" and undergoes random increases in luminosity, and may be a distant companion to the Alpha Centauri system as well.

Sirius is another binary system, and both stars are probably too hot and too large to have an interesting planetary system. Luyten is yet another binary, with the added distinction of having a white dwarf as one of the pair. Ross, Wolf and Lalande, though potentially interesting, are pushing the limits of distance.

---

[1] Stellar data obtained from *The Cosmos - Astronomy in the new millennium, 3rd ed.*, by Pasachoff and Filippenko

Bernard's Star is both quite close, not a member of a binary system, and may well support a solar system of its own. It is worth noting though, that recent astrometric data indicates that if it does have a planetary system, it would consist of a few, low-mass objects. Still, for the purpose of this thought experiment, we shall assume that this will be our target.

Now for our proposed starship. We are going to assume that we want the voyage to complete within a single human lifetime: fifty years. This obviates the issues of needing a generation ship, both in terms of life-support and social impact on the crew. This is not to say that the voyage will be for the faint of heart: the crew will initially be quite young and will arrive at the target when they are well into their seventies. And unless we postulate travel at a significant fraction of the speed of light (so that relativistic effects become pronounced), it shall be a one-way journey. I'll leave the social and mission-planning aspects of this to others.

Let us, for sake of argument, assume a crew of one-hundred. This will allow for some redundancy as well as for some flexibilities in their roles. And that size of crew is probably sufficient to allow for reasonable amounts of social interaction. Again, I'll leave the specifics of this to others.

So this places some constraints on the physical size and mass of the starship. If each of the crew needs one-hundred cubic meters of living space (the size of a small studio apartment) and a like amount of working and community space, then we must sustain a habitable volume of twenty thousand cubic meters. We'll allow for an additional ten thousand cubic meters for interconnects and access points. This volume must be sufficiently lit and suffused with breathable air and kept to a comfortable temperature within the depths of space.

Additionally, food and water must be provided. For sake of resource management, we will assume a fully-closed system that makes use of hydroponics to provide half of their food and all of their oxygen while reprocessing their waste. The remainder of their food and a certain reserve of water will need to be stored.

Also note that when they reach their destination, their needs do not stop. We will have to allow for the completion of their mission once the target system is reached. We shall assume a twenty-year mission on target for the purposes of this discussion.

An off-the-cuff estimate for their hydroponic needs is fifty cubic meters of “farm” per person, or about five thousand cubic meters in total, allowing for inefficiencies of design. The food stores not provided by their garden amount to about one liter per person per day for seventy years! Doing the math, this give us about three thousand cubic meters with reserves. Water reserves must be factored in as well, as closed systems (despite our best efforts) are never completely closed. We shall assume a water reserve equal to the food reserve. As it turns out, the water reserve (hopefully never needed) can have dual use, acting as additional shielding around the habitation module.

Power needs will likely be met via some sort of nuclear reactor; possibly via the propulsion system itself. For our community, we shall assume that ten megawatts will be required. This may sound excessive for just one hundred people, but note that this must run not only the life-support systems, but all of the scientific instruments as well.

One final constraint is that we will assume that there is some additional equipment to provide for rotating habitation modules. It is highly doubtful that a person could survive (let alone thrive) in a microgravity environment for their entire life.

So here we have some basic constraints and, using these, we can make some initial estimates as to the size and mass of the proposed starship.

For the superstructure, we shall assume a spherical habitation module. Given the approximate interior volume of fifty thousand cubic meters (again, allowing for inefficiencies), this means a sphere with a diameter of approximately forty-six meters (without regards for external instruments and shielding). This may seem a bit small, but realize that it would be about the same size as a fifteen story building.

As an initial mass estimate, we'll just take a stab at it and give the entire sphere an average specific gravity of one (specific gravity being the density as compared to water). Yes, the metal parts are much heavier, but there's a lot of empty space as well. So the habitation module (again, *sans* shielding) is approximately fifty thousand metric tons in mass (one metric ton equaling one thousand kilograms). Let's add in the shielding and make it five times the thickness of what is planned for some of the human Mars missions. Again, just pulling that number out of the air, but with a reason: the work of spacecraft design has already been done for us in the form of the NASA Mars Reference Mission.

Using these existing benchmarks, and adding a *huge* safety margin, gives us a shielding requirement of a meter-thick assemblage of metal and composite material with an average specific gravity of three. Again, returning to the calculator, this increases the mass of the habitation module by an additional twenty-four thousand metric tons. For those of you keeping score, our running total is up to seventy-four thousand metric tons, and this is just for structure and habitation (no power plant, propulsion system or fuel, as of yet).

We've an additional, specialized shielding requirement that normal planetary missions do not have. Because we intend to travel at a small (but still significant) fraction of the speed of light, any impact on the front of the vehicle could have catastrophic consequences. To mitigate this, I would propose a “bumper”, made of the same shielding material: a disk, four meters thick and fifty meters in diameter. We are actually going to need two of these bumpers: one for the front and one for the back. This may seem odd, but realize that the spacecraft will have to decelerate once it gets near the target system. For that phase of the journey, the ship will flip, end-for-end, such that the main engine is pointing forward. Each bumper has an additional mass of another twenty-four thousand metric tons, bringing our running total to a bit over one-hundred twenty thousand.

I would add one other thing to our running total of mass: space probes. I propose a large number of orbital probes and landers that could be deployed to points of interest. These would, of course, not be contained within the habitable volume, but would be arrayed along the outside, likely on the struts connecting the power and propulsion to the habitation module. Again, just for sake of argument, we'll assume twelve landers, each the size of the Mars Curiosity mission; and twelve orbital probes, each the size of the Cassini mission. This comes to three metric tons for each of the orbiters and 1.5 for each of the landers. Another sixty-thousand metric tons to bring our total to just about one-hundred eighty thousand.

Now to make this go.

Basic propulsion physics

There are three basic parameters that can be used to characterize a propulsion system: thrust, exhaust velocity, and specific impulse. The first is a direct result of Newton's laws of motion: force (thrust) equals mass times acceleration. From this, if you want a specific value of acceleration (say, to overcome gravity) then you need a value of force in proportion to the mass of the craft. Higher thrust means more acceleration which means that you reach your target velocity more quickly.

The amount of thrust that a rocket engine can produce is dependent on both the energy of the exhaust (the chemistry of the reaction, in the case of a solid- or liquid-fueled rocket) and the mechanical configuration of the engine itself. Thrust is generally measured in newtons (in metric units) or pounds (in English units). One newton is the amount of force required to impart an acceleration of one meter per second per second, to a one kilogram mass. One pound (force) is the amount required to support one pound (mass) against gravity.

Some example values for rocket engine thrust[2]:

Deep Space 1 ion engine	0.1 N	0.008 lbs
C6-3 model rocket engine	6 N	1 lb
Bell jet pack	1250 N	280 lbs
GE Honda HF-120 turbofan engine	9100 N	2050 lbs
Thiokol XLR99-RM2 (X-15) rocket engine	310,000 N	70,000 lbs
RS-68 (Delta IV) rocket engine	3,300,000 N	740,000 lbs
RS-25 Space Shuttle Main Engine	5,450,000 N	1,225,000 lbs
Space Shuttle Solid Rocket Booster	12,000,000 N	2,700,000 lbs

Specific impulse is the difficult one: it can be thought of as a measure of the efficiency of a propulsion system. In a car, you have MPG (or km/l) as the efficiency rating. That number will tell you how far you can get if you just have a gallon left in your tank. Obviously, this is absolutely meaningless in space.

Instead we have specific impulse, as measured in seconds. Though, again, the math can be a bit complicated, one can think of it as the number of seconds that one pound (mass) of fuel can produce one pound (force) of thrust. A one-pound solid rocket motor with a specific impulse of two hundred can produce one pound of force for two hundred seconds. Or, more likely, one-hundred pounds of force for two seconds. Some common figures for specific impulse are as follows[3]:

Deep Space 1 ion engine	3200 sec.[4]
C6-3 model rocket engine	82 sec.
Bell jet pack	122 sec.
GE Honda HF-120 turbofan engine	3000 sec.[5]
Thiokol XLR99-RM2 (X-15) rocket engine	279 sec.
RS-68 (Delta IV) rocket engine	410 sec.
RS-25 Space Shuttle Main Engine	453 sec.
Space Shuttle Solid Rocket Booster	268 sec.

Finally, the exhaust velocity of a propulsion system places a limit on the upper speed of the vehicle. Though a huge over-simplification (the precise physics of this can be complex), one can imagine that the maximum speed of a vehicle is approximately proportional to the exhaust velocity. In reality, one must analyze this on a case-by-case basis, taking into account the mass of the exhaust and the conservation of momentum. But in general, a higher exhaust velocity means a faster vehicle.

These three parameters can be applied to any propulsion system that creates thrust by throwing something out

---

[2] These figures come from the Wikipedia entries for the respective engines. They are for illustrative purposes and are not to be taken as absolutes.

[3] Again, these figures come from the Wikipedia entries for the respective engines. They are for illustrative purposes and are not to be taken as absolutes.

[4] The ion engine aboard the Deep Space 1 mission is incredibly efficient, but does not use chemical means of propulsion. Rather, it uses electrostatic forces to accelerate a working fluid directly, as opposed to burning it. Though this paper will focus on nuclear propulsion systems (simply because large ion engines are not practical) it is worth noting that this technology could be a viable alternative in smaller, unmanned missions.

[5] Note that the reason for the very high specific impulse of the turbofan engine is that it (unlike the other entries on the list) is an air-breathing engine. Because of this, most of the “fuel” is not carried on-board, but supplied by the medium through which it operates. As one can imagine, this increases the “efficiency” in the sense that most of the combustion reactants are not in-turn propelled by the engine, drastically saving weight. Were one to take into account the mass of the oxygen required for combustion, the specific impulse would be on par with other chemical rocket engines; *i.e.,* in the 100 to 400 second range.

the back. These are termed "reaction engines". Not all propulsion systems work in that way. In particular, solar and magnetic sails do not throw anything out the back, but work by deflecting the solar wind. Asking for the specific impulse of a solar sail is a lot like asking how many miles per gallon a sailboat gets.

Now for some equations of motion. Let us assume an initial period of acceleration, followed by a long coast through interstellar space, and then a period of deceleration to arrive on target. Given the fact that we want our trip to take fifty years, our coasting phase must occur at a minimum of about ten percent of the speed of light. Note that this does not account for the time spent accelerating and decelerating!

Were we to want to minimize the thrust required, we would have to spend half the voyage in acceleration, flip at the midpoint, and spend the other half in deceleration. This mission profile allows for the highest possible speed given a particular amount of thrust, assuming of course that the operational duration of the engine is sufficient to allow for continuous firing for the entire fifty year voyage.

That sort of continuous firing is an unrealistic expectation. Our best chemical rockets have only fired continuously for tens of minutes. Even the long-duration engines with which we have experimented *(e.g.* the ion engine aboard the Deep Space 1 probe) have only fired for months at a time and have experienced degradation during that time. Since we are discussing a mission of decades in duration, we have to account for such engine degradation so that our main engine is not "used up" in the acceleration phase.

The limitations of chemical rocket engines

Let us first analyze the mission profile, assuming that we will power the spacecraft with the best conceivable chemical rocket engine. This hypothetical engine will have the following parameters:

Fuel	liquid hydrogen and liquid oxygen in stoiciometric ratio
Vacuum Thrust	100,000,000 N (25,000,000 lbs.)
Specific Impulse	500 seconds
Exhaust Velocity	5000 m/sec.
Fuel Consumption	60,000 kg/sec.
Engine Mass	1000 metric tons
Fuel Containment[6]	0.1% of the mass of the fuel

These figures are extrapolated from the RS-25 engine and, knowing what we know about chemical rockets, this is probably the best that we can do. Now it gets interesting: the engine has to accelerate not only the mass of the payload, but the dry mass of the propulsion system, fuel containment, and mass of all of the remaining fuel as well. This results in a tricky-to-solve differential equation for the rocket's motion. Fortunately for us, computer simulations come to the rescue.

A few things about this simulation: First, we will assume that the engineering difficulties of firing a chemical rocket for extended periods of time have been solved. This is *not* a reasonable assumption. Chemical rocket engines are highly-overloaded systems. One can liken it to another system designed to contain and direct combustion – a firearm. In a gun, the useful lifetime of a barrel is measured in seconds. In much the same way, the combustion chamber and nozzle of a chemical rocket is designed to ablate and degrade in a controlled and predictable fashion. We shall cheerfully ignore this limitation for now.

Second, we shall assume that the starship starts from rest (zero initial velocity). Of course, it will be orbiting and so will have a non-zero velocity. We are also going to ignore any velocity gains from slingshot trajectories as

---

[6] Assuming an additional one-tenth of one percent of the total fuel mass for fuel containment and delivery is *very* rough approximation, but a not-unreasonable one. As the mass of the fuel increases (and, as we shall see, it increases considerably) the total mass devoted to containing it also increases, but not linearly. For smaller fuel systems, this is woefully inadequate, but for very large fuel payloads, it is more acceptable. We are making this assumption purely for the purposes of simplifying the example.

well as velocity losses from breaking orbit and associated gravitational interactions. Since our ultimate velocity goal is a significant fraction of the speed of light, such gains and losses are almost round-off error.

Finally, we will only model half of the trajectory. We would, of course, want to decelerate at the end of the journey. But for now, we shall just look at the difficulties involved in getting up to speed, not back down to orbital insertion.

For the purposes of this exercise, I created a program to properly model the dynamics of the system. It's a fairly simple, discrete solver for the differential equations of motion that breaks the problem down into what the system is doing on a second-by-second basis. At each iteration, it subtracts a little bit of fuel, calculates the kinematic parameters, and moves on to the next second.

For our initial run-through, we'll just take a rough guess and say that there is one thousand times more fuel than payload. The input parameters are then:

Dry Mass (payload plus engines):	181,000 metric tons
Propellant Mass:	181,000,000 metric tons
Propellant Delivery and Containment	181,000 metric tons

Note that this initial pass results in an incredible volume of fuel. Such an amount of fuel would occupy a spherical volume on the order of a kilometer in diameter. At this size, our starship could easily be mistaken for a medium-sized asteroid.

Dropping all of these parameters into the simulation yields the following:

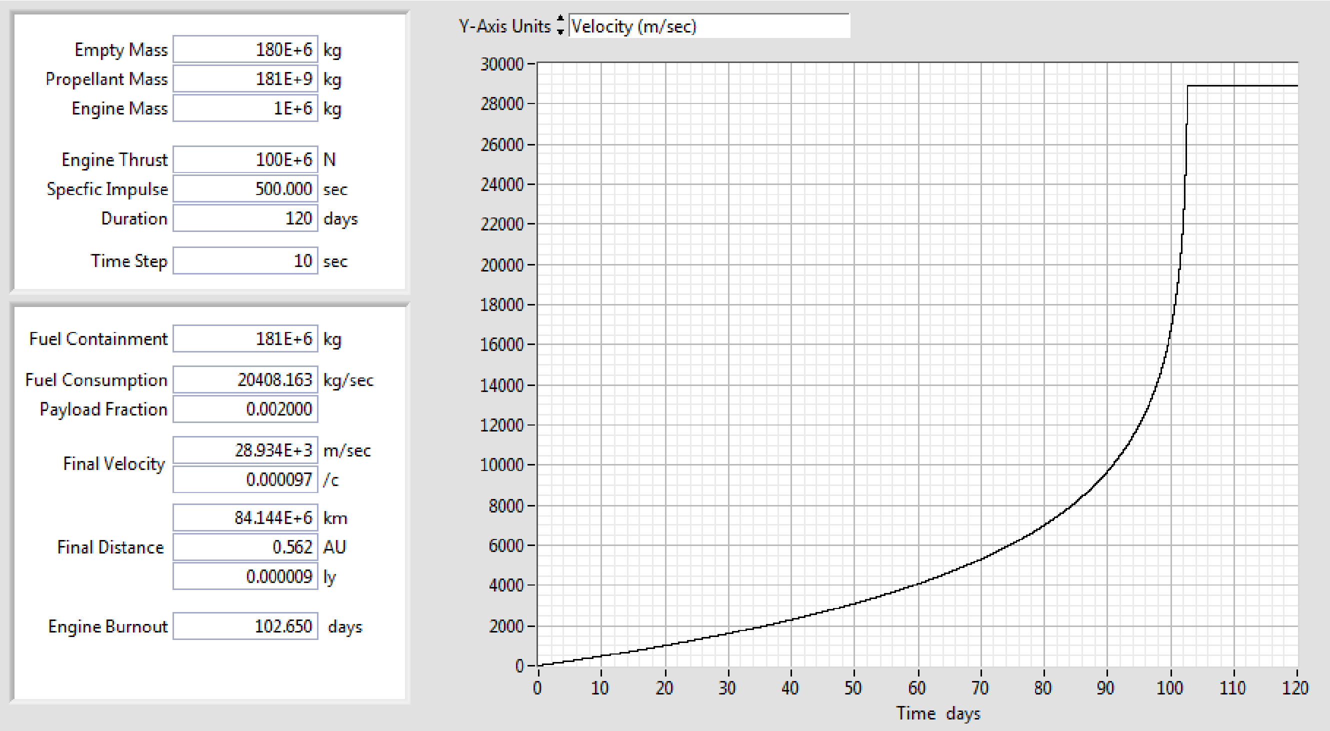

Empty Mass 180E+6 kg
Propellant Mass 181E+9 kg
Engine Mass 1E+6 kg

Engine Thrust 100E+6 N
Specfic Impulse 500.000 sec
Duration 120 days

Time Step 10 sec

Fuel Containment 181E+6 kg

Fuel Consumption 20408.163 kg/sec
Payload Fraction 0.002000

Final Velocity 28.934E+3 m/sec
0.000097 /c

Final Distance 84.144E+6 km
0.562 AU
0.000009 ly

Engine Burnout 102.650 days

Y-Axis Units Velocity (m/sec)

Figure 1. Chemical propulsion system – first pass

As the graph shows, the engines exhaust the fuel supply in a mere 103 days. At the end of that time, the ship is only traveling at 29 km/sec, or less than one-hundredth of one percent of the speed of light. In addition, since most of that time was spent just pushing fuel-mass around, the craft has only traveled about half the distance from the Earth to the Sun. At this rate, our voyage to Bernard's Star would take about 600,000 years. Not quite the short trip.

Let's change the parameters a bit: we'll increase the engine thrust and mass by a factor of one hundred , and the fuel load by a factor of one thousand.

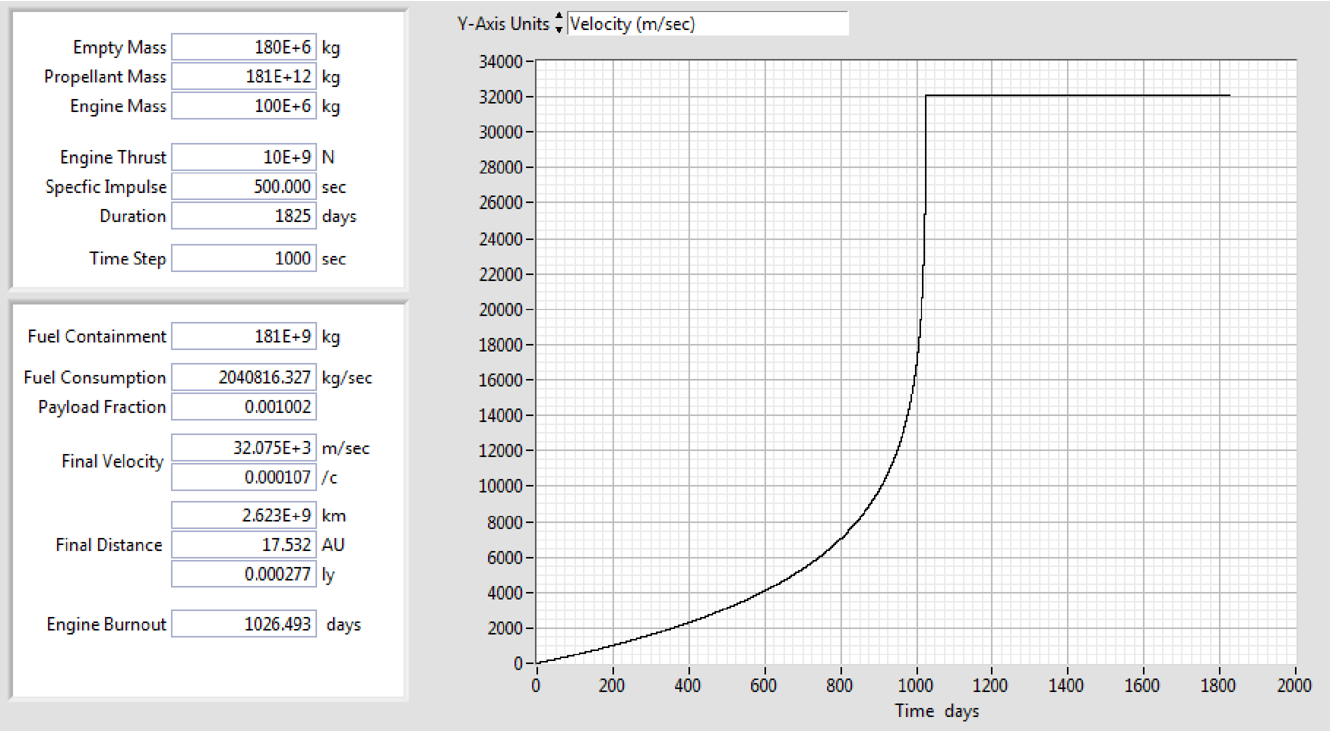

Empty Mass 180E+6 kg
Propellant Mass 181E+12 kg
Engine Mass 100E+6 kg

Engine Thrust 10E+9 N
Specfic Impulse 500.000 sec
Duration 1825 days

Time Step 1000 sec

Fuel Containment 181E+9 kg

Fuel Consumption 2040816.327 kg/sec
Payload Fraction 0.001002

Final Velocity 32.075E+3 m/sec
0.000107 /c

Final Distance 2.623E+9 km
17.532 AU
0.000277 ly

Engine Burnout 1026.493 days

Y-Axis Units Velocity (m/sec)

Figure 2. Chemical propulsion system – turned up to 11

We don't get much faster, but we do get a lot farther before engine burn-out. Now there is enough fuel to run the engines for nearly three years, continuously. But in that three-years time, we max out at approximately the same velocity as before due to the fact that the engines have to move so much more mass. So this is really no better, and in fact a little worse as it takes much longer to reach a speed that is just a little higher than before.

Also note that our fuel has now absolutely dwarfed the spacecraft. Containing nearly two hundred *trillion* kilograms of fuel would require a spherical container approximately seventy kilometers in diameter. Just to give you a sense of scale, Saturn and Jupiter have moons that are smaller than this.

Just to play with the numbers, let us take the original figures and, rather than increasing the thrust or the fuel, we shall increase the efficiency by a factor of ten, up to five thousand seconds.

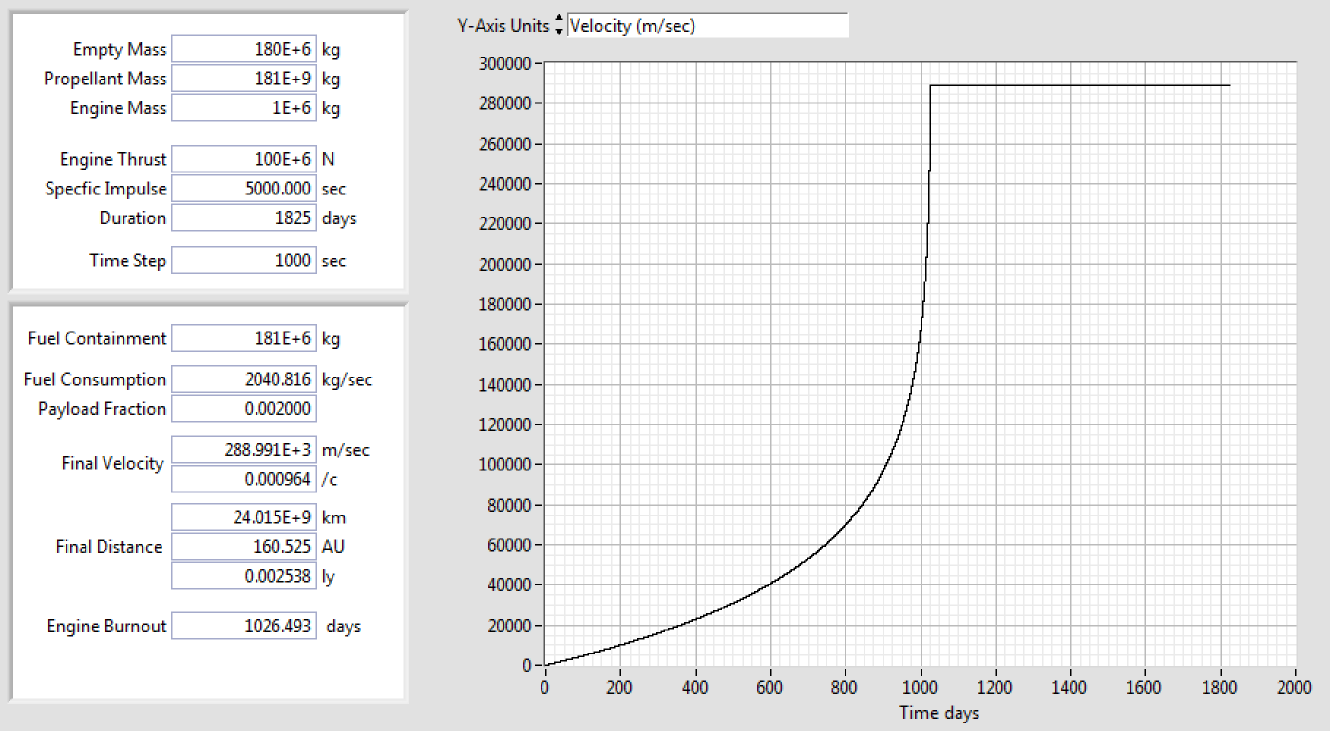

Empty Mass 180E+6 kg
Propellant Mass 181E+9 kg
Engine Mass 1E+6 kg

Engine Thrust 100E+6 N
Specfic Impulse 5000.000 sec
Duration 1825 days

Time Step 1000 sec

Fuel Containment 181E+6 kg

Fuel Consumption 2040.816 kg/sec
Payload Fraction 0.002000

Final Velocity 288.991E+3 m/sec
0.000964 /c

Final Distance 24.015E+9 km
160.525 AU
0.002538 ly

Engine Burnout 1026.493 days

Y-Axis Units: Velocity (m/sec)

Figure 3. Original chemical propulsion system parameters, but with ten-times the efficiency

Note what has happened: Though we have the same amount of fuel as in the initial simulation, it lasts ten times longer. This allows for a proportionately longer period of acceleration and results in a much higher speed. Though we're still short of what we'll need for the full mission, we are much closer with a final velocity of about one one-thousandth of the speed of light. So how can we increase the efficiency of a rocket engine? A chemical rocket engine generates its heat by the combustion process. But there are, of course, many other ways of heating (and thus, accelerating) the propellant.

A very brief introduction to nuclear physics – fission reactors

There are many different sorts of natural radioactive decay. Radioactive nuclei can spontaneously spit out charged particles of many sorts: electrons, photons, helium nuclei, etc. But there are a few elements that can spontaneously split apart entirely. Instead of just a particle or two, their decay mode splits them into two pieces with a few neutrons left over. Of the few elements that have a spontaneous fission decay mode, isotopes of thorium, uranium and plutonium are useful in nuclear reactors as fuel.

Here’s the idea: get a bunch of fissionable material together, introduce a few neutrons, and stand back. The fissionable atoms absorb the neutrons and, since they now have far too many, split into two nuclei and a bunch of other neutrons. These newly liberated neutrons go on to split other atoms, which in turn produce more neutrons, etc.

If this reaction goes too quickly, you get a weapon rather than a reactor. But by controlling the density of the fuel and by selectively adding neutron-absorbing and neutron-reflecting materials, the reaction rate can be controlled so

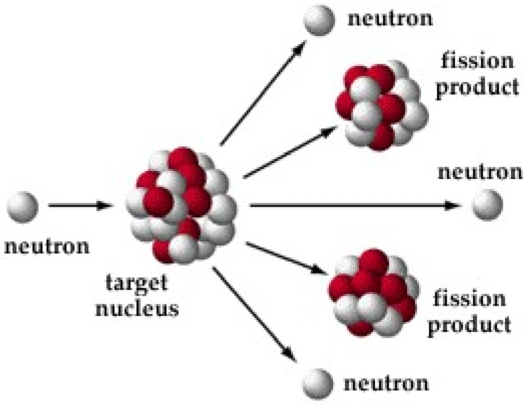

that useful heat can be removed without destroying anything. On the ground, this heat is used to make steam, which then spins a turbine. In space, this heat can be converted to mechanical power via a Sterling-cycle or Brayton-cycle engine, or can be converted directly to electricity using thermoelectric or thermionic converters, or to propulsion by heating and then ejecting a working fluid. Here comes the up side of all this: A fission reaction produces many, many times the heat of either a chemical reaction, or of natural radioactive decay.

In a normal (terrestrial) reactor, the core gets very hot and the heat is converted into electricity by way of a heat engine of some sort. As with all heat engines, this process is inefficient, leading to wasted heat. So let's flip the problem on its head: Instead of using the heat to generate power and throwing away the excess, we can heat a working fluid, throw it out the back and generate some power as well. Instead of a power reactor, we have a rocket that uses excess heat to generate power. Early on in the United States space program, a system was devised to do exactly that.

All rockets generate thrust by heating (i.e. accelerating) a working fluid and throwing it out the back. Chemical rockets do it be burning the fuel with the liberated energy heating the combustion products. Nuclear reactors can heat a fluid by simply mixing it with the reactor core. The reactor is cooled (this is good) and the working fluid is superheated (also good). The ROVER and NERVA projects, which took place from 1958 through 1972 (some sources cite different dates), produced working nuclear thermal rockets. The specifications for the constructed NERVA engine are a chemical propulsion engineer's dream[7]:

Engine Dry Mass:	34 metric tons
Thrust:	867,000 N
Specific Impulse:	835 sec.

The engine was test fired several times, including mid-firing stops and re-starts. It was run continuously for up to 65 minutes and produced full thrust during the entire time with minimal engine degradation. The nuclear fuel consumed was negligible.

So what happened to this technology? Clearly, this is far better than any chemical rocket and could easily take a manned spacecraft to any point in the solar system. The answer is twofold. First, there are environmental considerations. Though the liquid hydrogen fuel and exhaust could never become radioactive, the engine components did as the metallic parts absorbed neutrons. Mitigating this would require massive amounts of shielding as well as a means of storing, replacing and disposing of engine components. In addition, radiation damage causes metal to become brittle, leading to more structural failures. These are all technical issues that, though difficult, could be overcome.

The second point is political. No one in congress wanted nuclear rockets anywhere near their constituents. Environmental action groups were lined up, waiting to protest. As with many things nuclear, the fear outpaced the reality. A properly designed NERVA type engine could never produce the same levels of contaminants as even a small power reactor. Though the danger is real, it can be managed effectively.

The NERVA type engine is starting to make a comeback. Variants are being designed that nearly eliminate the possibility of nuclear contamination. In fact, Pratt and Whittney introduced its Triton rocket design in 2005 – a new, (what would be) off-the-shelf, nuclear thermal rocket based on the old NERVA designs.[8]

Of course, the NERVA design in general is based on a terrestrial power reactor. If we are less concerned with power and want to devote as much energy as possible to propulsion, we can do much better than this. Recall from basic thermodynamics that a heat engine's efficiency increases with temperature. In a conventional nuclear reactor, the temperature at which parts start to melt limits the core temperature. At about 1300 degrees

---

[7] Figures from *The Nuclear Rocket*, by James Dewar

[8] See AIAA 2004-3863: *TRITON: A Trimodal capable, Thrust Optimized, Nuclear Propulsion and Power System for Advanced Space Missions,* by Joyner, *et al*

centigrade (2300 degrees Fahrenheit) the fuel rods start to weaken. Because this is considered a bad thing (*c.f.*: the China Syndrome), the temperature is limited to less than this. But suppose we could run much hotter…

At about 2300 degrees centigrade, the core is a molten mess. If a conventional power plant were to ever get to this temperature, another Three-Mile Island or Chernobyl would be the result. But a space-borne reactor has an advantage over a terrestrial one: no gravity to get in the way. In the micro-gravity environment of space, a liquid or gaseous blob of uranium can be stable without needing to interact with a container. The advantage of this: more heat. A gas-core reactor can reach temperatures of tens of thousands of degrees, giving a specific impulse of around eight thousand seconds.

This may all sound like pie-in-the-sky, and for the most part, it is. There is one exception to this: A gas-core design called a "nuclear light bulb" has received some serious consideration. This is a quartz bulb filled with highly enriched uranium gas and surrounded by a graphite reflector. A working fluid is passed around the quartz bulb to remove heat as fast as it can be produced. Though not as hot as a true gas-core reactor, it can be considerably hotter than a reactor with solid fuel rods (about 8000 degrees as one has to keep the quartz from melting).

United Technologies performed a long-term study of the nuclear light bulb concept for NASA. Their final 1979 report cited the following performance parameters (all of which are theoretical, as a prototype was never built) as compared to the NERVA engine:

	NERVA	Nuclear Light Bulb
Empty Mass:	34 metric tons	32 metric tons
Vacuum Thrust:	867,000 N	409,000 N
Specific Impulse:	835 seconds	3200 seconds

Note that if you were to make a reasonable extrapolation of nuclear light bulb technology, it begins to look similar to our most recent simulation. Clearly, we are in the ballpark on this:

Propulsion System Mass:	3,000 metric tons
Engine Thrust:	40,000,000 N
Specific Impulse	4000 seconds

Though many technical hurdles remain, the nuclear light bulb concept is superior (in theory) to the traditional nuclear thermal rocket. The main issue facing the nuclear light bulb is handling such high temperatures. Once we can do that, probably by making use of intense magnetic fields, then fusion and antimatter rockets are not far behind.

More nuclear physics – fusion reactors

Nuclear fusion: the "holy grail" of nuclear power research and perpetually twenty years in the future. When people speak of mining the "vast reserves of energy in the form of helium-3" on the surface of the moon, it was to nuclear fusion that they refer. So why is fusion better than fission?

A typical fission reaction spits an atom of a fissile material (plutonium, uranium or thorium) using a stray neutron. This liberates both energy and more neutrons that can in turn split other atoms. The reaction looks like this:

Energy In	Reaction	Energy Out
0	U^{235} + n ▷ (some fission products) + 2 n	200 MeV[9]

---

[9] From *Nuclear Principles in Engineering*, Tatiana Jevremovich

Note that there is a column labeled “Energy In”. In a fission reaction, this is effectively zero as such a reaction requires no energy to get it going. Also you may note the unfamiliar units for energy: units of electron-volts (or in this case, million-electron-volts) are routinely used to measure small quantities of energy, particularly when discussing reactions in particle physics. Just to give you a sense of scale, the kinetic energy of a single molecule of atmospheric gas at standard temperature is on the order of 1/40 of an electron-volt. Now compare that to what is available from chemical combustion[10]:

Energy In	Reaction	Energy Out
0	$U^{235} + n \triangleright \text{(some fission products)} + 2\,n$	200 MeV
0	$2\,H_2 + O_2 \triangleright 2\,H_2O$	5.9 eV

Clearly, the energy yield of a fission reaction is orders of magnitude higher than that of a chemical reaction, and of course, much much hotter than room temperature. More heat liberated in the reaction equates to a much hotter exhaust and a much higher engine efficiency.

Fusion reactions are something of the inverse of fission: rather than splitting a single, heavy atom apart, the reaction forces two, lighter atoms to combine. Energy is released as the new atom has a lower-energy (*i.e.*: more stable) configuration than the starting atoms. Of course, this comes at a cost, and that cost is that the fusion reaction requires energy in to the system in order to overcome the natural repulsion of the two nuclei.

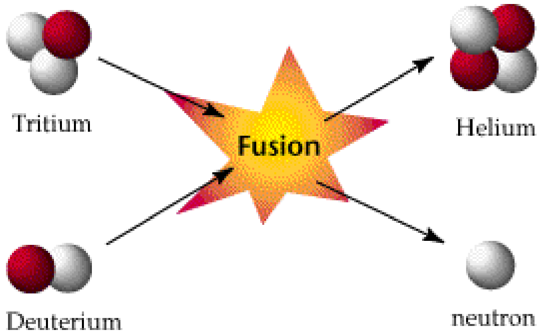

This input energy is in the form of heat. In practice, an aggregate of gasses is heated (lasers and radio waves being the most common means) and contained at high pressure. At a particular point, their kinetic energy is enough to overcome the electrostatic repulsion and to cause them to fuse together. The energies available are on the order of[11]:

Energy In	Reaction	Energy Out
3.5 keV	$H^2 + H^3 \triangleright He^4 + n$	17.6 MeV
30 keV	$H^2 + H^2 \triangleright He^3 + n$	3.3 MeV
30 keV	$H^2 + H^2 \triangleright H^3 + H$	4.0 MeV
300 keV	$H^2 + He^3 \triangleright He^4 + H$	18.3 MeV

It is worth noting that every 10 keV (thousand electron-volts) corresponds to a temperature of approximately 100 million degrees. This is the primary reason that fusion reactions, unlike fission, are not self-starting; and this is why we don’t yet have commercial fusion reactors.

---

[10] Figures from *Chemistry, 12th ed.*, by Brown, *et al*

[11] Figures are from *Fusion Research (vols. 1-3)*, by Dolan

In addition to heating issues, fusion reactors also have containment issues. Most terrestrial fusion experiments use magnetic fields to attempt to contain the fusion plasma, with the classic analogy being, "containing jell-o with rubber bands". Since any contact with a physical containment vessel would both cool the plasma and damage the container, intense magnetic fields are really the only viable option. Note that in limiting the discussion to space-borne propulsion systems, we can gloss over many of the containment issues. These devices will be operating in a vacuum in the absence of strong gravitational fields and will also be designed to jet some of the plasma as well. For our purposes, containment is a much more manageable issue.

Also notice that the fission reaction produces hundreds of MeV's (million electron-volts) of energy while the fusion reactions produce tens. Why then do we even bother with fusion? Because uranium is approximately one hundred times more massive than an equivalent amount of hydrogen. The energy released by fusion is therefore three to five times higher than that released by fission given an equal weight of fuel. A fission rocket needs both the heavy elements to produce the reaction and a lighter working fluid to produce thrust. A fusion reactor does away with the heavier fuel entirely, generating thrust and power directly from the plasma of the reaction. From the standpoint of propulsion system design, this mass savings is very important.

There are many concepts for fusion rockets, but I'll confine the current discussions to those that have been most highly developed. The oldest of these ideas are the Orion and Daedalus concepts. For both of these concepts, not only is containing the plasma completely unimportant, it is absolutely impossible.

The Orion concept is a truly scary contraption designed in the fifties and sixties by General Atomics. The basic point is that you load up a rocket the size of a large building with hundreds or even thousands of hydrogen bombs. A contraption similar to a tennis ball machine fires the bombs out the back at the rate of about one every second. These explode (of course) with the shock wave being absorbed by a large, shock absorber mounted, pusher plate. The bombs are designed as shaped charges to deliver most of their energy to the spacecraft, propelling it forward in a jerky fashion. This series of explosions would give the spacecraft an average thrust in the millions of newtons with a specific impulse of about fifty thousand seconds.

This was (and still is) a very serious concept with quite a lot of research around it. Though riding a shock wave in that fashion sounds like a risky proposition to say the least, the simulations and even physical experiments that were performed indicate that it is a practical means of propulsion.

It does tend to fall short for traversing interstellar distances, however, as the number of nuclear warheads that would have to be carried becomes prohibitively large. One warhead weighing approximately one hundred kilograms, fired once per second for about a year, yields a "fuel consumption" rate of about three million metric tons per year of firing. Computer simulation gives numbers that would get us to another star within a reasonable time, but at a horrible cost in terms of starship mass.

The Daedalus concept came out of the British Interplanetary Society during the seventies and is similar to Orion in that it uses pulses of nuclear explosions. In contrast to the Orion project, this concept went with a more refined approach of firing fusion fuel pellets out the back and igniting them with lasers or electron beams. Each pellet produces far less impulse, but they are ignited hundreds of times per second. Average thrust is lower, but efficiency can be much higher – up to five hundred thousand seconds by some accounts. In addition, since the fuel pellets are not bombs themselves, there is none of the danger (whether it real or imagined) of a catastrophic accident involving the ignition of one or more warheads. In contrast to the Orion rocket, the Daedalus rocket would be smaller and lighter, and could therefore reach a much higher velocity.

It can be seen from the aforementioned tables that the specific energy of fusion reactions is higher than anything else within our technological grasp.[12] Another useful fact about fusion reactions is the abundance of fusion fuels. Fusion burns hydrogen and helium – the two most abundant elements in the universe. It is estimated that, taken together, hydrogen and helium comprise about 99.999% of the mass of the universe. The amount of uranium

---

[12] Matter-antimatter reactions non-withstanding. In point of fact, antimatter can yield on the order of 10^{17} Joules per kilogram of energy. This is, of course, a separate matter entirely (pardon the pun).

and thorium can be measured in parts per trillion. On the surface of the Earth alone, there is sufficient deuterium to supply mankind's energy needs for the next million or so years. Known stores of uranium may take us to the next hundred. Cost is, of course, inversely proportional to abundance.

	Nuclear Light Bulb	Orion Project[13]	Daedalus Project[14]
Empty Mass:	32 metric tons	5,000,000 metric tons	3,000 metric tons
Vacuum Thrust:	409,000 N	2,500,000 N	7,00,000 N
Specific Impulse:	3200 seconds	50,000 seconds	100,000 seconds

Note that due to the combination of high average thrust and incredibly high specific impulse, these concepts are the best-developed ideas for sending humans to neighboring stars systems. Though the Orion concept warrants no further investigation for the purposes of interstellar travel – it's mass requirements are simply too high to be practical in this regard – the Daedalus propulsion system does need further exploration.

Very basic plasma physics

In order to understand what takes place in the fusion plasma, one must understand the properties of plasmas in general. Simply put, a plasma is state of matter in which the electrons are not bound to the nuclei of the atoms. All electrons and nuclei exist in an excited, gaseous state in which they can interact, but not recombine. Note that the existence of a plasma state does not necessarily mean that the gas is also ionized. Plasmas certainly can be ionized, but are often electrically neutral when considered in bulk.

Because plasmas have all of their charge carriers free, they are highly conductive, having an essentially zero resistance[15] for most purposes. Due in part to this high electrical conductivity, plasmas have a high heat capacity and can be considered to be good reflectors of radiant energy.

Plasmas are of course, extraordinarily hot. The fusion plasmas that we will be examining have temperatures in the tens or even hundreds of millions of degrees, Kelvin. This high temperature gives rise to the single largest problem in fusion research: containment. If a plasma comes into contact with something cold (and relative to a hundred-million degree plasma, everything else is cold), it sheds its energy instantly, becoming a normal gas.

And keeping the plasma from coming into contact with anything is a nearly impossible problem. Due to the high temperature of the gas, its particles have a high kinetic energy, and so are trying to fly away from each other as quickly as possible. If the particles manage to do this, it is a near certainty that they will hit something cold. The upshot of all of this is that, in the atmosphere, the lifetime of a hot plasma can be measured in microseconds.

One important fact worth mentioning at this point is that high temperature does not equate with high heat. This may seem a contradiction at first, but consider these cases:

i. A single neutron with a kinetic energy of 10 MeV (temperature of 10^{11} Kelvin)
ii. Gas inside a fluorescent light bulb at about twenty-thousand degrees Kelvin
iii. Cigar smoke with a temperature of six hundred degrees Celsius
iv. Boiling water with a temperature of one hundred degrees Celsius
v. Concrete on a hot summer day with a temperature of fifty degrees Celsius

---

[13] Figures are extrapolated from Freeman Dyson's paper, *Interstellar Transport*. This is an extreme “super-Orion” version of the concept and, in practice, the propulsion system alone would be the size of a small city.

[14] Figures from *A Technical Review of the Daedalus Propulsion Configuration,* by Richard Obousy

[15] Note that this is not the same as being superconducting. Plasmas do have a very small resistivity, though it is many orders of magnitude smaller than the resistivity of metals.

The neutron will not cause any significant temperature rise in anything that it may hit (though it will "cool off" instantly). The heat from cigar smoke or from a fluorescent tube can be felt, but with little, if any, harm. Boiling water can be handled for brief moments, but more than instantaneous exposure is very painful. And of course, everyone knows how painful hot concrete can be, especially when running from the beach back to the car.

In all of these cases, it can be seen that the absolute temperature of the object in question has little to do with how much heat it transfers. The amount of heat transferred is dependent on the heat capacity and density of the substance as much as it is the temperature.

The plasmas that we will be discussing will have very high temperatures and good heat capacities, but very low densities. In fact, were it not for the fact that they have unbound charge carriers, we would call them vacuums rather than plasmas. Typically, we will be discussing systems containing about 10^{13} particles per cubic centimeter. Contrast this with the density of the atmosphere (at sea level and room temperature), 2.5×10^{19} particles per cubic centimeter. Our plasmas can then be said to be at about one millionth of an atmosphere in pressure.

So then, we can think of a plasma as a very hot, rarefied gas. Like any other rarefied gas, it can be described as either a collection of charged particles, all trying to get away from each other, or as a continuum, exhibiting a form of collective behavior.

Nuclear fusion propulsion

Nuclear fusion reactions, as stated, require a combination of high-temperatures and high pressures in order to ignite. This combination of temperature and pressure must meet what is know as the “Lawson Criteria”. The specific Lawson number is different for each reactant, but in general terms, the pressure and temperature must be high enough that the nuclei in the plasma are able to overcome their electrostatic repulsion and combine. Again in general terms, this equates to a temperature of around ten million degrees and a particle density of about twenty times that of lead. Just for sake of illustration, for common reactions, the Lawson criteria is[16]:

$n\tau = 10^{14}\ \text{cm}^{-3}\bullet\text{sec}$ for deuterium-tritium
$n\tau = 10^{16}\ \text{cm}^{-3}\bullet\text{sec}$ for deuterium-deuterium

The temperatures are generated by any number of means: laser irradiation, radio-frequency or microwave heating, or direct electromagnetic acceleration of the plasma. The pressure is achieved by successfully confining the plasma in a small volume. This is the difficult part.

Gravitational confinement

In a star (nature's own fusion reactor), confinement is achieved by gravity. Though there are many mysteries associated with star formation, astronomers believe that they know the basics. As a cloud of gas (mostly hydrogen) collapses under its own gravity, its volume decreases and hence its density increases. As you may recall, the ideal gas law states:

$$PV = nRT$$

The product of the partial pressure and volume of an ideal gas is proportional to the product of the temperature and number density of that gas. In our collapsing cloud, the volume is decreasing and density and pressure are increasing. The temperature must then increase as well. When the temperature gets above the fusion ignition temperature of the gas, the newly formed star ignites. This is often referred to as "frictional heating", due to the fact that it arises as a result of the gas molecules bumping into each other. Obviously, this only works on stellar scales. Gravitational confinement is not an option for a man-made reactor.

---

[16] Francis Chen, *Introduction to Plasma Physics*

On Earth, we've a number of options for confinement. The first fusion reactor was a hydrogen bomb. In this case, the fusion fuel is confined by the inward radiation-pressure of an exploding fission bomb. Though highly efficient, it only works once. This is the Orion concept and will not be further explored.

Magnetic confinement

One of the most researched confinement schemes is magnetic confinement (MCF). Examples of these reactors are the tokamaks that are being built in France and Germany. The idea behind magnetic confinement fusion is that the plasma is contained by intense magnetic fields. Though this sounds as though it would be foolproof, the reality is that plasmas are extraordinarily difficult to contain and strong magnetic fields are difficult to create and control. Still, this concept is worth further exploration.

And now for a bit of theory: The force that a particle feels in a magnetic field is equal to the cross product of the current and the field and is therefore perpendicular to both the field direction and the charged particle's motion. This is unlike the force from an electric field, as that force is independent of the movement of the charge and is in the same direction as the field. The magnetic field then, attempts to pull a charged particle in motion into an orbit.

So this then is the problem. Ions in a system may be going more or less the same way and so feel more or less the same force from the applied magnetic field. But, if the ions are hot, their random thermal motion is just another added on to their velocity vector. This then adds a small random term to the force that they feel from the applied magnetic field. So then, unlike the situation with an electric field, the magnetic field does not remove random motions from a collection of particles. Rather, it imparts random forces to those particles in response to their random motion. The overall effect of all of this randomness is that the plasma, though nice and orderly, will tend to squeeze out of a magnetic field. In fact, it is from this oozing motion that plasmas get their name.

How do we combat this? There are a few options. We can use more orderly plasmas, with a lot of motion in a single direction; we can use electric fields as well as magnetic; or we can use lots of magnetic fields and hope for the best.

If we start by giving the plasma a net momentum in a single direction, then the applied magnetic field will induce a net force. This net force, being constant rather than random, will effectively overwhelm the net force produced by random thermal motions. Note that the net force resulting from thermal motion will still exist and the plasma will still try to leak out of the containing magnetic field because of it. But every particle will also feel that constant force and so leakage will be controlled.

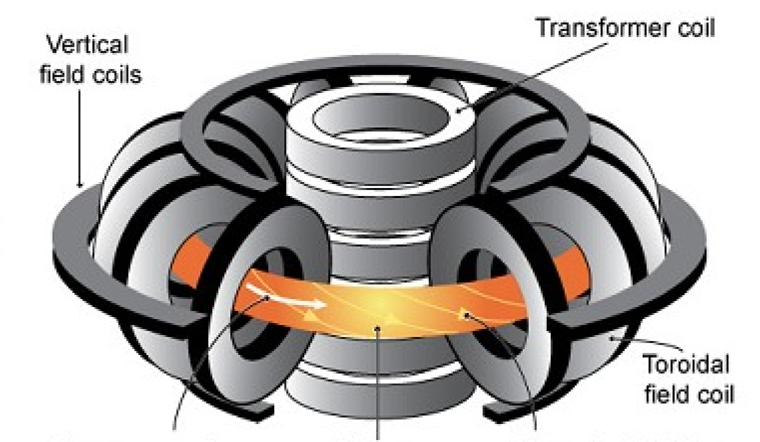

Tokamak Reactor Schematic
(image courtesy of General Fusion Inc, ©2010)

Using multiple magnetic fields is tactic used to achieve a measure of control over the plasma. One of the more popular fusion reactor geometries is the tokamak. This is a doughnut-shaped chamber that makes use of two sets of magnets. The first set applies a magnetic field in the z-direction, normal to the plane of the doughnut. This causes a net circulation of plasma around the interior.

After giving the plasma a net motion, another field is added, this time from magnets coaxial to the doughnut. This produces a motion in the plasma that is normal to the walls of the chamber and toward the center. This is the field that we don't want leaking, since it keeps the plasma from hitting the walls of the chamber. Excess leaking from this confinement field will cause too much heat to be transferred from the plasma to the chamber walls. This will have little effect on the chamber, but it will cool down the plasma, quenching fusion reactions before they have a chance to start.

Magnetic confinement comes in a variety of flavors, ordered by their geometry. As mentioned, a tokamak reactor confines the plasma within a doughnut-shaped series of superconducting coils. A sphereomak reactor (and its many derivatives) is similar in geometry but for the fact that the doughnut is deformed in to more of a spherical shape.

Field-reversed configuration is a linear geometry. The plasma is confined to a linear tube with superconducting coils banding it. At the ends of the tube are two additional, powerful magnets that confine the plasma within. This concept is particularly suited to propulsion, as one of the "mirror" magnets can be made "leaky" such that some of the fusion plasma escapes. The escaping plasma can then be used for thrust and/or power generation. The other mirror magnet can also be made a bit leaky in order to inject more fuel in to the system.

The VASIMIR concept, as researched by Ad Astra in Houston, is a practical embodiment of the field-reversed configuration with the caveat that fusion need not occur. Rather, it uses the confinement to increase the efficiency of the radio-frequency heating. The plasma is contained until it is extremely hot and then some is vented, producing thrust. With higher containment fields, and a lot more heating, this could conceivably become a fusion rocket.

Inertial electrostatic confinement

Another means of confining a plasma is via electric fields. This is embodied in the inertial electrostatic confinement fusion (IEC) concept. First pioneered by Dr. Philo Farnsworth (inventor of color television) in the 1950, this uses a series of electrically-charged grids to confine the plasma. Particle guns fire charged nuclei toward the center of the reactor, passing through several of these grids. These particles are then trapped in the center and compressed by the inertia of the beams themselves.

The "polywell" concept combines IEC and MCF by supplementing the electrostatic grids with strategically-placed magnetic fields. This concept is also well suited to the needs of propulsion as the fields and grids can be made "leaky" such that a plasma beam is ejected in a controlled fashion.

Of course, there are issues with IEC-type devices. First and foremost, is space charge. In a vacuum tube, when ions pass through a grid, most are accelerated and then confined (which is what we want), but some are deflected away and others are captured adding to the space charge on the grid. This "space charge" is simply an anomalous charging of the grid resulting from ions passing by or striking it. This is undesirable, as an excess of space charge will tend to deflect the ions coming from the particle guns.

Unfortunately, the more current in the ion stream, the more of a problem this is. To overcome this problem, we need a "virtual" grid that can act as though it carries charge but doesn't physically exist in the path of the ion stream. Some success has been gained by combining

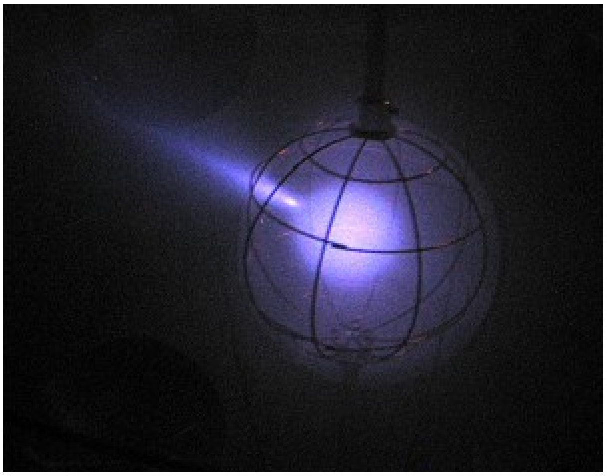

Photograph of an inertial electrostatic confinement fusion reactor operating in its "jet" mode in order to release a portion of the fusion plasma

(photo courtesy of MIT Space Systems Lab, 2010) electrostatic and magnetic confinement as with the polywell concept.

The second problem with this approach is heating of the grid. The plasma temperature within the central grid can be several hundred million degrees Kelvin. And of course, this plasma is trying its best to get past the grid and escape. This may seem like a high temperature, but bear in mind that this is in a near vacuum. The resultant heating of the grid is only a few thousand degrees kelvin. This heat must still be dealt with and some possible schemes include water or gas cooled grids.

Another problem is damage to the grid from ions and neutrons. Though as mentioned, the grid picks up only a few thousand degrees of heat, this comes as a result of high speed ion impacts. By way of analogy, we can apply a hundred or so pounds of force to a brick wall by leaning against it. This is a nice, gentle pressure that does no damage to the wall. We can also apply a hundred or so pounds of force to the same wall by firing machine gun rounds at it. Though the pressure is identical, the latter is caused by a few, highly energetic events, and the wall will eventually be destroyed by the impacts. Currently, grids are made of tungsten or molybdenum, both of which can take a lot of punishment. Even so, a better solution is needed, particularly as the ion current and fusion power is increased.

Inertial confinement

Inertial confinement (ICF) does away with magnets and grids entirely. In this scheme, pellets of fusion fuel are bombarded with lasers or electron guns. These then simultaneously compress and heat the pellet to the point at which the Lawson criteria is met and fusion occurs. This is the engine configuration of the Daedalus and Icarus projects.

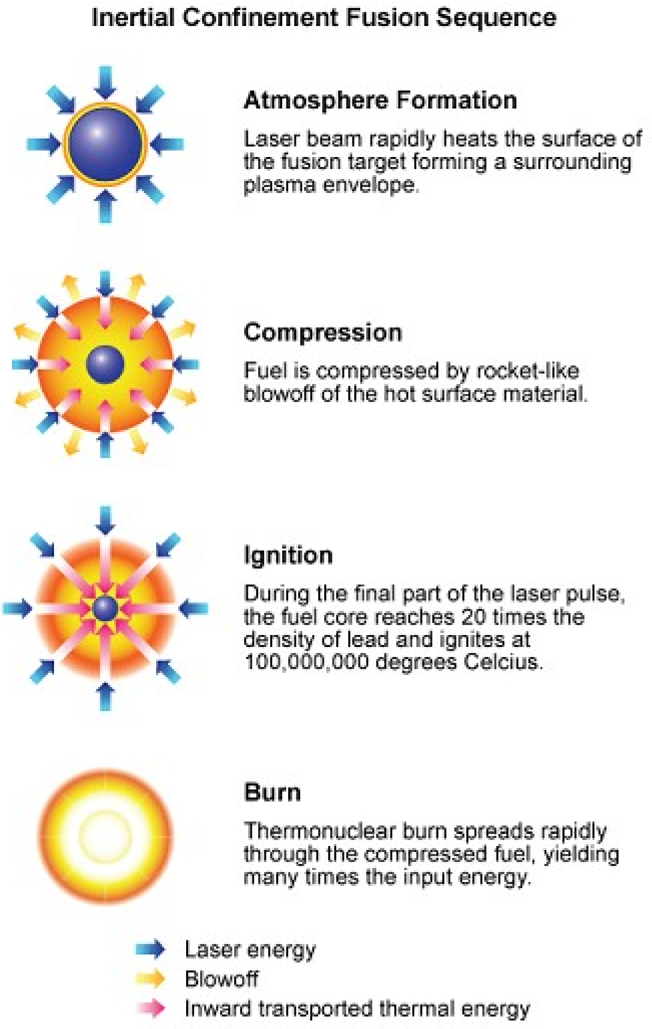

Laser beam rapidly heats the surface of the fusion target forming a surrounding plasma envelope.

Fuel is compressed by rocket-like blowoff of the hot surface material.

During the final part of the laser pulse, the fuel core reaches 20 times the density of lead and ignites at 100,000,000 degrees Celcius.

Thermonuclear burn spreads rapidly through the compressed fuel, yielding many times the input energy.

(image courtesy of General Fusion Inc, ©2010)

Though it may seem odd that lasers can compress a pellet of fuel, there is real precedent for this. A hydrogen bomb works, not by compressing the fusion fuel with the mechanical shock of the exploding fission trigger, but with x-rays emitted by the fission trigger. In much the same way, in the ICF reactor, the fuel pellets are contained within gold-plated, glass envelopes. When struck by lasers or electron beams, the gold vaporizes creating x-rays which in turn compress the fuel.

When enough fusion reactions have occurred such that their energy counterbalances the x-ray's energy, the plasma stops its inward compression and expands outward. By the time expansion begins, most of the fuel has been burned in fusion reactions. All of this occurs in a time scale so fast that it cannot be accurately measured. As expansion continues, the plasma cools and the energy can be collected for use.

As an interesting historical note, it was thought that this could not be a method of heating and compressing a plasma. Compressing something using x-rays seemed very counter-intuitive to most physicists due to a well-known phenomenon called the "Rayleigh-Taylor instability." This is just a way of saying that a light fluid cannot support a dense fluid against its pressure. If you pour water on top of oil, "fingers" of water seep into the oil displacing it, until all of the water is on the bottom and the oil on top.

Similarly, it was thought that x-rays and gas would be like fluids of different densities. In other words, that x-rays would displace the gas rather than compress it. It wasn't until the hydrogen bomb development projects of the 1950's that it was found out to be otherwise. Hydrogen bombs make use of this x-ray compression, this time from a fission trigger, in order to ignite the fusion reaction. Though the Rayleigh-Taylor instability does play a role in plasma physics, it does not affect x-ray plasma compression in the slightest.

As mentioned, both the Daedalus and Icarus concepts make use of ICF engines. In the former, a ring of lasers ignites fuel pellets as they are ejected out of the spacecraft. In the latter, intense electron beams do the dirty work. In both cases, pellets are ejected at a rate of about 250 per second. This high rate of fire imparts massive amounts of thrust. The high temperature of the fusion plasma equates to a very high specific impulse, and the high firing rate means that there is no need for massive pusher plates and shock absorbers as with the Orion concept.

Both of these suffer from a significant drawback in that it is very difficult to extract useful, electrical energy from the fusion reactions. Since the reactions take place well away from the ship itself, the plasma cannot be intercepted in order to generate power. Both of these need an additional power source (probably a fission reactor) in order to power the ship and drive the ignition sources.

Radiation-pressure confinement

As x-rays are compressing the plasma in inertial confinement fusion, could it then be said that it is radiation pressure that actually confines the plasma? As it turns out: not really.

Radiation pressure compresses the plasma, but it does not confine it against its own force of expansion. Once fusion reactions begin to occur, the outward pressure from those reactions quickly overwhelms the inward pressure of the x-rays. This is why we must call it inertial confinement. Would it be possible to use radiation pressure to confine a plasma?

The answer is a reserved "yes". Though it is true that intense x-rays or other forms of radiant energy can transfer momentum to a plasma, the intensity required is fantastic. To illustrate, consider a single photon having a momentum of:

$$p = h / \lambda$$

In which the momentum is given by Planck's constant divided by the wavelength of the photon. This photon must counteract the kinetic energy of a single nucleus, given as:

$$K_{NUCLEUS} = kT$$

Whereas the kinetic energy is equal to the temperature (in Kelvin) times Boltzmann's constant. So we have the relation:

$$K_{NUCLEUS}^2 = c^2 p_{PHOTON}^2$$

Or:

$$\lambda = hc / kT$$

For an aneutronic (one that does not produce neutrons) fusion reaction, the thermal energy must be at least about 35 keV on the average, with some particles having a thermal energy as high as 50 keV. This corresponds to a wavelength of about 0.25 angstroms (one ten-billionth of a meter), or well into the gamma ray part of the spectrum. Since no one has yet invented the gamma ray equivalent to a laser, we must use lower energy photons with a higher photon flux.

There do exist lasers with wavelengths around 1500 angstroms. This corresponds to an energy of about 8 eV (pretty weak, in terms of nuclear scales). Were five thousand such photons incident on the same nucleus, we could push it around. A single nucleus has a radius of less than the photon's wavelength, so the size of the spot to which we can focus is limited by diffraction. This is approximately equal to two-thirds of the wavelength or about 1000 angstroms in our case. Since a spot of 1000 angstroms in diameter can contain tens of thousands of nuclei in cross section at the density of solid matter, we must multiply our irradiance figure by at least that amount in order to hit most of them with most of our photons.

At this spot then, we must bring to focus about fifty million photons. This must then be done over the entire surface area of our target. If our target is one centimeter in diameter, this equals 40 billion spots, each 1000 angstroms in diameter. To cover the surface of our fuel pellet with photons is just a start. We must keep doing it, so that the plasma is contained for a sufficient length of time so that most of the fusion fuel is burned.

If we use pulsed lasers to do this, then each time they fire they project 50 millijoules of energy on to the target. Not much energy, right? But, this must occur at least as fast as the ion transit time from one side of the fuel pellet to the other. Given our same 50 keV ions and our pellet size of one centimeter, our ions are traveling at about 2 million meters per second and so will make it from one side of the pellet to the other on the order of a few nanoseconds. So we have to fire our lasers at least one billion times per second, giving a power on target of about 50 million watts.

This is of course, a "back of the envelope" calculation, but our answer's order of magnitude is correct. Even if we estimated high by two orders of magnitude, this still means that we need megawatt lasers in the far ultraviolet to soft x-ray region of the spectrum. So the bottom line is this: Radiation confinement is a nice idea, but technologically unfeasible at this time.

Speculative fusion concepts

There are a few, highly advanced fusion systems worth mentioning at this point. So long as we are speculating, why not consider some "alien technology"? By this I mean, technology that would exist if there were no technological limits, only scientific ones.

As it turns out, nature provides us with an example of a nearly ideal fusion system: ball lightning. We know very little about ball lightning, but from what we can infer, it is possible maintain a stable, high-density plasma in the atmosphere for several minutes or better. Once we are able to duplicate this ball lightning effect, small, high-power fusion reactors operating within the atmosphere may become possible.

Another hot area of research is in the use of femtosecond lasers[17] to ignite self-sustaining reactions. A femtosecond laser can produce, at a point, temperatures hot enough to ignite an ICF reaction. This happens but for an instant, and then the energy released by the fusion event drives the plasma, igniting other fusion events. This is sometimes referred to as "laser catalyzed fusion", because the laser gets the reaction going which then continues as a chain-reaction. This will become a viable means of fusion once the reaction can be contained such that the energy from a reaction can be absorbed by the plasma rather than lost by escaping particles.

Focus and pinch devices

What we really want is a single device that can operate as efficiently as an inertial confinement system and with the control of a magnetic confinement system. Ideally, this system should have the versatility to make use of a wide variety of fuels, be easy to manufacture and maintain, and easily scale to high energies. Such systems exist in the various pinch and focus devices. Our discussions will be limited to one particular implementation of this: the dense plasma focus (DPF) concept.

---

[17] This being a laser whose pulse only lasts for about a femtosecond, or one quadrillionth of one second.

Dense plasma focus (and indeed, all pinch and focus) devices rely on the fact that an electrical current induces a magnetic field. Though they operate in pulsed mode (similar to ICF devices) they have more in common with magnetically confined systems. A basic pinch device consists of a set of coaxial electrodes, between which, fuel is injected and an intense electric arc is struck. Their operation can be divided into several stages.

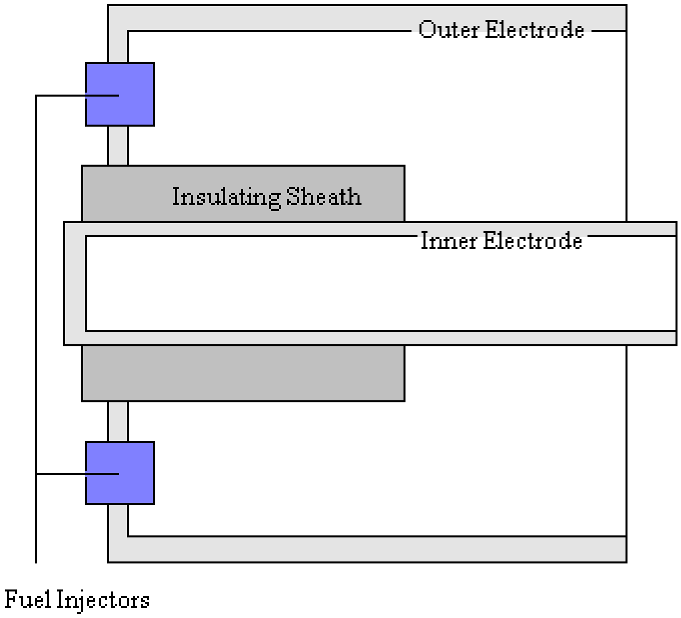

Dense plasma focus (DPF) device cut-away view

The first stage is termed “injection”. At the beginning of operation, the fusion fuel is injected in to one end of the electrodes. This is a high-pressure pulse of gas, consisting of the proper mixture of elements that we wish to fuse (duterium-tritium, for instance).

At the “breakdown” stage, the device is hit with a voltage spike between the anode and cathode. In past experiments, values of between ten thousand and one million volts have been used. This is sufficient current to break down the gas, converting it to a plasma, and inducing a current of up to ten million amps. In total, this corresponds to an energy of tens of mega-joules put to the electrodes.

During the “rundown” stage, a high current induces an enormous magnetic field, accelerating the plasma to an extremely high velocity (around Mach-50) in a few tens of nanoseconds. As the plasma travels down the barrel, it gathers and heats the gas ahead of it (this is termed, the "snowplow effect"), increasing the density of the resultant pulse.

The complicated part comes during the “pinch” phase. At this point, the electric arc separates from the ends of the coaxial electrodes. The magnetic field collapses, trapping the gas in a shrinking bundle of electric and magnetic field lines. This has the effect of compressing the plasma to the point of fulfilling the Lawson criteria and inducing a fusion reaction. Like an ICF device, this plasma “pellet” undergoes fusion and is ejected out the back at high velocity, producing thrust. But unlike ICF, it is not x-rays that compress the plasma, but the collapsing magnetic field.

Stage one: gas injection

In prior DPF experiments, a mixture of deuterium and tritium has been used due to its low ignition temperature and because the 14 MeV neutrons produced in the reaction are a good indication of the rate of fusion. In most experiments, the gas has not been injected into the device. Rather, the chamber is filled with the gas mixture at a partial pressure of between one millitorr and one torr[18].

A few exceptions exist, in which the mixture is injected into the chamber with each pulse. One notable experiment, the MARAUDER apparatus at Kirtland Air Force Base[19], uses high-speed valves to inject the gas into the vacuum with each pulse of the capacitor bank. These valves are designed to inject a gas soliton (a “clump” of gas whose structure is similar to that of a smoke ring) at the moment of breakdown. On the same apparatus, a post-pinch injector has been added to inject a neutral gas into the plasma focus. Both sets of injectors have been used repeatedly with good results in terms of gas efficiency and jet power out.

---

[18] These are measurements of pressure. One standard atmosphere is equal to about 760 torr.

[19] J.S. Brzosko, et. al., "Comments on the Feasibility of Achieving Scientific Break-Even with a Plasma Focus Machine”

Ideally, the injector valves should operate fast enough such that all of the gas that they emit is used in the initial breakdown, or caught in the rundown phase. This is not a difficult requirement, as gas injected, even at a high velocity, will still likely be caught by the arc as it travels down the electrodes. As a typical electrode length is fifty centimeters, with a rundown time of less than one hundred nanoseconds, the arc can catch up to any gas released within a millisecond of discharge and traveling at less than about ten thousand meters per second.

For a real-world fusion rocket, injector performance is critical because of the need for fuel conservation and because it is impossible to simply keep the electrode cavity filled with fuel. Aside from timing considerations, the only issue is that the injectors be non-metallic in order to simplify the electromagnetic design of the apparatus. Along these lines, a research group at Ohio State University[20] has designed high-speed injector valves of machined acrylic. These look very promising, at least in simulation.

Stage two: breakdown and the power systems

Dense plasma focus devices require a large amount of electrical current to be delivered across a small gap in a very short amount of time. From the point of view of the power system, this is a very demanding application. Most power supplies are designed to deliver lower amounts of current for extended periods, rather than short, high-power pulses.

Pretty much the only thing that can deliver such a pulse of current is a large capacitor bank. Capacitors store energy in the space charge between two metal plates, which are kept separated by an insulator. They are charged by a continuous power-supply, and then rapidly discharged across the gap of the DPF device, dumping all of the energy built-up during their charging. The capacity of the bank is measured in “Farads”, and the energy contained in such a bank is given by:

$$E = CV^2 / 2$$

Whereas:

E = energy in Joules
C = capacitance in Farads
V = charging voltage in Volts

Because the discharge time of a capacitor is directly proportional to the capacitance, we generally want the capacitance to be low and the charging voltage to be high. Current terrestrial experiments typically use a capacitance of 100 micro-farads and a charging voltage of 250 thousand volts, yielding a stored energy of about three million joules. For a spacecraft propulsion system, the capacitance will probably be about the same, but with a much higher charging voltage, perhaps in the tens of millions of volts.

This capacitor bank adds a lot of weight to the system, but at this stage of our technical understanding, there is really nothing else that can substitute for them. Capacitor banks have a few advantages in this application: they have no moving parts and they last for tens of millions of charge-discharge cycles. Were we to find some other means of storing and rapidly discharging energy, it would need to have those characteristics as well.

As soon as the current spike from the capacitor bank hits the electrodes, breakdown of the injected gas occurs along the cylindrical insulating sheath. The gas is immediately heated to a plasma as the current (essentially a bolt of lightning) travels from the outer to the inner conductor. As one can well imagine, at this stage, both the electrodes and the insulating sheath take a lot of abuse. This is one of the main points of failure in the DPF device and a lot of research is being conducted in to how to mitigate the damage.

Efforts to reduce this damage have met with some success. Notably, the previously mentioned MARAUDER apparatus at Kirtland Air Force Base uses specially shaped electrodes in place of the insulating sheath. This approach has been successful at eliminating this point of failure, though more research is needed so that the

---

[20] Pavlos Mikellides, et. al., "Design of a Fusion Propulsion System -- Part I: Gigawatt-Level Magnetoplasmadynamic Source*”*

efficiency of the device is not impacted.

Also of some use has been a magnetic coil placed within the device near the point of breakdown. This creates an axial magnetic field that protects the insulator from the breakdown current by somewhat diffusing the discharge plasma. This does, however, impact the device efficiency as well. Both the MARAUDER apparatus and the PACO device at the Universidad Nacional del Centro in Buenos Aires have made use of this technique.

Stage three: rundown

This is perhaps the most important, and the most boring, phase of DPF operation. It is during this phase that plasma sheath acceleration occurs. By virtue of the induced magnetic force, the plasma sheath lifts off of the insulator and begins its run down the electrodes away from the gas injectors. This current sheath has a velocity on the order of hundreds of kilometers per second as it travels down the barrel.

As it moves into unheated gas, it creates a shock wave that gathers and heats the gas ahead of it. This shock front is a fraction of a millimeter in thickness and travels with a Mach number near one hundred. Gas atoms pass through this shock wave within one nanosecond, creating a plasma layer with a temperature of hundreds of thousands of degrees.

Stage four: pinch and fusion

When the plasma sheath arrives at the end of the electrodes, it breaks away from them, and its direction of travel changes from axial (along the electrodes) to radially inward. This inward force is governed by the same magnetic field equations, except for the facts that the magnetic field is now dependent on the distance from the center-line, and that the current is traveling radially as well as axially. As the plasma sheath collapses toward the center, it has a velocity in the millions of meters per second.

The collision of everything at the focus causes temperatures in the tens of millions of degrees and pressures approaching the particle density of a solid. At this time, the magnetic field tries to collapse but “kinks up” as the plasma gets in the way. The plasma then holds at this point for up to about several nanoseconds, due to the magnetic pressure outside of the focus. It is at this time that the Lawson criteria are fulfilled and the fusion reaction initiates.

When the fusion reactions occur, they destabilize the plasma focus leading to decompression. At decompression, neutral gas can be injected into the focus in order to carry off the heat of the nuclear reaction and increase the thrust of the engine at the expense of specific impulse. Following focus decompression, the plasma temperature drops back down to a few hundred thousand degrees. If desired, the plasma can be passed through an MHD generator in order to extract electrical power.

At the point of fusion, the energy from the fusion event itself is up to five-hundred times the amount of energy that went in to initiate the reaction. In the case of a DPF propulsion system, all of this energy is liberated as heat to drive the exhaust.

DPF devices for propulsion

Ultimately, this comes down to a question of engine thrust and efficiency. Various studies indicate that a DPF-based propulsion system could be scaled up to the point where it would have a similar specific impulse and thrust to the Daedalus engine: 100,000 seconds and 7 million Newtons, respectively. Though a single DPF thruster could probably not be scaled to this level, there is nothing wrong with using a number of them. This then leads to the question as to why we should pursue a DPF thruster rather than the more-studied Daedalus system? There are a number of reasons that DPF thrusters are more attractive from an engineering standpoint.

First, there is the issue of fuel efficiency. Even though the Daedalus and DPF engines have similar figures for specific impulse, the DPF engine wins this one. The reason for this is that there is a hidden overhead built in to the design of the Daedalus engine. In ICF-based engines, the fuel is contained within a “fuel pellet” of gold-plated glass. These pellets are approximately one centimeter in diameter and contain a mixture of gasses to serve as the fuel for the fusion reaction. The pellets themselves contribute next to nothing to the output power of the engine: they are only needed to serve as the trigger to ignite the reaction. Unfortunately, the mass of the pellet accounts for better than half of the mass of the fuel, all while making essentially no contribution.

A DPF device makes use of fuel injectors to deliver the gas directly to the point of ignition without the need for such a pellet or its associated overhead. Once we eliminate the pellet, we also eliminate fuel loss due to pellet breakage and defects as well as all of the moving parts associated with storing, transporting, and ejecting the pellets. This more than doubles our effective fuel efficiency and eliminates a lot of the mechanical complexity of the system. In effect, our figure for specific impulse just doubled.

Second, the engine design is far simpler. Unlike the Daedalus system, there are no lasers or electron beams and their associated optics that are required for ignition. This eliminates many layers of complexity and many potential points of failure. High-powered lasers operating continuously for years at a time would almost certainly require re-alignment and suffer from failures in their optics, crystals or drive systems.

A DPF device requires just two things: a power plant and a capacitor bank. Capacitors have been around for hundreds of years and, by this time, are absolutely reliable in their operation. They have no moving parts, require little in the way of cooling, never have alignment issues, and have been shown to operate continuously for tens of millions of cycles. Again, this is another level of complexity that can be completely eliminated from the system.

Finally, the design of the DPF device itself is simple. A dense-plasma focus device consists of metal rods, ceramic and quartz insulators and... that's about it. In point of fact, with just a few thousand-dollars in equipment, you could build a small DPF thruster in your garage. Indeed, there is a fledgeling home-brew community that is actively engaged in researching this technology. There exist a number of instances of small, home-built DPF devices that have ignited fusion reactions. No ICF device can say the same.

So let us go back to our simulation and hypothesize a device with the following parameters:

	Nuclear Light Bulb	Daedalus Project	DPF Engine
Empty Mass:	32 metric tons	3,000 metric tons	3,000 metric tons
Vacuum Thrust:	409,000 N	7,000,000 N	7,000,000 N
Specific Impulse:	3200 seconds	100,000 seconds	200,000 seconds

These figures represent a natural extrapolation of known technology and such a system could be built without requiring much in the way of technological breakthroughs. Again, for the sake of our simulation, we'll return to our original numbers for the spacecraft, modified for the new engines:

Dry Mass (payload plus engines):	183,000 metric tons
Propellant Mass:	183,000 metric tons
Propellant Delivery and Containment	183 metric tons

Again, the Daedalus engine has more overhead than the DPF-based engine, but this will be accounted for by modifying their specific impulse figures. First, the Daedalus engine:

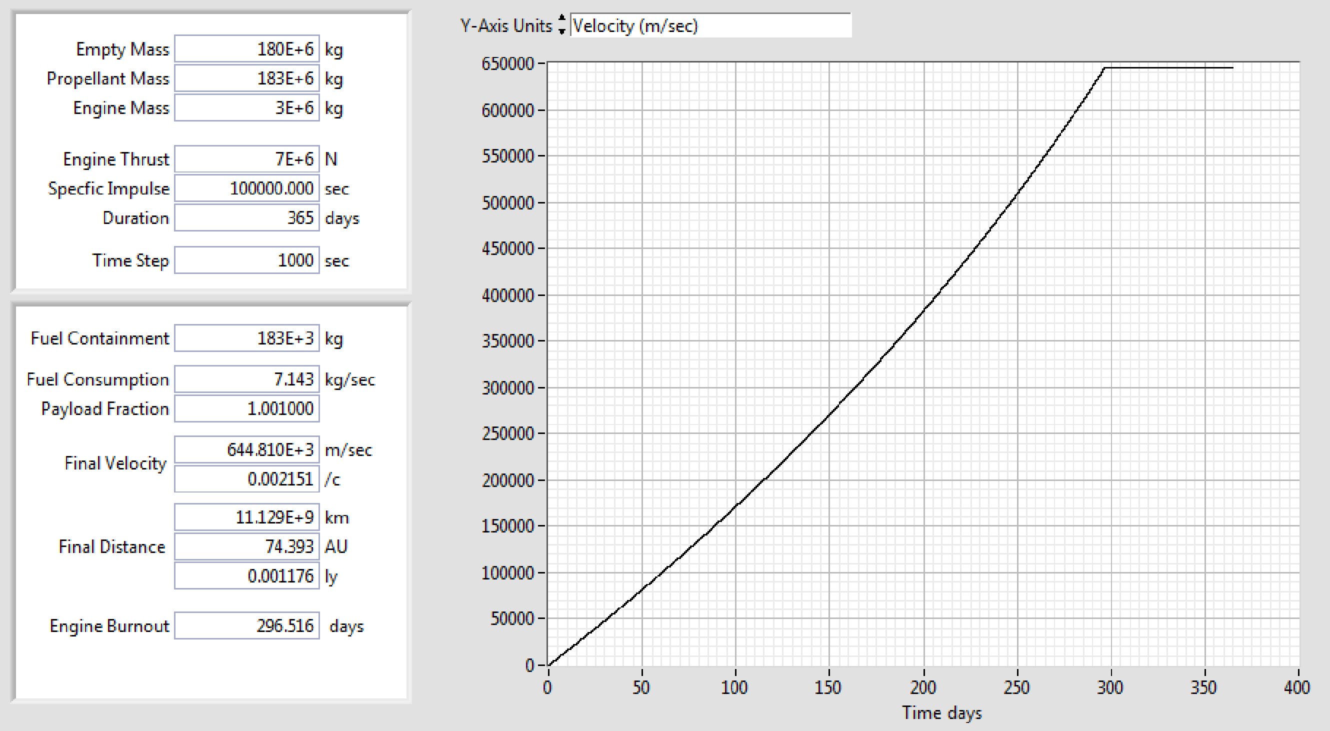

Empty Mass 180E+6 kg
Propellant Mass 183E+6 kg
Engine Mass 3E+6 kg

Engine Thrust 7E+6 N
Specfic Impulse 100000.000 sec
Duration 365 days

Time Step 1000 sec

Fuel Containment 183E+3 kg

Fuel Consumption 7.143 kg/sec
Payload Fraction 1.001000

Final Velocity 644.810E+3 m/sec
0.002151 /c

Final Distance 11.129E+9 km
74.393 AU
0.001176 ly

Engine Burnout 296.516 days

Y-Axis Units Velocity (m/sec)

Figure 4. Analysis of a Daedalus-class propulsion system

Note that this is far better than any chemical thruster and even beats-out the Orion system. After about one year, we would reach a cruising-speed of about 0.2% the speed of light. Were we to decrease the payload and increase the fuel, 2% of light speed would be within our reach. Now let us perform the same analysis for the DPF-based engine:

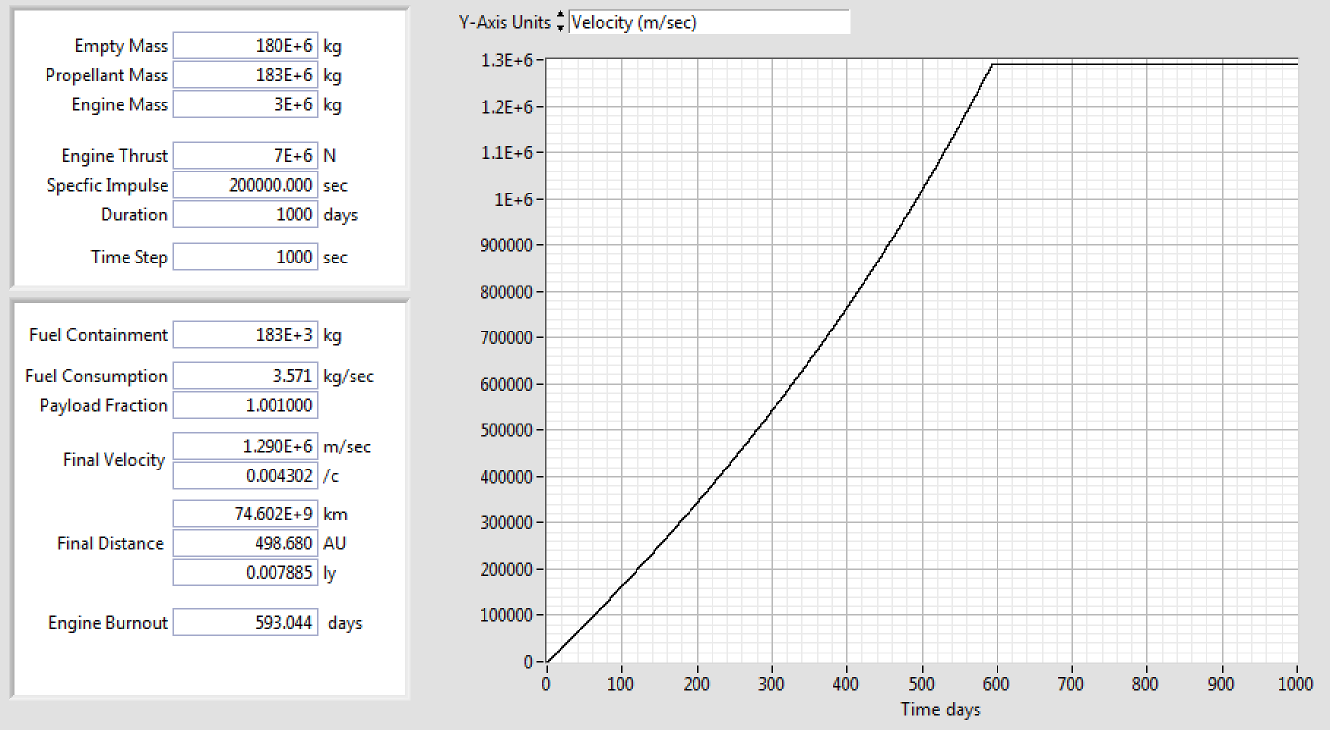

Empty Mass 180E+6 kg
Propellant Mass 183E+6 kg
Engine Mass 3E+6 kg

Engine Thrust 7E+6 N
Specfic Impulse 200000.000 sec
Duration 1000 days

Time Step 1000 sec

Fuel Containment 183E+3 kg

Fuel Consumption 3.571 kg/sec
Payload Fraction 1.001000

Final Velocity 1.290E+6 m/sec
0.004302 /c

Final Distance 74.602E+9 km
498.680 AU
0.007885 ly

Engine Burnout 593.044 days

Y-Axis Units Velocity (m/sec)

Figure 5. Analysis of a dense plasma focus-based propulsion system

As you probably could have imagined, if we double the effective specific impulse, we double the firing duration and approximately double the cruising-speed. Now our baseline is up to nearly 0.5% the speed of light. Again, were we to decrease the payload and increase the fuel, 5% of light speed would be within our reach.

But this still falls short of our 10% light-speed target velocity. Though I will not discuss such technologies in this paper, there are some ideas that we can pursue in order to increase our velocity. One interesting area of research is to use antimatter to catalyze a fusion reaction. This could augment our existing DPF device and could enable specific impulse figures that approach 500,000 seconds. Of course, generating (or gathering) and storing antimatter are problematic at best.

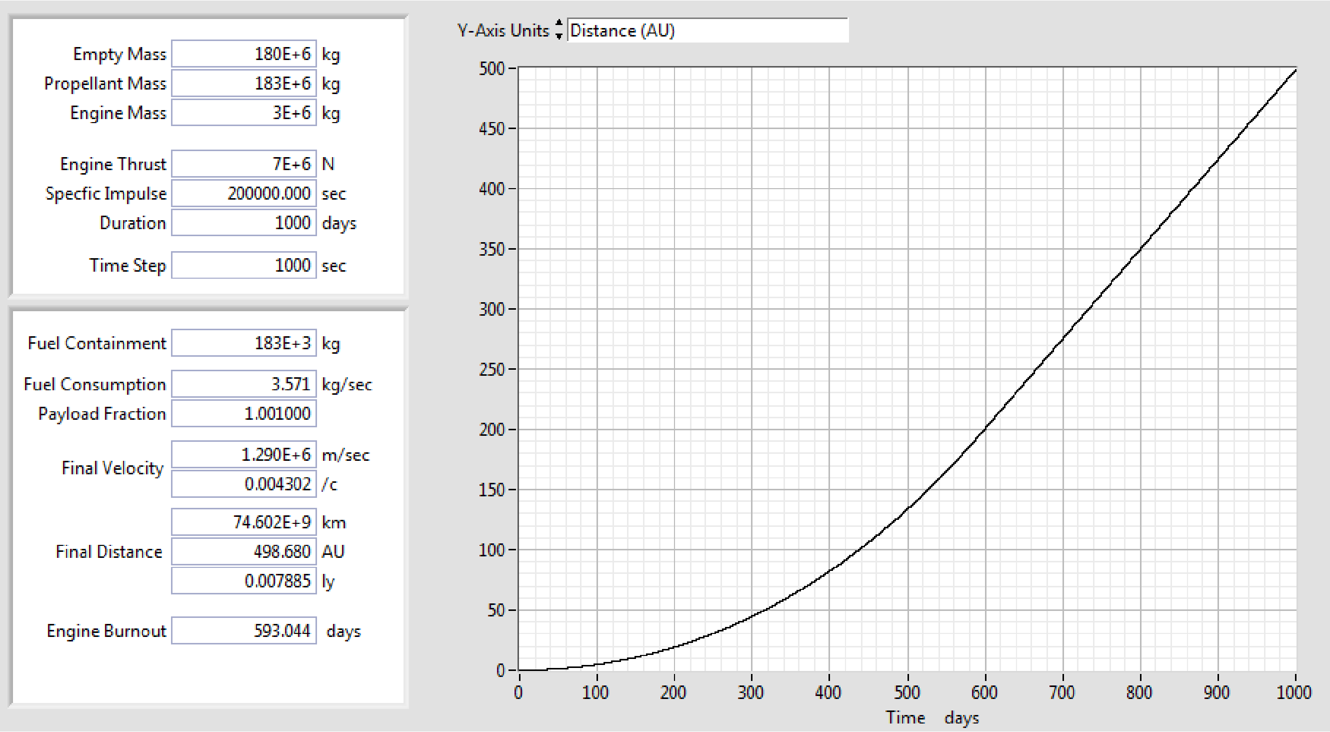

Empty Mass 180E+6 kg
Propellant Mass 183E+6 kg
Engine Mass 3E+6 kg

Engine Thrust 7E+6 N
Specfic Impulse 200000.000 sec
Duration 1000 days

Time Step 1000 sec

Fuel Containment 183E+3 kg

Fuel Consumption 3.571 kg/sec
Payload Fraction 1.001000

Final Velocity 1.290E+6 m/sec
0.004302 /c

Final Distance 74.602E+9 km
498.680 AU
0.007885 ly

Engine Burnout 593.044 days

Y-Axis Units Distance (AU)

Figure 6. Graph of distance traveled using a DPF-based propulsion system

We could also increase acceleration during the early phase of the trip by augmenting our propulsion system with solar or magnetic sails. Looking at the graph of distance vs. time, you will notice that much of the first year is spent within 50 astronomical units of our Sun. At this distance, a solar or magnetic sail could be deployed to take advantage of that fact and to radically increase the acceleration of the craft. In point of fact, such a sail would likely be a part of the design anyway, in order to allow for breaking at the target system without having to expend (and thus, carry) additional fuel.

Conclusions

As you can see, a dense plasma focus-based propulsion system has a number of advantages over an ICF-based system and very few disadvantages. Currently, the only mark against it is the difficulty in creating large capacitor banks for terrestrial experimentation. Though my simulations are not nearly rigorous enough to take in to account all possible variables, I believe that they give hints as to the potential of DPF-based propulsion. I hope that, in the coming years, I can further refine my simulations in order to come to a more-complete understanding of this system. And I at least hope that we don't become myopic in our visions, considering only systems that have been well-researched while neglecting other, potentially viable solutions.

About the author

Jim Cavera is currently the senior systems engineer for AIM-USA, a small, German-based avionics manufacturer. For the past twenty years, he has worked in the field of instrumentation, avionics, and test systems. He has experience at various NASA facilities as an independent contractor, designing equipment for robotic and manned space missions. Currently, he is a senior member of the American Institute for Aeronautics and Astronautics and, for the last two years, has sat on their Nuclear and Future Flight Technical Committee.

His educational background includes degrees in physics and optical engineering with additional graduate work in both nuclear and aerospace engineering at Purdue University. While there, his master's work was an exploration of dense-plasma focus devices for manned space missions.

Bibliography

B.P. Bromley, J. DeMora, B. Casey, P. Humm, E. Kirstein, A. Siddique, J. Javedani, G.H. Miley, S. DelMedico, "Plasma Focus Spacecraft Conceptual Design Study," CONF-950110, American Institute of Physics, 1995

T.L. Brown, H.E. LeMay Jr., B.E. Bursten, C.J. Murphy, P.M. Woodward, *Chemistry – The Central Science, 12th ed.,* Prentice-Hall, New York, 2012

Francis Chen, *Introduction to Plasma Physics and Controlled Fusion, vol. 1: plasma physics 2nd ed*, Springer Science, New York, 2006

C.K. Choi, L.T. Cox, G.T. Nakafuji, "Fusion Plasma Thruster Using a Dense Plasma Focus Device", *Fusion Energy in Space Propulsion*, AIAA Progress in Aeronautics and Astronautics, Vol. 167, AIAA, 1995

J.H. Degnan, et. al., "Compact Toroid Formation, Compression and Acceleration," *Physics of Fluids B*, Vol. 5, No. 8, American Institute of Physics, August 1993

James Dewar, *The Nuclear Rocket*, Apogee Books, Burlington, Ontario, 2009

T.J. Dolan, *Fusion Research – Principles, Experiments and Technology, vols. 1-III*, Pergamon Press, 1982

Freeman Dyson, “Interstellar Transport”, *Physics Today*, October, 1968

David Griffiths, *Introduction to Electrodynamics*, 2nd ed., Prentice-Hall, New York, 1989

J.D. Huba, *NRL Plasma Formulary*, Plasma Physics Division, Naval Research Laboratory, Washington DC, 2000

Tatjana Jevremovic, *Nuclear Principles in Engineering*, Springer Science, New York, 2005

C.R. Joyner, J.E. Phillips III, R.B. Fowler, S.K. Borowski, *TRITON: A Trimodal-capable, Thrust Optimized, Nuclear Propulsion and Power System for Advanced Space Missions*, AIAA publication 2004-3863, presented at the AIAA Joint Propulsion Conference, Fort Lauderdale, Florida, 2004

T. Kammash, P. Zarchan, ed., *Fusion Energy in Space Propulsion*, AIAA Progress in Aeronautics and Astronautics, Vol. 167, AIAA, 1995

T. Kammash, "Principles of Fusion Energy Utilization in Space Propulsion", *Fusion Energy in Space Propulsion*, AIAA Progress in Aeronautics and Astronautics, Vol. 167, AIAA, 1995

E. Kuffel, W.S. Zaengl, *High Voltage Engineering Fundamentals*, Pergamon Press, Oxford, 1984

K.F. Long, R.K. Obousy, A.C. Tziolas, A. Mann, R. Osborne, A. Presby, M. Fogg, *Project Icarus: Son of Daedalus – Flying Closer to Another Star*, Presented at the British Interplanetary Society, Daedalus After 30 years Symposium, September 2009, revised May 2010

Ian McNab, ed., *Pulsed Power*, Institute for Advanced Technology, University of Texas at Austin, 1996

P.G. Mikellides, P.J. Turchi, I.G. Mikellides, “Design of a Fusion Propulsion System – Part 1: Gigawatt-Level Magnetoplasmadynamic Source”, *AIAA Journal of Propulsion and Power vol. 18 no. 1*, January, 2002

P.G. Mikellides, P.J. Turchi, I.G. Mikellides, “Design of a Fusion Propulsion System – Part 2: Numerical Simulation of Magnetic Nozzle Flows”, *AIAA Journal of Propulsion and Power vol. 18 no. 1*, January, 2002

M.M. Milanese, R.L. Moroso, J.O. Pouzo, "Dynamics of the Ionizing and Magnetic Fronts in the Radial Compression Stage of a DPF Current Sheath," *IEEE Transactions on Plasma Science*, Vol. 21, No. 4, August 1993

Richard Obousy, *Project Icarus: A technical review of the Daedalus propulsion configuration and some engineering considerations for the Icarus vehicle*, paper presented at the Advanced Space Propulsion Workshop, 2010

Jay Pasachoff and Alex Filippenko, *The Cosmos – Astronomy in the New Millennium, 3rd ed.*, Thomson, Brooks and Cole, Belmont, California, 2007

Robert Peterkin, "Direct Electromagnetic Acceleration of a Compact Toroid to High Density and High Speed," *Physical Review Letters*, Vol. 74, No. 16, American Institute of Physics, April 1995

George Sutton, Oscar Biblarz, *Rocket Propulsion Elements*, 7th ed., John Wiley and Sons, New York, 2001

Yu.N. Zhukov, A.I. Markoliya, A.F. Popov, A.F. Chachakov, "Physical Processes in the Final Phase of the Plasma Focus," *Technical Physics*, Vol. 46, No. 11, 2001, pp. 1379–1383, Translated from *Zhurnal Tekhnicheskoï Fiziki*, Vol. 71, No. 11, 2001 MAIK "Nauka/Interperiodica," pp. 32–36.