

GUIDED IMPACT FUSION

v1.0

Colin Jack

ColinBJ@gmail.com

© 2009-2011

Comments welcome. For permission to distribute this document, please email me. Technical information disclosed is protected by patents with a base priority date of 4th January 2010.

CONTENTS

Abstract	2
1. INTRODUCTION / MANAGEMENT SUMMARY	3
2. BASELINE DESIGN	6
3. DEVELOPMENT PATH	28
Appendix: Electrical force calculations	31
References	33

ABSTRACT

There is a straightforward and cost-effective way to generate energy from fusion. The enabling technology has existed for at least a decade: it is an indictment of the way science currently progresses that it has been overlooked.

The basic method is familiar: a hollow fuel capsule implodes within a hohlraum. However the hohlraum is heated not by lasers, but by the impact of charged micropellets fired at ultravelocity from a modified particle accelerator. This technique has long been used to test spacecraft micrometeoroid shields, and has been suggested for fusion. The key novel step is that it is now possible to track and guide each pellet individually during flight, using COTS-available technology. This opens up options never before considered:

- The pellets catch up together over a long flightpath, so peak power level can be multiplied a millionfold. A train of pellets launched over a period of milliseconds arrives at the hohlraum within a span of nanoseconds: an accelerator of modest power can provide a larger peak input than is possible with lasers.
- The pellets are progressively discharged as they travel, so mutual repulsion at convergence is eliminated.
- The pellets impact the hohlraum in a precisely specified pattern, whose variation with time is also tailored to cause optimal fuel capsule implosion.

The method is ideally suited to standoff operation. Detonation can take place within a disposable projectile, within a lithium waterfall which extracts the energy while breeding replacement tritium. There is no need for a large vacuum chamber.

The only net fuel input is deuterium. Capital cost is modest. Equipment life is indefinite. It will be possible to retrofit existing coal-fired generating plant for fusion.

Overall length of the accelerator and standoff pipe is substantial, several kilometres. However even if the whole length has to be placed in a tunnel, its cost is small compared to that of a power station.

1. INTRODUCTION/MANAGEMENT SUMMARY

Igniting inertial confinement fusion should not be difficult in principle. A quantity of energy comparable to that in a hand grenade must be injected into a volume of space the size of a raindrop, in the time it takes light to cross a room. However, achieving this with lasers has turned out to be very challenging.

Now a better way has been found: harnessing the kinetic energy of pellets fired from a linear accelerator. The linear accelerator is situated a kilometre or more from the fusion site and fires a graduated-speed train of pellets into a vacuum pipe connecting the two, with the later ones travelling faster.

The course of each pellet is individually monitored and corrected several times during flight, so that all catch up together in a precisely specified pattern as they approach the target. The pellets are progressively discharged by electron beam as successively finer course corrections are applied, so that mutual repulsion as they converge is never excessive.

The pellets strike a foil to produce a tailored pulse of X-rays which implodes a fuel capsule, just as planned for laser-driven fusion, but with far more power available.

Key advantages of the approach include:

- A linear accelerator of relatively modest power can be used. An initially long train of pellets fired over a period of milliseconds catch up together to arrive within a span of nanoseconds, multiplying the peak power rate a millionfold.
- The spacing and timing of pellets arriving at the hohlraum can be set with exquisite precision, independent of the distance between accelerator and target.
- Standoff is achieved with an inexpensive pipe, which contains a modest grade of vacuum compared to that required in a fundamental particle accelerator. If the pipe is given a radius of a

few centimetres, comparable deviations due to earth movement and wind can be tolerated: the pipe can be mounted in open air.

- A vacuum chamber is not required for the fusion. Detonation takes place within a sacrificial projectile which has been fired into a chamber containing a lithium waterfall.
- Almost all neutrons and kinetic energy are absorbed to heat the lithium: virtually none reaches the chamber wall, which has indefinite working life. Energy is extracted by circulating the lithium through a heat exchanger using electromagnetic pumps.
- All tritium required is produced in the lithium. High energy neutrons interacting with ${}^7\text{Li}$ can generate both tritium and a second neutron, so the fuel cycle can be closed with less than 100% capture efficiency.
- The scaling laws to reach higher ignition energies are favourable. Even if a very large ignition pulse turns out to be required, it can be provided.
- Energy release in the gigajoule range can be easily contained. Repeat rate can therefore be modest, minimizing the number of targets used.
- There are no military applications. The system is too long to fit aboard any vehicle, even a large oceangoing vessel, and incapable of detonating anything larger than a fuel capsule. Simpler ways to manufacture neutrons and tritium exist.

The key enabling technologies are:

- Operating at modest frequency compared to a particle accelerator, the pellet accelerator can be powered by cheap solid state MOSFETs providing variable frequency power.
- Accelerator tube is similar to the 'dielectric wall' currently under development for medical use.
- Los Alamos have demonstrated that a modified particle accelerator can fire a continuous stream of pellets even from an impure source

without performance degradation: a chicane weeds out pellets whose charge/mass ratio is outside tolerance.

- Strong pellets capable of bearing high charge/mass ratio can be manufactured cheaply and simply.
- Inexpensive CCD or CMOS cameras along the beamline measure the position of each pellet as it passes to sub-micron accuracy: picosecond-pulse lasers control exposure.
- Small numbers of rapidly switchable electrode pairs along the beamline tweak the pellet trajectories. Inexpensive solid state power switches are available operating at ~20 GHz.
- Fuel capsule design is identical to that already developed for laser-driven fusion.

An electricity generating station in the gigawatt range will have a capital cost similar to a coal-fired unit, but will operate with negligible fuel cost. Power generated will be cheaper than from any other source. It will be possible to retrofit most existing coal-fired stations for fusion.

R & D required is modest. The technology of linear accelerators, including the modifications required to fire charged pellets, is mature. All components required for pellet tracking and steering are COTS-available. The expertise necessary for development, electrical engineering closely related to particle accelerator design, is widely available.

There is an enormous margin, two or more orders of magnitude, between the maximum impact energy that can be delivered by the accelerator and that expected to be necessary. In stark contrast to every other fusion system proposed to date, technical risk is close to zero.

Full IP protection will be available during development. 24 of 29 claims of patent PCT/GB2011/000009, describing the apparatus and methods herein with a base priority date of 4th January 2010, have been accepted by the international examiner.

The goal of unlimited cheap, clean energy is easily achieved. It is a grave reflection on the system by which science is currently progressed that the technique has been feasible for over a decade, yet has been overlooked.

2. BASELINE DESIGN

To minimize technical uncertainty and facilitate comparison, the Baseline Design uses a hollow spherical fuel capsule identical to those developed for laser driven fusion, and driven by an identical X-ray pulse. The capsule contains deuterium-tritium fuel. The X-rays evaporate material from its outer wall, causing it to implode by reactive force, compressing and igniting the fuel.

In light of the National Ignition Facility's performance problems, a larger capsule is assumed, as designed for NIF's intended successor LIFE. Moreover an increase in capsule absorbed energy from LIFE's baseline value of 0.77 MJ to 1.0 MJ is assumed. Design parameters for NIF, LIFE and Guided Impact Fusion GIF are compared in Table 1.

Note that for GIF, the enormously conservative assumption is made that just 6% of the energy entering the hohlraum reaches the capsule. This is in contrast to the highly optimistic 25-33% projected for NIF and LIFE.

TABLE 1


		NIF	LIFE ^[1]	GIF
capsule diameter	mm	2.2	4.1	4.1
hohlraum temperature	eV	300	250	250
hohlraum received energy	MJ	1.8	2.2	16
capsule received energy	MJ	0.2-0.45	0.77	1.0
fusion thermal output	MJ	16	200	200

Because electrical acceleration of pellets is far more efficient than electrical generation of laser energy, guided impact fusion achieves a

favourable overall output ratio. Thermal energy available after capsule detonation is 200 MJ from fusion, plus 40 MJ from fission in the lithium blanket surrounding the reaction, plus the 16 MJ kinetic energy originally input = 256 MJ. This generates over 115 MJ of electricity, of which about 25 MJ is required to generate the next pellet pulse. Net electricity output is 90 MJ: 11 detonations per second can drive a 1-GW power station.

The temporal shape of the X-ray pulse required is shown in Figure 1: a long initial constant pulse, whose energy rate then rises exponentially to a peak of some 500 times the initial value. This is equivalent to black-body radiation temperature rising from 50 to 250 eV.

FIGURE 1 Compression/ignition pulse variation with time^[2, p50]



W power received by capsule, watts

eV temperature, 1 eV ~11,600 degrees Kelvin)

t1, t2 ~ 20 nanoseconds each: power doubles every ~2 ns during rise

The GIF hohlraum is shown in Figure 3, beneath the LIFE hohlraum drawn at the same scale in Figure 2 for comparison. The GIF hohlraum is asymmetric because pellets approach from one direction only (the right).

FIGURE 2 LIFE hohlraum (major features to scale)




FIGURE 3 GIF hohlraum (major features to scale)

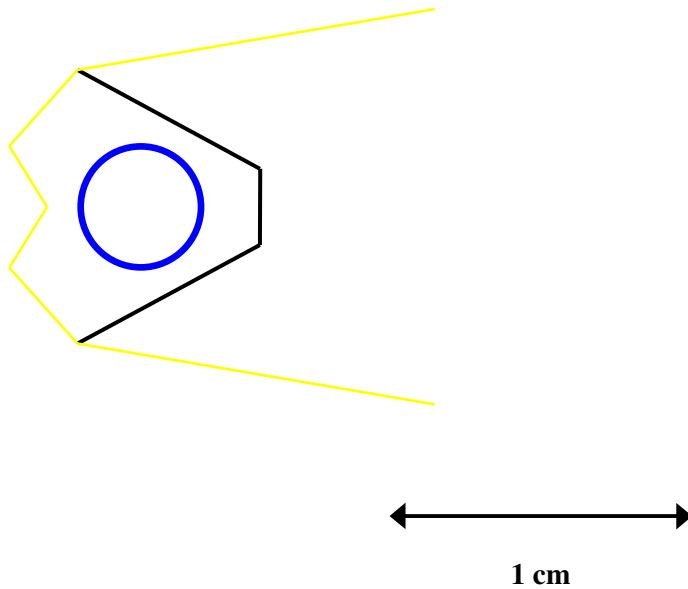




FIGURE 4 Compression sequence

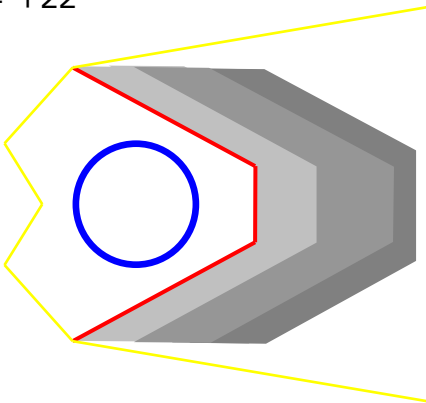
$t = -15$



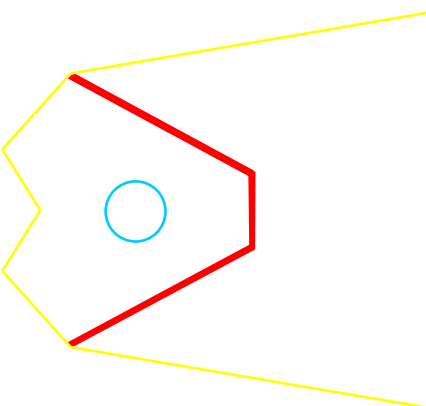
$t = +2$



$t = +22$



$t = +42$



To release their energy, the incoming pellets need to strike a stationary mass: this is provided as the membrane shown in black. Figure 4 shows the sequence: times are nanoseconds before/after first pellet impact. A suitable pellet speed is ~ 500 km/sec, so each pellet moves about 0.5 mm per nanosecond. The pellet speeds actually range from 450 km/sec to 550 km/sec, so the pellet cloud contracts as it travels, as shown.

A first wave, comprising about 7% of the total pellet mass, strikes the membrane at $t=0$, heating it by shockwave to form a 'collision plasma' at 50 eV. At this temperature radiative loss is comparatively slow, so few or no pellets need to strike over the next 20 ns as the fuel capsule outer layer becomes preheated and the second wave, the main pellet cloud, approaches. The internal density distribution of the main pellet cloud is chosen such that the collision plasma temperature then increases as shown in Figure 1.

During this time, the collision plasma moves leftward relative to the original membrane position due to conservation of momentum. Almost all of this movement results from first wave impact, as the bulk of the second wave arrives only in the final nanoseconds. Total movement over the 40-nanosecond period is ~ 0.3 mm. By this point, the fuel capsule surface is already imploding at a rate greater than the collision plasma is approaching. There is no need for large clearance between the membrane and the fuel capsule because (a) the pellets can be distributed in a precise pattern to heat the membrane very evenly, unlike the case with a laser pulse heating a hohlraum wall; (b) the rocket-exhaust evaporation from the surface of the fuel capsule itself protects it from interaction with the collision plasma.

The fuel capsule experiences a temperature environment, an X-ray 'bath', identical to that planned for laser-driven fusion. The capsule implodes and ignites under its own momentum.

After colliding with an equal mass of membrane, half of the original pellet energy is contained in the linear momentum of the collision plasma. The

rest, released by collision at 250 k/sec w.r.t. the centre of mass, would heat the plasma to 450 eV if no thermal energy escaped. In practice, useful energy is received by the capsule until the plasma temperature falls to about 10% below the peak temperature required, to 225 eV, so half the plasma thermal energy is radiated at a useful rate. One-quarter of the original kinetic energy is so emitted, 4 MJ. Of this about half, 2 MJ, is ultimately radiated leftward.

From this point on, the right and left halves of the fuel capsule should be considered separately. We require each hemisphere to receive 0.5 MJ. The right hemisphere receives essentially all its radiation directly from the closely adjacent collision plasma: 0.5 MJ. The left half receives radiation mainly indirectly, after bouncing off the hohlraum wall. Hohlraum wall losses are conservatively put at two-thirds, so the left hemisphere must be sent 1.5 MJ. The total requirement is 2 MJ, as provided.

The energy flow per detonation cycle is summarised in Table 2. 11 detonations per second provide 1 GW grid power.

TABLE 2 Energy per detonation

Thermal input	MJ
Kinetic energy of pellets	16
Fusion energy	200
Fission in lithium waterfall	40
Total	256
Electricity output	
Electric generation efficiency	45%
Gross electric output	115
Returned to accelerator	25
Net output	90

Electric acceleration can be 50-90% efficient, while even future generation lasers of suitable type are unlikely to be more than 18% efficient^[3], so overall efficiency of the pellet system remains better than lasers, despite the high ratio of kinetic energy supplied to X-ray energy received by the capsule.

Total pellet mass required, delivered at 500 km/sec average speed, is 128 mg. The mass required for the membrane is the same, generating 256 mg total collision plasma. Its area is 1 cm².

The opacity of the collision plasma is an issue. It must not be so transparent that it fails to radiate as a black body, nor so opaque that energy is significantly delayed in escaping.

Realistic pellet material choices are very strong hard substances: as explained in the next section, diamond, Widalite aluminium-lithium alloy, or maraging steel would all be feasible to accelerate. The Rosseland opacity of the corresponding plasma at the maximum and minimum temperatures relevant, 50 and 250 eV, is as shown in Table 3.

TABLE 3 Rosseland opacity of plasma (cm²/g) [4][2,p359]

Material	Main elements	Opacity 50 eV	Opacity 250 eV
Diamond	Carbon 100% (z=6)	1000	10
Widalite 049	Aluminium 92% (z=13) Copper 6% (z=29) Lithium 1.3% (z=3)	8000	20
Maraging steel	Iron 70% (z=26) Nickel 18% (z=28) Cobalt 8% (z=27)	12000	200

The membrane contributes 125 mg/cm^2 to the collision plasma; the pellet contribution rises from 10 mg/cm^2 in the preheat period to 125 mg/cm^2 as peak temperature is reached. A suitable material choice for the membrane is carbon of thickness 0.35 mm.

A suitable choice for the pellets is either synthetic diamond or Weldalite. If diamond, collision plasma optical thickness goes from ~ 130 initial to ~ 2.5 final; if Weldalite, from ~ 200 initial to ~ 4 final. Radiant energy need escape only relatively gradually during the preheat phase, but should be squeezed out rapidly as temperature approaches the maximum: either diamond or Weldalite pellets are therefore acceptable. Maraging steel is too opaque.

2.1 PELLETS

Charged pellets can be accelerated to extreme speeds in a modified particle accelerator, a technique used since the 1960s to simulate the effect of ultravelocity micrometeoroid impacts on spacecraft. A landmark Los Alamos design^[5] included chicane-style wiggles to exclude all material not of the correct mass and charge, so that the accelerator remained clean of conducting dust and performed perfectly in continuous operation, even when the pellet source was a very imperfect monodisperse: commercially available iron microspheres of variable size contaminated with nanoparticles.

To minimize accelerator length, pellet charge/mass ratio should be maximized. The pellets should be given positive rather than negative initial charge, to avoid field effect electron current leakage, and to permit later controlled charge reduction by electron addition. The charge which can be placed on them is then limited by two closely related effects: field effect evaporation of atoms from the surface, and burst-apart due to mutual charge repulsion. At ordinary temperatures burst-apart is the limiting factor.

The additional stress due to applied acceleration $\sim 10^7 g$ is small compared to self-repulsion. However, as they proceed through a 2-phase accelerator, the pellets experience force cycling rapidly from maximum between electrodes to zero within an electrode. This does no harm to an electron or atomic nucleus, but repeatedly flexes a larger object, and induces an alternating voltage across it. The pellets must therefore satisfy the following criteria:

1. High specific strength, for high charge/weight ratio
2. High stiffness, so do not melt due to flexural heating
3. Must not be damaged by induced voltage or current
4. Must be reasonable cost

Achievable charge/mass ratio, as calculated in the Appendix, is:

$$Q/M = \frac{1.26 \times 10^{-5} \sqrt{\sigma}}{R\rho}$$

where R is the radius, ρ the density and σ the permissible tensile stress.

To satisfy criterion 2, the stiffness/density ratio should be sufficient that the propagation time of a compression/tension wave across the pellet is less than the accelerator cycle time at the maximum frequency used.

To satisfy criterion 3, the pellet should either be conducting so that surface charge can flow freely, or else have dielectric strength greater than the peak accelerator voltage gradient experienced.

Realistic choices are shown in Table 4. Diamond would be ideal. However at the moment UNCD (ultra-nano-crystalline diamond, much stronger than standard CVD) is still viewed as an experimental material, manufactured in small batches. By contrast metal microspheres of the required size can be made simply by spraying droplets of molten metal which solidify as they fall, the technique originally used to make ball bearings.

Maraging steel is unsuitable due to the high z-number of its constituent elements, which would give rise to an unacceptably opaque plasma as described in the previous section.

Aluminium-lithium alloy is therefore chosen. The Weldalite™ grade shown reaches its maximum strength without working, by heat treatment alone

(hence the trade name) so is ideally suited to making strong microspheres.

TABLE 4 Pellet candidate materials

material	density	strength GPa	relative accelerator length
Diamond UNCD	3.5	3.0 ^[6]	1
Maraging steel	8.1	2.4 <i>yield</i>	2.2
Al-Li Weldalite 049-T8	2.7	0.69 <i>yield</i> ^[7]	2.4

Pellet properties are summarised in Table 5. 200,000 pellets are used per pulse.

TABLE 5 Pellet properties

Material	Al-Li Weldalite 049-T8
Density	2.7 g/cc
Yield strength	0.69 GPa
Diameter	77 μm
Capacitance	4.28×10^{-15} F
Mass	0.64 μg
Kinetic energy (average)	80 J
Charge	1.67 nC
Charge/mass ratio	2.6 C/kg
Stress	0.46 GPa (=> safety factor 1.5)
Potential	390 kV
Surface field	1.0* V/Å

*This is comfortably below the level, 2-4 V/Å, at which field effect evaporation from an aluminium surface becomes significant.^[8]

2.2 ACCELERATOR

To strike the hohlraum together, pellet launch speeds must vary from 450 to 550 km/sec from start to end of the burst fired. The higher speed corresponds to acceleration through 58 GV potential difference.

In some ways the accelerator resembles a fundamental particle accelerator. However the pellet speed is <1% lightspeed, so the electrodes do not act as RF resonant cavities: the tube walls should be insulating material rather than metal. Of current designs, the accelerator most closely resembles the 'dielectric wall' under development at Lawrence Livermore^[9] to fire relatively low energy nuclei for medical applications.

The volts per metre achievable are limited by:

- Vacuum breakdown field of electrodes
- Bulk dielectric breakdown strength of insulator
- Surface effect breakdown strength of insulator

Treated electrodes can withstand surface field strengths of^[10]:

Copper	134 MV/m
Molybdenum	266 MV/m
Tungsten	406 MV/m

Bulk insulating materials such as Teflon and Rexolite are vulnerable to surface flashover, even in vacuo, at field strengths of a few MV/m. However research for the dielectric wall accelerator has shown that insulator suitably interspersed with conducting material, for example a sandwich structure with dielectric layers alternating with thin layers of conductor, or bulk dielectric with metal inclusions, performs much better, to >100 MV/m without either internal breakdown or surface flashover.^[9] This applies even with metal layers or inclusions whose voltage 'floats', unconnected to any external source.

Peak inter-electrode gradient can therefore be at least 100 MV/m. Allowing a safety factor of 1.5, and taking into account that a pellet passing along the tube experiences an average field ~0.6 of the peak field

strength, 40 MV/m working value is achievable, hence accelerator length 1.45 km.

An appropriate drive frequency at the relatively low pellet speed (1/600 lightspeed) is ~ 100 MHz. Exotic sources such as klystrons are not required: at this frequency solid state components work well and are cheaper. An example of a suitable RF power MOSFET is the IXYS IXZ2210N50L whose cost in bulk is $\sim \$50/\text{unit}^{[11]}$. When used to provide power in pulses of length 1-10 ms, it can deliver ~ 18 J per pulse, irrespective of the exact pulse length. (IXZ210N50L Safe Operating Area graph^[12], multiplied by 1.8 for model 2210 relative to 210.) Thus RF MOSFET cost is $\sim \$3/\text{J}$.

Allowing for 20% losses downstream of the MOSFETs, total energy 20 MJ must be provided at RF frequency: total MOSFET cost will be \$60 million.

Drive frequency of the accelerator is increased uniformly from 90 MHz to 110 MHz from the time at which the first pellet leaves the accelerator to the time at which the last does so. Electrode interval is 2.5 mm at the fast end, with electrode voltage cycling between ± 83 kV. To achieve this, the output from the MOSFETs is fed into compact air-core transformers.

At the low-speed end of the accelerator there will be a minimum acceptable pellet separation of about 0.5 mm, due to inter-pellet repulsion (as calculated in the Appendix) and practical fabrication constraints. In place of what would be the first 15m of the accelerator, a low-power leader section 120m long is therefore provided, with fixed electrode separation 0.5 mm and highly variable drive frequency from 2 MHz to 90 MHz. This section is fed with pellets which are precharged to 40 kV and fired in at 1 km/sec from a source at 2 MV potential over a period of 0.1 seconds. The 2 MHz pellet feed rate required is an order of magnitude lower than the rate at which a modern inkjet printer head ejects droplets of similar size.

After 0.1 seconds, the leader contains 200,000 pellets in a 100m length line, levitated by a DC vertical field component. The variable frequency electrodes then accelerate this line en bloc to 45 km/sec as the front pellets enters the main accelerator at 90 MHz, still at 0.5 mm separation

as required. During this process the pellet voltage is raised to 390 kV by offsetting the local electrode voltages to this level: note that electrons can flow easily from a pellet to the electrode it passes through, but not vice versa, as the pellet acts as a discharge point source while the treated electrode surface is smooth.

Focussing to keep the pellets centered during acceleration can be provided by electric or magnetic fields, e.g. quadrupole magnets as used in particle accelerators.

Pellets are strength-tested by charging them to slightly above operational voltage before firing. In-tube pellet failure is therefore unlikely. Any failure which does occur has the potential to become contagious. However worst-case energy release is 16 MJ. The tube can be wrapped in a Kevlar blanket so that external equipment is not affected and the damage is kept localised.

Chicane wiggles can be incorporated at multiple points in the system: slight bends with lateral electric fields which divert pellets of exactly the correct charge/mass ratio into the next section, but allow any other material to fly on into open ended 'dump tubes' in which the plasma from their impact is safely contained. If desired, pellets can be imaged in flight by camera systems similar to those described below: the divert electrodes are then made switchable so that only pellets of the correct shape and size, and which have no extraneous material nearby, are deflected into the next accelerator portion, all other material being rejected.

The overall electrical efficiency of the accelerator can be well over 50%. Pellet trains of kinetic energy 16 MJ are provided at rate 11 Hz, so accelerator consumption is ~250 MW, 20% of the 1.25 GW gross output of the power station it drives.

Unlike lasers, the repeat firing rate and service life of the accelerator are essentially unlimited, and its maximum power delivery rate is not constrained by laser-plasma interactions. Its overall cost, calculated in section 2.5, is vastly less than that of equivalent lasers.

2.3 STANDOFF PIPE


200,000 pellets emerge from the accelerator in a train of initial length 1 km. All pellets reach a common point 5 km downstream at the same instant, after passing along a vacuum pipe of this length.

Containing a relatively soft vacuum by particle accelerator standards, the standoff pipe does not need to be housed in a building, and comprises a simple pipe, mounted on pylons or stilts likely shared with a power transmission line. It is given a much larger internal radius than the maximum pellet deviation from the beamline, so that it can tolerate corresponding lateral displacements due to wind etc. of several centimetres, and also so that its vacuum can be maintained by pumping from a limited number of points along its length. The vacuum pumps are mounted in Portakabin-size units spaced every mile or so, which also contain the pellet course correction systems.

The pellets must be steered with very high precision, preferably $\sim 1 \mu\text{m}$. To this end their individual trajectories are measured and corrected at several successive points during the flight. This is demanding because of the large number of pellets and the high rate of pellet flow past a given point, increasing linearly from $\sim 100 \text{ MHz}$ at accelerator exit to theoretical infinity at the collision point.

Tracking stations comprise lines of paired cameras, with a pulsed laser providing light via an optical fibre as seen in Figure 5: cameras A and B point at the beamline C (dotted circle in which individual pellets are solid black); photons from leaky optical fibre D set in reflecting trough E illuminate the beamline. The cameras use standard CCD or CMOS chips with microscope-style lens turrets. Exposure is controlled by the laser, not by shuttering: inexpensive chips with shutter rate $\sim 25 \text{ Hz}$ have similar readout rate, a few tens of Mbytes/sec, to expensive fast-shuttered units.

FIGURE 5 Tracking camera pair



Each camera can report the positions of many pellets. To avoid confusion from overlaps, pellet trajectories are spread, and timings chosen, so that each pellet imaged appears in a different part of the camera's field of view: for example a camera's 1000x1000 pixel field is subdivided into 1000 squares each of 30x30 pixels, at the nominal centre of which a given pellet is expected to be seen. 200 camera pairs are then required at each tracking station, in order that each pellet can be imaged once from two different angles.

If three uniformly spaced course correction stations are provided, each making successively finer steering corrections with the first close to the accelerator exit, the final station will be 1.7 km upstream of the collision point and see a peak pellet passage rate of 330 MHz.

Read-out is a potential bottleneck. However both CCD and CMOS chips can perform on-board processing: the first stage of this involves pixel binning of the rows and columns of each subsquare if CCD chips are used, or massively parallel ADC if CMOS chips are used. Minimum output required is 2 bytes per pellet imaged, representing its X,Y position within its subsquare: a modest 3 MHz byte rate per camera.

The lasers are standard laboratory desktop units providing 10 μ J output pulses of duration 1 picosecond at repeat rate up to 2 MHz. These are available COTS from various suppliers, e.g. the Tangerine from Amplitude Systemes^[13]. A 10 μ J laser pulse contains 3x10¹³ photons, so a single illuminating laser can potentially serve all cameras. A pellet moves 0.5 μ m in 1 picosecond.

After measurement, individual trajectories are tweaked as required in the vertical and horizontal directions and in speed. Lateral steering could be done using either electric or magnetic fields. At relativistic speeds in a fundamental particle accelerator, magnetic fields are more effective. However at 1000 km/sec, the lateral force from a 1 Tesla magnetic field is equivalent to only a 1 kV/mm electric field: a switchable electric field of at least 40 kV/mm across the beamline can easily be provided. Electric fields are therefore used. To apply independent corrections to each pellet, electrode length should be about half the pellet separation, which reduces from 5 mm at accelerator exit to 1.7 mm at the last correction station.

A 40 kV/mm lateral field would deflect a fully charged pellet ~ 0.4 microradian (0.2 m/s) per millimetre of electrode length. In practice, to minimize inter-pellet interaction in the beamline, pellet charge is reduced by a factor 10 immediately on accelerator exit and a further 3 after the first course correction, as described in the Appendix. Reducing the pellet charge by factor n increases the total electrode length needed proportionately. However, many consecutive independently switchable electrodes can be provided, so there is no practical limit to the size of course correction that can be applied. Three sets of electrodes are needed at each station, to permit vertical, horizontal and speed adjustment. In practice maximum correction needed will be a few microradians at the first station, much less at the second and third, so the total number of steering electrodes and associated switches needed is trivial compared to the accelerator.

The pellets are discharged in stages by electron guns supplying total current a few tens of milliamps. Precise charge/mass ratio of a pellet at any stage can be measured from its deflection by a known DC field.

(The charge can be fine-tuned by providing a spatially modulated flux of electrons downstream of the DC field, such that the maximally deflected pellets, with highest initial charge/mass ratio, subsequently experience the highest electron flux. Alternatively, the flow from an electron gun can be actively switched to adjust the per-pellet dose required, or the system can simply log the exact charge/mass ratio of the pellet, so that the next

course correction can be modulated to supply exactly the field strength required.)

The easiest way to deliver a precisely metered correction to each pellet is to switch the deflecting electrode on or off at a precisely chosen moment during pellet passage. Switching rate of at least 20 GHz, as compared to maximum pellet passage rate of 330 MHz, is available. As successively finer trajectory adjustments are made, ultimate positional accuracy is limited only by the precision of the tracking cameras. Using visible light, ~ 0.2 microns is achievable.

It is not anticipated that course corrections close to the impact point will be needed. However, independent last minute corrections to each pellet trajectory could be made as little as 250m upstream from the impact point if required.

2.4 REACTION CHAMBER

A particular advantage of the system is that a vacuum chamber is not required for the fusion. By using a sacrificial projectile to protect the pellets during the final part of their journey, the detonation can be made to take place in a void surrounded by a lithium waterfall which should provide 50-100 cm thickness of liquid lithium: 1-2m waterfall thickness if the lithium droplets are 50% space-filling.^[14]

80% of the fusion energy released goes to heat the liquid lithium by direct neutron interaction, with no explosive release. The remaining 20% is in the form of energetic α -particles. This vaporises the nearest part of the projectile, the hohlraum and its mount. However because this constitutes a small mass, ~ 1 g surrounded by near-vacuum, although the total energy here (including the original pellet kinetic energy) is 60 MJ, the associated momentum pulse is only ~ 350 kg.m/s. This high-speed plasma loses almost all its kinetic energy as it strikes consecutive droplets of the lithium waterfall, and while the associated outward momentum is not lost, it reaches the chamber wall as a relatively lengthy pulse of modest overpressure rather than an intense shock wave.


The lithium is circulated through a heat exchanger by electromagnetic pumping. Lithium has high boiling point and heat capacity, 1347°C and 3.85 J/gK, with density 0.5 g/cc: in theory a flow of as little as 2 m³/sec could extract enough heat to generate 1 GW electricity. In practice a flow rate about ten times larger will be used, both for reasons of thermodynamic efficiency, and to provide the 'waterfall' which protects the chamber walls. 1 metre net lithium thickness gives a stainless steel chamber an indefinite working life.

Almost all neutrons produced are absorbed by the lithium to breed tritium. Initial high energy neutrons from the DT reaction can be absorbed by ⁷Li to produce both a tritium atom and a lower energy neutron, which can breed a further tritium atom from ⁶Li. Natural lithium comprises a mixture of these isotopes, 7.5% ⁶Li to 92.5% ⁷Li: this ratio can easily be altered by fractional distillation. Thus it is possible to fine-tune the system so that exactly 100% of the tritium required for continued operation is produced.

(A further 'tritium supplement' comes from deuterium-deuterium reactions. Two types of DD reaction occur: one produces tritium directly, the other releases a neutron which breeds tritium on hitting the lithium blanket. Their combined rate is about 1% that of the DT reaction.)

FIGURE 6 Reaction chamber

(not to scale: actual pellet number will be thousands, each ~0.1mm diameter)



The chamber is shown in Figure 6, as seen from above in cross-section: the lithium flows vertically down into the page, to collect in a sump. The sacrificial projectile comprises a hollow tube open at the rear, with the hohlraum containing the fuel capsule fixed at the forward end. The projectile is fired into the reaction chamber at ~ 100 m/s. The detonation pellets subsequently travel through the high vacuum preserved in the projectile interior to strike the hohlraum foil.

The outer walls of the projectile can be notched as shown to create toroidal cavities which act as a cascade of cold traps. Traces of lithium vapour escaping past the projectile could compromise the good vacuum behind and within it: this arrangement ensures that none does so. The rotating barrel used to insert projectiles acts as a barrier between the chamber and the standoff pipe at other times: additional doors, cold traps for metal vapor and electrostatic traps for light gas atoms can be provided.

The injection tube allows the projectile to pass through the lithium waterfall without losing speed. While the chamber itself has indefinite life, the injection tube is inexpensive and can be replaced at regular intervals. Nevertheless it is well protected at the moment of detonation: neutrons and blast wave must pass through a long slant length of projectile wall before reaching any part of it.

The sacrificial projectile can be made of lithium or aluminium-lithium alloy, which is recovered from the melt as it circulates, and poured into moulds to make more projectiles at the required rate. Unburned deuterium and tritium is recovered from the molten lithium as it circulates, a simple chemical separation as there is no other hydrogen present.

Overall size of the chamber is determined by the gap size between the central point and the inner side of the lithium waterfall. At the inner side of the waterfall, each centimetre thickness of lithium intercepts $\sim 1.5\%$ of the fusion energy flux.^[14, interpolated from Figure 6] It is desirable that this is not sufficient to boil the lithium: energy deposition can be up to 1.5 kJ/cc if the coolest lithium emerging from the heat exchanger is sprayed at the inner side of the waterfall. For 200 MJ thermal energy fusion detonations

in the baseline system, this implies an inner waterfall surface area of 2,000 cm², inner radius 13 cm. In theory, chamber diameter could be as little as 2 metres. However, centrally released alpha-particle energy is sufficient to vaporise some 10 kg of lithium per pulse from cold, much more if the inner lithium is already raised to near boiling point by neutron energy. While such evaporated vapour quickly recondenses on other droplets, a more reasonable minimum chamber diameter is 4.5m, providing 1m net thickness of liquid lithium at 50% volume fill plus a 25cm inner waterfall radius. This will give a stainless steel containment chamber indefinite working life.

11 detonations per second provide 2.2 GJ of fusion energy. However an additional 20% is released by fission within the lithium, as a side effect of neutron capture, and a further 176 MJ of kinetic pellet energy is recovered, giving total thermal output 2.8 GJ. 45% efficient electricity generation yields 1.25 GW, of which ~250 MW is required to drive the accelerator, giving net electric output 1 GW as specified.

2.5 OVERALL SIZE AND COST

Electrical engineering

The largest single cost element of the accelerator is the RF power MOSFETs: total approximately \$60 million, as calculated in section 6.2.

Since duty cycle is low, ~2%, time-averaged power is only about 1/3 of the maximum the chip can provide: power supply and cooling cost will be moderate.

Raw materials for the accelerator tube itself will also be relatively cheap: volumetric cost of Rexolite insulator is \$80/litre, virgin grade Teflon \$60/litre, copper \$60/litre.

To allow for installation, control systems and auxiliary equipment such as the feeder system and trajectory adjustment stations, total electric system cost is nevertheless put at three times the MOSFET cost: \$180 million.

Civil Engineering

A 1.6 km long building is required to house the accelerator and its feeder – about the same length as the longest airport buildings – providing an environment of dry air at constant temperature. Single-storey industrial buildings typically cost ~\$1500/m² footprint^[15]. On this basis a building consisting of a 4m wide corridor would cost \$6,000/m. However the building must run straight and level for the length of a mile. Ground preparation cost may therefore be higher than normal. A double-track railroad on cheap land without geophysical complications costs^[16] ~\$5,000/m. Total cost of the accelerator building, including foundations and construction, is therefore put at \$10,000/m = \$16 million.

Standoff pipe cost is assumed similar to a 69kV single-circuit overhead transmission line^[17]: 5 km @ \$300/m = \$1.5 million.

Total accelerator system cost is therefore put at **\$200 million**.

The capital cost per GW of a conventional power station is around^[18]:

\$1.5 billion – coal-burning

\$1.5 billion – conventional hydroelectric


\$0.5 billion – gas turbine

The capital cost of a 1-GW fusion station will be about 15% greater than an equivalent coal-fired plant generating electricity from steam.

(The additional element is the accelerator cost only : the lithium chamber and its associated recycling equipment are assumed to replace a conventional furnace and coal handling equipment of comparable or greater cost.)

Fuel cost will be negligible, as compared to ~\$250 million/annum for a coal-fired plant. The fusion add-on will recoup its capital cost within one year.

FIGURE 7 Overall system length



Existing coal-fired plants, which already have steam turbines, generators and cooling towers, can be retrofitted for fusion.

Even if constraints of topology or surrounding urban development make it necessary for the entire length of the accelerator and standoff pipe to be housed in tunnels dug for the purpose, the additional cost of tunnelling will be ~\$36,000/m for the accelerator, ~\$20,000/metre for the standoff pipe, total \$160 million.

3. DEVELOPMENT PATH

Guided impact fusion will be a highly attractive option for all large scale energy generation, competitive on commercial grounds alone.

The necessary expertise for fast-track development, scientists and engineers with experience of building large particle accelerators, is already available. All components of the accelerator and guidance system are COTS-available.

There is therefore a strong case for proceeding immediately on a commercial basis, bypassing the glacially slow pace of academic research and government-funded development.

A laboratory-scale version can prove all aspects of the concept, including production of fusion neutrons. Given a baseline system accelerator, further optimisation can be done mainly in software: by making changes to the pellet trajectories, and small variations in hohlraum design. There is scope for further improvement.

3.1 REDUCING THE IGNITION PULSE ENERGY REQUIRED

The 16 MJ kinetic input assumed for the baseline system is almost certainly vastly pessimistic, for the following reasons:

- More than 50% of the energy radiated from the collision plasma can be sent forward toward the inner hohlraum and fuel capsule. The plasma has a z-gradient due to the forward part being mainly carbon, the trailing part mainly aluminium, so energy will preferentially escape forward. Part of that which escapes rearward is returned from the outer hohlraum; another part is absorbed by oncoming pellets, usefully preheating them to plasma before returning the energy on collision. As an additional measure, the final pellets could be made of maraging steel: the higher z-number of its ingredients will 'cap' rearward escape of X-rays.
- Increasing the speed of the pellets is beneficial in two ways. As the total pellet mass becomes smaller relative to the mass of the

stationary membrane, less energy is wasted in the final linear kinetic energy of the collision plasma. As the ratio of the energy input to the total collision plasma mass becomes higher, less energy is wasted in the final heat of the collision plasma. Both effects can be achieved with increased pellet speed, allowing a larger fraction of their kinetic energy to be usefully radiated.

- Research (theoretical and experimental) for both fast ignition and z-pinch fusion shows that implosion of a fully spherical fuel capsule is not necessary: imploding a hemisphere backed by a 'glide plate' of dense metal, or indeed a cone of almost any apex angle, works almost as well. If a hemisphere or cone whose curved base faces the oncoming pellets is used, coupling between the collision plasma and the capsule is excellent: very little X-ray energy is wasted.
- The mass distribution of the pellet cloud can be tailored in software alone, with the trajectory of each of the thousands of pellets set independently. It will be possible to produce a pulse shape far more precisely tailored to the optimum than is possible with lasers.

It may well be possible to increase the ratio of capsule-received energy to hohlraum-received energy from the 6.25% assumed for the GIF baseline design to something nearer the 25-30% aspiration for the LIFE and NIF laser-driven designs.

3.2 INCREASING THE IGNITION PULSE ENERGY AVAILABLE

Past experience of fusion development suggests that it will be wise to be able to provide much more than the theoretically calculated input energy if necessary. Greater power available for the compression-ignition pulse also allows simpler, cheaper and fewer fuel capsules and hohlraums to be used.

The cost of the accelerator RF MOSFETs increases linearly with the total energy provided in the input pulse. However the cost of other elements increases much more slowly. For example to provide 8 times the input energy using the same number of pellets at the same delivery speed,

pellet diameter doubles. This halves the charge/mass ratio possible for the pellets, only doubling the civil engineering cost element of the accelerator.

To maintain a favourable energy output ratio, the thermal energy obtained per detonation would also need to increase eightfold, to ~ 2 GJ. However this does not increase the cost of the containment chamber proportionately. The thickness of lithium waterfall needed is almost constant regardless of detonation size. It is only the inner radius of the waterfall that needs to increase, to avoid neutron energy alone being sufficient to vaporise the inner layer of lithium. If detonation energy increases 8 times, the inner area of the waterfall should also increase by 8, hence the inner radius by a factor of ~ 3 , to 75cm. Overall chamber diameter becomes 5.5m: total chamber volume is less than double that of the baseline system.

It would be therefore be possible to increase the input pulse power by a factor of 10 or more while keeping capital cost acceptable.

There is thus an enormous margin, two orders of magnitude, between the maximum energy that could be delivered by the accelerator and the minimum calculated to be required. In stark contrast to every other fusion system proposed to date, technical risk is close to zero.

APPENDIX: ELECTRIC FORCES WITHIN AND BETWEEN PELLETS

Maximum charge which can be placed on a pellet is limited by burst-apart due to self-repulsion. As the pellet is a sphere, charge will distribute itself evenly over the surface. The force calculation is then mathematically identical to the well known case of the self-gravity of a thin spherical shell. A point mass m at a distance R from the centre of a spherical shell of mass M experiences a gravitational pull of GmM / R^2 in the space outside the shell, and zero everywhere inside it. An average particle of the shell itself thus experiences $GmM / 2R^2$. The mass per unit area is $M / 4\pi R^2$, giving an inward surface pressure of $GM^2 / 8\pi R^4$.

The corresponding outward pressure on a spherical shell carrying charge Q is $kQ^2 / 8\pi R^4$, where $k = 1/4\pi\epsilon_0$. This must not exceed the safely usable strength σ of the material, so maximum charge permissible is $5.3 \times 10^{-5} R^2 \sqrt{\sigma}$. The mass of the sphere is $(4/3)\pi R^3 \rho$, so the maximum charge/mass ratio is:

$$Q/M = \frac{1.26 \times 10^{-5} \sqrt{\sigma}}{R\rho}$$

where σ is the usable tensile strength of the material after allowing a suitable safety factor.

The maximum charge which can be held is also limited by field effect evaporation (assuming a positively charged sphere: field effect evaporation typically requires a much higher field than field electron emission). Capacitance of the pellet is $4\pi\epsilon_0 R$, potential $V = Q / 4\pi\epsilon_0 R$, so surface field strength is $Q / 4\pi\epsilon_0 R^2$.

As regards inter-pellet interactions, the repulsion between two point charges q at separation r is kq^2/r^2 where $k \approx 9 \times 10^9$. At distance ~ 5 mm (the separation with which they leave the accelerator) the force between two pellets each carrying charge 1.67 nC is 1.0 mN, which accelerates each at 1.6×10^6 m/s². Theoretically, the opposing forces from nearest neighbours ahead and behind should cancel. However there is a stability problem in that a pellet displaced slightly from the nominal beamline will experience a lateral force proportional to the displacement.

Within the accelerator, this stability problem does not arise, as the electrode-induced acceleration is $\sim 10^8 \text{m/s}^2$, two orders of magnitude greater than the inter-pellet repulsion. However once in the standoff pipe, any small lateral displacement will tend to grow exponentially. At 5mm linear separation, a pellet displaced $1\mu\text{m}$ from the beamline will experience lateral force $0.4 \mu\text{N}$, accelerating it at 600m/s^2 : the time constant for the lateral displacement to double $\tau \sim 50 \mu\text{s}$, during which the pellet travels a linear distance of 25 metres.

τ scales as $q^{-1} \cdot d^{1.5}$ where q is pellet charge and d is pellet separation. Although the unwanted lateral acceleration can be calculated from the relative pellet position as seen by the monitoring cameras, hence allowed for in calculating the course adjustment required, it is desirable to keep τ no larger than the approximate time of flight between consecutive course correction stations. If pellet charge is reduced by factor 10 immediately on accelerator exit, a further 3 following the first course correction, and to zero after the final course correction, this is achieved.

REFERENCES

1. A1.5 Fusion Performance, P. Amendt, LLNL-TR-477871 (2011)
<http://e-reports-ext.llnl.gov/pdf/482669.pdf>
2. The Physics Of Inertial Fusion, S. Atzeni & J. Meyer-Ter-Vehn
OUP (2004)
3. Compact, Efficient Laser Systems required for Laser Inertial Fusion Energy;
to appear in Fusion Sci. Technol. (2011)
<http://yag1.kaist.ac.kr/hjkong/>
4. Opacity Calculations Of Low Z Plasmas for ICF, J Rubian et al (2008)
35th EPS Conference on Plasma Physics, ECA Vol.32D, P-2.147
http://epsppd.epfl.ch/Hersonissos/pdf/P2_147.pdf
5. Hypervelocity microparticle characterization, G.C. Idzorek, Los Alamos (1996)
<http://www.osti.gov/bridge/servlets/purl/397131-QIEjdF/webviewable/397131.pdf>
6. The Nanoscale Strength of Ultra Nano Crystalline Diamond, N. Pugno (2005)
Rev.Adv.Mater.Sci. 10 (2005) 156-160
www.ipme.ru/e-journals/RAMS/no_21005/pugno.pdf
7. <http://www.matweb.com/search/datasheetText.aspx?bassnum=MA9WELD8>
8. Field-evaporation from first principles; C. Sanchez & A. Lozovoi
<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.62.9026&rep=rep1&type=pdf>
9. The Dielectric Wall Accelerator, G. Caporaso, Y. Chen & S. Sampayan
(2009) LLNL-JRNL-416544
<http://e-reports-ext.llnl.gov/pdf/377059.pdf>
10. Investigations Of DC Breakdown Fields, T. Ramsvik et al. (2006)
Proc. EPAC 2006, Edinburgh, Scotland
<http://accelconf.web.cern.ch/AccelConf/e06/PAPERS/MOPLS095.PDF>
11. <http://uk.farnell.com/ixys-rf/ixz2210n50l/mosfet-n-rf-754/dp/1347729?Ntt=ixz2210n50l>
12. http://www.ixysrf.com/pdf/switch_mode/IXZ210N50L.pdf
13. Amplitude Systemes online catalog
<http://www.amplitude-systemes.com/>
14. Reactor Concepts for Laser Fusion, W Meier & J Maniscalco
<http://www.osti.gov/bridge/servlets/purl/5232953-TPTXcg/5232953.pdf>
15. Building Costs In Illinois
<http://www.ildceo.net/NR/rdonlyres/2E961958-4860-4526-806B-2705C5F42339/0/BuildingCosts.pdf>
16. Calculating Rail Construction Costs In Light Of The Eddington Report
<http://melbpt.wordpress.com/2008/04/26/calculating-rail-line-construction-costs-in-light-of-the-eddington-report/>
17. Underground Electric Transmission Lines
<http://psc.wi.gov/thelibrary/publications/electric/electric11.pdf>
18. Estimated Capital Cost of Power Generating Plant Technologies
<http://www.jcmiras.net/surge/p130.htm>