Ion beam fusion

Roger O. Bangerter

Phil. Trans. R. Soc. Lond. A 1999 **357**, 575-593
doi: 10.1098/rsta.1999.0342

Ion beam fusion

By Roger O. Bangerter

Ernest Orlando Lawrence Berkeley National Laboratory, University of California,
Berkeley, CA 94720, USA

Fusion faces three types of challenge: scientific, economic, and environmental. Many scientists believe that both inertial and magnetic fusion can meet the scientific challenge in the sense that they can produce net energy. The National Ignition Facility is expected to demonstrate ignition of an inertially confined plasma within a decade. The economic and environmental challenges will remain.

This paper describes research on ion beam inertial fusion. Ion accelerators are, in many ways, ideally suited to the requirements of fusion power production. They can be durable, reliable, and efficient. They can easily achieve the required pulse repetition rates. Cost and beam quality are the principal issues. The cost of a fusion accelerator, if it were built with today's technology, would be acceptable for large power plants (several gigawatt electric). Cost reductions that will allow good economics at smaller plants appear possible.

The environmental issue is activation produced by neutrons from deuterium–tritium fuel. For ion beam fusion, it appears possible to shield the structure of the fusion chamber with a neutronically thick liquid layer. This method of protection greatly relaxes the requirements on materials and endows deuterium–tritium fusion systems with many of the advantages of advanced, aneutronic fusion systems.

The inertial fusion community in the United States has recently proposed a new programme for the development of inertial fusion as a commercial energy source. This paper gives a brief description of the proposed programme.

Keywords: fusion energy; heavy-ion fusion; inertial fusion energy; accelerator for fusion; heavy-ion accelerator

1. Introduction

The worldwide fusion programmes, both magnetic and inertial, have made remarkable progress during the past decade. It now appears likely that both kinds of fusion will eventually produce net energy. The National Ignition Facility described in the paper by Kilkenny *et al*. (this issue) is expected to demonstrate the scientific feasibility of inertial fusion (ignition) within a decade. Unfortunately, ignition is only one step on the road to a commercial fusion power plant. After ignition is achieved, important economic and environmental challenges will remain. We must develop inexpensive drivers (accelerators or lasers) having long life, good reliability, high pulse repetition rates, and high efficiency. We must also develop fusion chambers that are durable, safe, and environmentally benign; and we must learn to mass produce low-cost targets.

This paper is intended to provide an overview of commercial power production based on inertial confinement fusion. It emphasizes accelerators, particularly heavy-ion accelerators, rather than lasers because, in many respects, heavy-ion accelerators

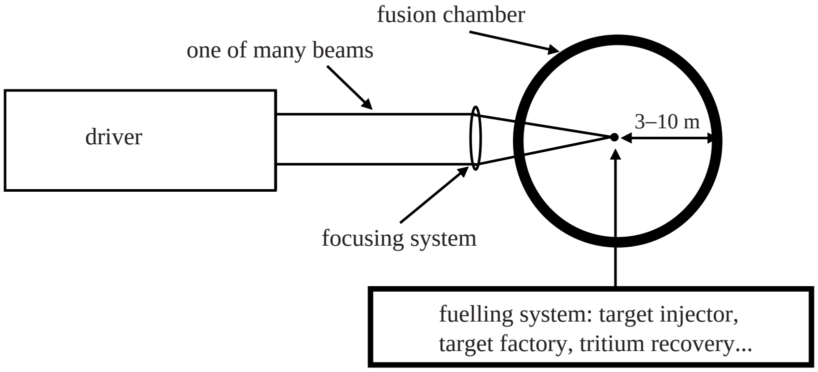

Figure 1. A block diagram of an inertial fusion power plant. The main components are a driver, a focusing system, a chamber, and a fuelling system to produce the targets and inject them into the centre of the chamber.

are well matched to the requirements of inertial fusion power production. Inertial fusion power production is a broad and complex topic. Many details are beyond the scope of this paper. Readers who are interested in details may refer to the literature.†

Sheffield (this issue) discusses several future energy sources. It appears likely that these sources will produce electricity for \$0.05 kWh^{-1} or less. If fusion is to become a major source of energy, it must be competitive. In this regard, it is instructive to compare fusion to fission. Fission is a competitive source of energy that is, in some respects, similar to fusion. The direct capital cost of the best fission plants is approximately \$1 $\mathrm{W}_\mathrm{e}^{-1}$. In this paper, we adopt \$0.05 kWh^{-1} as a goal for inertial fusion energy (IFE). As is the case for fission, achieving this goal demands a capital cost of approximately \$1 $\mathrm{W}_\mathrm{e}^{-1}$. It is important to emphasize that projections and goals for fusion are tentative. By the time fusion is available, the environmental and economic situation may be very different.

2. IFE power plants: science, engineering, economics, and environment

Figure 1 shows the heart of an inertial fusion power plant. It consists of four main components.

1. A driver to provide the energy needed to compress and ignite the target.
2. A focusing system to focus the driver beams onto the target.
3. The fusion chamber.
4. A fuelling system (the targets and a factory to make them, an injector to shoot them into the chamber, a tritium recovery system, etc.).

---

† There have been several international symposia on heavy-ion fusion. The proceedings of these symposia are excellent references for readers interested in the details of topics discussed in this paper. The most recent symposium was held in Heidelberg, Germany, in September 1997 (Barletta *et al.* 1998). The most recent proceedings already in print are those edited by Barnard *et al.* (1996). See also Lawson (1988). Lindl (1996) is a good reference on targets, and Hogan (1995) is a book on nearly all aspects of inertial fusion energy production.

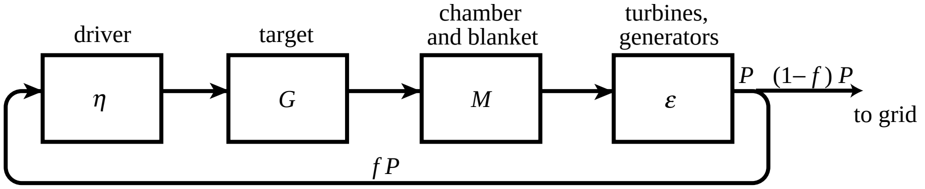

$$fP\eta GM\varepsilon = P$$

Figure 2. Simplified energy flow diagram of an IFE power plant. One can readily see that $fP\eta GM\varepsilon = P$, where f is the recirculating power fraction, P is the gross electrical power, η is the driver efficiency, G is the target gain, M is the blanket multiplication factor, and ε is the conversion efficiency of the system. The net electrical power is $(1-f)P$. This diagram ignores the power required to drive auxiliary systems such as pumps and the fuelling system.

These components must satisfy several important constraints.

(a) *The efficiency-gain* (ηG) *product.* Figure 2 is a simplified energy flow diagram of an IFE power plant. One can readily see that $fP\eta GM\varepsilon = P$, where f is the recirculating power fraction, P is gross electrical power, η is driver efficiency, G is target gain, M is the blanket multiplication factor, and ε is the conversion efficiency of the system (e.g. turbines and generators) that converts the fusion and blanket energy into electricity. The target gain G is the thermonuclear energy produced by the target divided by the energy needed to drive the target. The blanket multiplication factor M is usually close to unity and conceptual power plant designs seldom have ε greater than 0.5. Figure 2 ignores the power that is required to drive auxiliary systems such as pumps and the fuelling system. It is anticipated that this power will be a small fraction of P. To produce net power, f must be less than unity, much less for economical power production. In other words, we cannot afford to use a substantial fraction of the power that we produce to drive the driver. For example, to achieve $f < 0.2$ with $M = 1$ and $\varepsilon = 0.5$, ηG must be greater than 10. The requirement $\eta G > 10$ is often quoted as a requirement for economical power production, but for $\eta G = 10$, only 80% of the power is available to sell. The cost to the consumer can be substantially lower if ηG is large (50–100).

(b) *Targets.* The large hydrodynamic codes that are used to simulate targets have been very successful in predicting the behaviour of target experiments performed using existing drivers. At driver energies of several megajoules and peak powers of several hundred terawatts, these same codes predict target gains of the order of 100 for both laser and ion-driven targets. In addition to megajoules of energy and hundreds of terrawatts of power, there are other target requirements. The beams must be focused to radii of a few millimetres and they must deposit their energy in a mass of the order of 100 mg or less. Furthermore, the beam–target interaction must not produce excessive preheat.

Advanced concepts such as fast ignition, discussed in the paper by Willi (this issue), may lead to higher gains or lower energy requirements. Some of the implications of fast ignition will be discussed § 5. Target cost is an important issue. Today, most targets are built individually. The cost of a single target often exceeds $1000. For power production, the targets must be mass produced and they must be much less expensive. Assume, for example, that a target has a gain of 100 at an input energy of 3.6 MJ. The yield is 360 MJ = 100 kWh. If 50% of the energy is converted to

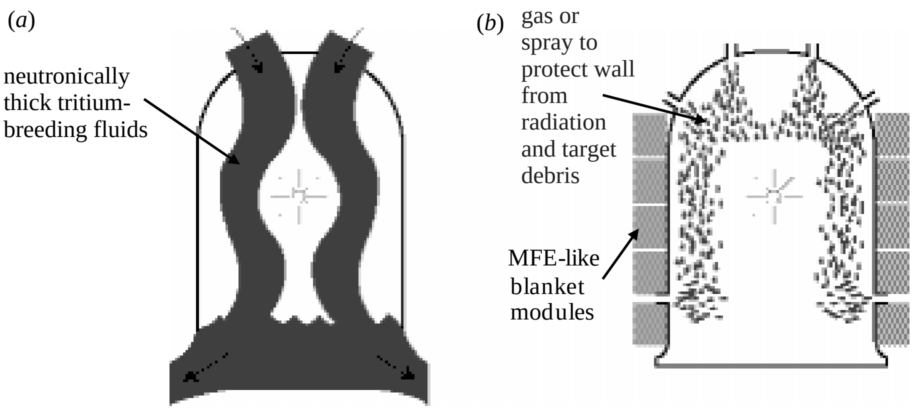

Figure 3. Two principal classes of inertial fusion chambers: (*a*) thick fluids; (*b*) thin fluids.

electricity, we obtain 50 kWh. At \$0.05 kWh^{-1}, the total value of electricity is \$2.50. If the cost of target fabrication is to contribute less than 10% to the cost of electricity, the target cost cannot exceed \$0.25. Target fabrication will be discussed in § 4.

(c) *Chambers.* Scientists who are not familiar with inertial fusion often ask how the fusion chamber can survive the target explosion. Detailed numerical simulations of chamber behaviour indicate that the pressures that are generated are not excessive. A simple estimate gives the same result. The pressure in the chamber is comparable to the energy density. A typical target is expected to have an energy yield of several hundred megajoules. Inertial fusion chamber designs usually have a radius of several metres corresponding to a volume of several hundred cubic metres. Thus, the energy density is approximately $10^6\ \text{J m}^{-3}$ corresponding to a pressure of only 10 atmospheres.

Figure 3 shows two general classes of chambers: (1) chambers with neutronically thick fluid layers (Moir *et al.* 1994), and (2) chambers without these layers. The fluid layers are envisioned to consist of molten metals such as lithium or lithium–lead, molten salts such as lithium–beryllium–fluoride (Flibe), or flowing granules. The layer thickness is of the order of 1 m. The layers serve multiple functions. They protect the first structural wall from neutrons, photons, charged particles, target debris, and high pressures. In deuterium–tritium systems, they also breed the needed tritium. Finally, they serve as a coolant that is passed through a heat exchanger to make steam (or perhaps high-temperature helium) to drive a turbine.

The ability of liquid layers to protect the structure from neutrons is particularly important. In fusion systems without such protection, neutron damage limits the life of the structure to a few years. Calculations show that thick liquid wall protection can extend the life to 30 years or more (Sahin *et al.* 1996). Moreover, at the end of this time, the chamber components could qualify for shallow burial disposal under current regulations in the United States (Lee 1994). Simply stated, thick fluid walls endow deuterium–tritium fusion with many of the advantages of advanced, aneutronic fusion systems.

If thick liquid wall protection encounters an unforeseen engineering difficulty, it appears possible to use most of the chamber and blanket concepts that have been suggested for magnetic fusion. In this case, it is necessary to protect the wall from the intense pulse of photons, charged particles, and debris coming from the target.

Several methods of protection have been suggested. These include thin liquid layers on the chamber wall, liquid sprays (illustrated in figure 3), solid sacrificial shields surrounding the target, and distance. The last option is simple but it leads to large chambers.

(d) *Focusing systems.* The optical elements that focus the beams onto the target must survive in the fusion environment. The optical element closest to the target is subjected to an intense flux of neutrons and usually an intense flux of charged particles, photons, and target debris. It may also be subjected to high-velocity solid particles, liquid drops, and pressure pulses. Since lasers normally employ solid lenses, it has proved challenging to develop concepts for laser optics that might survive in the fusion environment. Ions are usually focused by magnetic fields. The conductors that produce these fields can, in principle, be shielded from the fusion environment. This possibility provides a conceptual solution to the optics protection problem, but detailed engineering studies have yet to be performed.

After the beams leave the final lens, lasers appear to have an advantage. Photons are not charged so there is no mutual repulsion. Ions, on the other hand, interact with each other by producing electric and magnetic fields. This interaction leads to important constraints on accelerator design. This topic will be discussed in § 4.

(e) *Drivers.* In addition to providing beams that satisfy the target requirements described above, a driver must satisfy several other requirements.

There are two economic requirements. The cost of a full-scale driver for a power plant must be acceptable and the cost of the research programme leading to that driver must also be acceptable. The driver is expected to account for about half the total cost of an IFE powerplant. Therefore, to be competitive with other energy sources, the cost of the driver cannot exceed approximately \$0.50 $\mathrm{W_e^{-1}}$. Most conceptual fusion power plants have a capacity of approximately 1 $\mathrm{GW_e}$ requiring a driver cost of less than \$500 million.

The second economic requirement, acceptable cost for the research programme, has become unusually important. The Superconducting Super Collider was cancelled and the International Thermonuclear Experimental Reactor (ITER) design is currently being revised to reduce cost. In response to this situation, the inertial fusion community in the United States has adopted a cost goal of $\$2 \times 10^9$ for an Engineering Test Facility; a facility with capabilities comparable to or exceeding those of ITER.

To be competitive with other energy technologies, the driver reliability must exceed 90%. The driver must have a lifetime of 30 years or more (unless it costs so little that it can be replaced). At a typical pulse repetition rate of 5–10 Hz, a lifetime of 30 years corresponds to nearly 10^{10} shots.

It has proved challenging to develop drivers that can satisfy all the requirements listed above. The paper by Kilkenny *et al.* (this issue) discusses laser drivers. Accelerator drivers, a main topic of this paper, are discussed in the next section.

3. Accelerator drivers

In 1974, A. W. Maschke of Brookhaven National Laboratory and Ronald Martin of Argonne National Laboratory independently suggested that conventional high-energy accelerator technology could be adapted to inertial fusion. The conventional accelerators in existence at that time had already demonstrated many of the requirements for inertial fusion power production. The better accelerators were reliable and

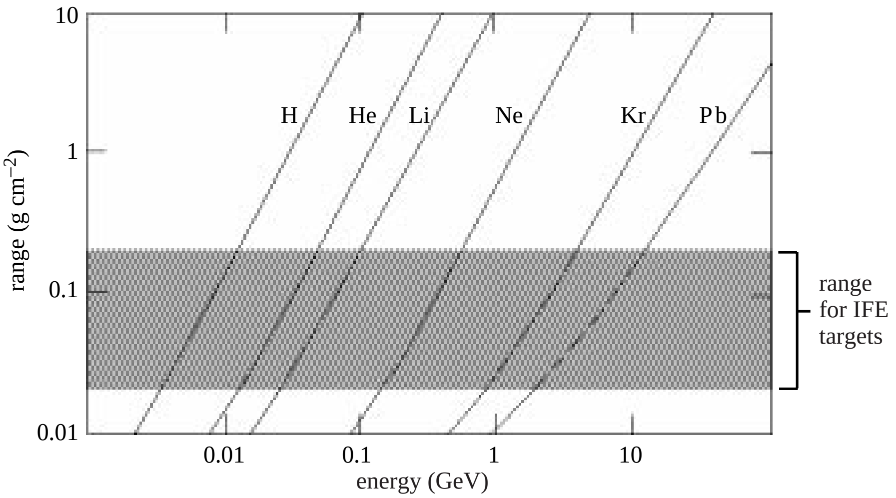

Figure 4. Ion range as a function of kinetic energy for a variety of ions. The shaded band indicates the region appropriate to fusion.

durable. They operated nearly around the clock for decades. They could produce pulses at high repetition rates and they could be efficient. The beams could be focused over many meters to small spot sizes. The larger accelerators produced of the order of 1 MJ of energy per pulse and the Intersecting Storage Rings at CERN had a recirculating DC beam power of the order of 1 TW. Magnetic focusing provided a plausible solution to the optics-protection problem. In summary, conventional ion accelerators appeared capable of meeting many of the requirements described in § 2. Nevertheless, there were some important challenges. The large conventional accelerators accelerated protons or electrons. The penetration depth (range) of these particles is too large for inertial fusion targets and the peak power that had been demonstrated was still two orders of magnitude too low to drive a target. On the other hand, nuclear and particle physicists usually preferred a high ratio of average beam power to peak beam power, so little effort had been devoted to obtaining high peak power.

If the range of the particles is too large, the beams will heat too much mass (more than 100 mg) and the specific energy density (energy divided by mass) will be too small to produce the implosion velocity required for ignition. Maschke suggested that the range problem could be solved by accelerating heavy ions rather than protons. Heavy ions have a much smaller range than protons or other light ions for a given kinetic energy. Figure 4 illustrates this fact and shows that it is possible to use heavy ions having kinetic energies as high as 10 GeV. Unfortunately, the use of heavy ions exacerbates the peak-power problem. During the next few years following Maschke's suggestion, simple estimates, detailed theory, and numerical simulations indicated that it was possible to produce the requisite power using heavy ions. The heavy ions, at a given kinetic energy, are more difficult to focus than protons, but the focusing still appeared possible with some safety factor. Nevertheless, demonstrating that high-powered beams can actually be focused remains the most important scientific issue.

There has been a great deal of confusion in the fusion community regarding the cost of accelerators. When heavy-ion fusion was first suggested, several large accelerator facilities had just been built throughout the world. Funding for fusion was increasing rapidly and it appeared reasonable to propose large, expensive research facilities for fusion or for high-energy physics. In this environment, several of the early heavy-ion fusion leaders suggested that the world should begin building a heavy-ion fusion driver based on conservative design and existing radiofrequency (RF) technology. Projected costs for such a facility were usually greater than 10^9 dollars. The proponents felt that the facility would enable the rapid development of fusion energy and that it would be a good investment. Economic studies suggested that power plants based on the existing technology (Badger *et al.* 1985) could produce electricity economically at large plant capacities (several $\mathrm{GW_e}$).

The philosophy just described was not warmly received by the fusion community. Lasers for inertial fusion were being developed rapidly and the difficulty of the challenges facing laser fusion and other fusion energy options was not widely appreciated. Moreover there was no consensus in the accelerator community regarding the type of accelerator to be built. Denis Keefe of Lawrence Berkeley National Laboratory had suggested that induction linacs might be better for inertial fusion than the more conventional RF accelerators. Large light-ion accelerators based on pulsed power technology were under construction. The light-ion accelerators promised a less expensive route to fusion than conventional accelerators. In addition, much of the target physics was classified. In this environment it was difficult to convince governments and the scientific community to invest in a new, untried approach to inertial confinement, particularly an approach that was expected to cost more than $\$10^9$.

4. Progress in heavy-ion fusion

Despite inadequate funding there have been several noteworthy achievements in heavy-ion fusion. The most modern simulation codes have been used to design a variety of targets specifically for heavy-ion fusion. For example, Max Tabak and Debra Callahan-Miller have recently designed an indirectly driven (X-ray driven) target that gives a calculated energy gain of 68. This target requires 5.9 MJ of 4 GeV lead ions (or other ions having the same range as 4 GeV lead ions). By reducing the size of the X-ray cavity, they have produced a preliminary design that gives a calculated energy gain of 139 with an input of 3.1 MJ of 3.35 GeV lead ions (Callahan-Miller & Tabak 1999). If this design survives further scrutiny, it will enable the construction of smaller, less expensive accelerators. These X-ray target designs allow two-sided illumination (rather than the spherical illumination required for most directly driven targets). Two-sided illumination simplifies accelerator design and leads to attractive chamber concepts. There has been some concern that X-ray targets may be more expensive to fabricate than directly driven targets because X-ray targets are enclosed in a radiation cavity (hohlraum). Fortunately, the hohlraum is relatively easy to fabricate compared to the fuel capsule itself, so it does not add much to the total cost of the target. In fact, it appears likely that directly driven targets will have to be injected into the chamber in a sabot (a protective container) to survive the trip to the centre of the chamber. The sabot may be comparable in complexity and cost to the hohlraum. For X-ray targets, the hohlraum protects the fuel capsule and a sabot is not needed.

Although X-ray targets are not likely to be significantly more expensive to fabricate than directly driven targets, low-cost target fabrication remains an important challenge. This challenge has two parts. The targets contain tritium so the power plant must have equipment that removes unburned tritium from the chamber and extracts tritium from the breeding material. IFE may have some advantage relative to magnetic fusion energy (MFE) in this part of the fusion cycle because the fraction of the tritium burned in a target (approximately 30%) is usually much higher than the fraction burned in a confinement time in magnetic fusion. This fact means that less tritium must be removed from the chamber and recycled. The second part of the challenge is target fabrication. Here magnetic fusion appears to have a clear advantage. While studies (Meier *et al.* 1992) suggest that it is possible to fabricate the precise targets needed for inertial fusion at acceptable cost, far more work is needed.

There has also been significant progress in another important area of target physics, namely the beam–target interaction. There was originally a great deal of concern about this topic because experiments with lasers had revealed an unusually interesting and complex set of phenomena that, under the wrong conditions, led to excessive preheat and poor implosion efficiency. A worldwide effort involving analytic theory, numerical simulations, and experiments has now convinced most scientists that the beam–target interaction for ions is relatively benign (see Hewett *et al.* (1991), and references therein; see also Barnard *et al.* (1996)).

In addition to progress in targets and chamber research, there has also been progress in accelerator research. Lacking the resources to build a large accelerator for heavy-ion fusion, research programmes around the world turned to theory, numerical simulation, experiments on existing accelerators, and small-scale experiments. Given the extensive experience with accelerators for high-energy and nuclear physics, it is important to ask why additional research has been needed to adapt this technology to inertial fusion. As previously mentioned, existing accelerators have not been designed to deliver the high peak power required to drive an inertial fusion target. High peak power pushes accelerators into a previously unexplored region of parameter space.

Accelerators require a focusing structure, usually a sequence of alternating-gradient quadrupole magnets, to prevent unwanted beam expansion due to (1) the pressure arising from the beam's transverse temperature and (2) the pressure from space-charge repulsion. Accelerator scientists usually express beam temperature in terms of emittance, a quantity proportional to the product of the beam size and the square root of the beam temperature. The beam is characterized as emittance-dominated or space-charge dominated according to which pressure is the larger. Space-charge dominated beams are generally more difficult to control because of nonlinearities and potential instabilities. In most conventional accelerators, the beams are usually emittance-dominated. In proposed fusion accelerators, the beams are usually strongly space-charge dominated. To a large extent, this difference is the feature that distinguishes fusion accelerators from more conventional accelerators.

There are three broad classes of fusion accelerators: (1) light-ion diodes, (2) heavy-ion RF accelerators, (3) heavy-ion induction linacs. Although light-ion fusion originally promised low cost, it has proved challenging to focus the beams to small spots. The heavy-ion programme has hoped to circumvent focusing difficulties by accelerating the ions to higher kinetic energy (approximately 10 GeV compared to 30 MeV for light ions). Ions with higher kinetic energy can be focused more accurately. The angular spread among the ions in a beam is equal to the random transverse velocity divided by the directed longitudinal velocity. High velocity decreases the angular spread so the ions can be aimed more accurately. More importantly, far fewer ions are required to achieve the energy needed to drive the target and each ion is less easily deflected from its desired trajectory. At kinetic energies of 10 GeV or higher, it may be possible to focus the beams without neutralizing their mutual repulsive forces with electrons or a plasma. The possibility of unneutralized focusing has been one of the main considerations in setting the kinetic energy at 10 GeV. Even at 10 GeV, it is necessary to use 10 or more beams to achieve the required power.

The RF systems usually consist of an RF linac injecting into several storage rings. The storage rings are required to amplify the linac current of approximately 1 A to the tens of kiloamperes required by the target (10 kA at 10 GeV = 100 TW of beam power). Multi-beam induction linacs can carry large beam currents so the storage rings are not required. Most of the research on RF accelerators has been done in Europe and Japan. The programme in the United States has emphasized induction accelerators

There have been several noteworthy achievements in the RF research programmes. In England, Continental Europe, and Japan, existing accelerators have been used to perform important experiments on a wide variety of issues such as resonances, instabilities, and beam–target interaction physics. These experiments have been augmented by excellent numerical simulation. It is assumed that RF accelerators are sufficiently well understood to perform detailed conceptual design studies of systems for fusion. The HIBALL power plant designs (Badger *et al.* 1985) and a more recent European study of a heavy-ion ignition facility are good examples of detailed studies. In the opinion of this author, these studies have produced some important conclusions regarding the RF systems.

1. It is difficult to produce powerful beams that can be focused to the sizes required for fusion. It is possible to solve this problem if one is willing to build a sufficiently conservative (expensive) machine.

2. The cost of an RF system based on modest extrapolations of today's technology is sufficiently high that economical power production will only be possible at large plant capacities (several $\mathrm{GW_e}$).

The induction accelerator research in the United States has taken a different approach. There are no large ion induction accelerators. Consequently, the experimental programme has performed a set of scaled experiments designed to give some experience with all components of a heavy-ion fusion power plant (except conventional components such as turbines, generators, etc.). Figure 5 is a block diagram of a generic system based on induction technology. It is probably more complicated than an actual system in the sense that it illustrates all beam manipulations that are currently under consideration. For example, figure 5 shows the transverse merging of four beams to form a single beam. Not all accelerator designs require merging. To give the reader a feeling for the scale of the US experiments, figure 6 shows MBE-4 (Multiple Beam Experiment with 4 Beams). It is the largest US experiment to date.

Note that not all the experiments shown in figure 5 are accelerator experiments. In particular, there have been several scaled experiments on the hydrodynamics of fusion chambers. A target injection experiment is currently in progress at Lawrence

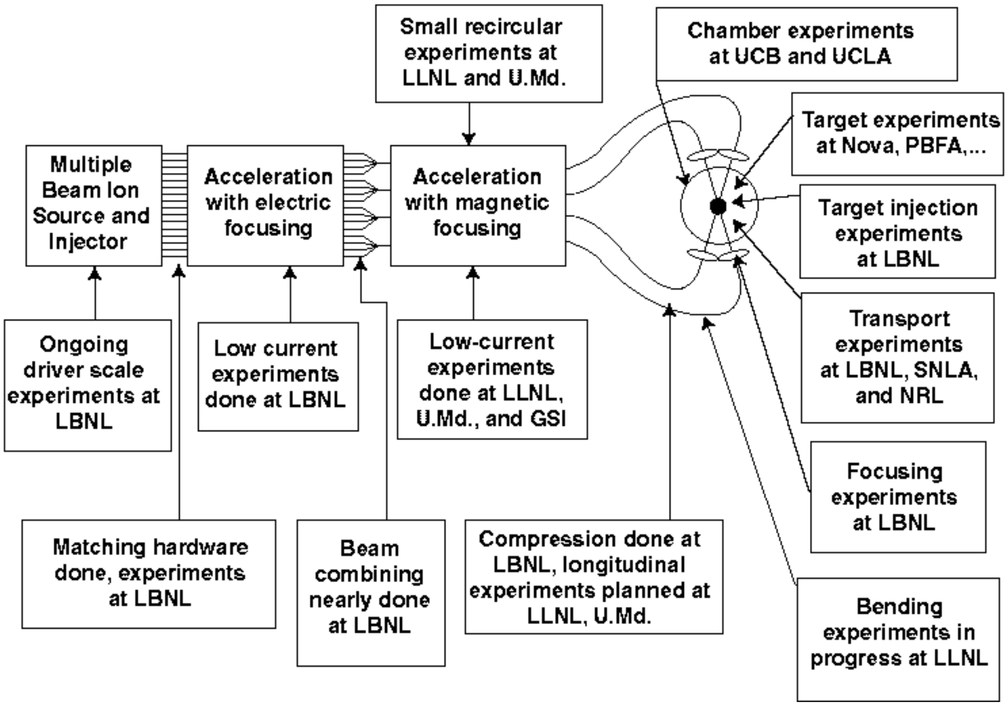

Figure 5. A block diagram of an IFE system driven by a generic induction accelerator. Scaled experiments are addressing nearly all subsystems and beam manipulations have been completed or are in progress at various US institutions.

Figure 6. MBE-4 (Multiple Beam Experiment with 4 Beams). This is the largest US experiment in heavy-ion inertial fusion.

Berkeley National Laboratory. A gas gun propels a surrogate target (no fuel) into a chamber at a speed of nearly 100 m s^{-1}. Photodiodes track the target trajectory and predict its position in the chamber before it actually arrives in the chamber. The current accuracy of prediction is 100 mm, an accuracy that appears to be good enough for X-ray targets.

The experiments illustrated in figure 5 are rapidly nearing completion. Details are beyond the scope of this paper, but the general conclusions are similar to the conclusions regarding RF systems. It appears possible to produce, accelerate, and focus powerful beams, but cost reduction is highly desirable.

Neither the RF systems nor the induction systems presently produce enough beam power to perform definitive experiments on focusing. Building machines capable of definitive focusing experiments is a critical programmatic need.

5. The future

As noted above, the cost of heavy-ion accelerators is an important issue. Calculations show that conservatively designed accelerators can produce electricity economically at plant capacities of several GW_e; but at 1 GW_e, the driver cost should not exceed $500 million.

In light-ion fusion the beams are neutralized by electrons or a plasma as they approach the target. In the early days of the heavy-ion fusion programme, neutralization was poorly understood and the physics appeared difficult. For this reason, among other reasons, most accelerator designers adopted an ion kinetic energy of approximately 10 GeV. The accelerator designers always recognized that the accelerators would cost less if they used lower kinetic energy, a higher-charge state, or both; but these options exacerbate the space-charge problem.

All large accelerators that have been built cost of the order of $\$10^5\ \mathrm{m}^{-1}$ or more. Induction accelerators that have been built to date usually have maximum acceleration gradients of 1 MV per metre. Thus, to achieve 10 GeV with singly charged ions; the machine must be of the order of 10 km long. At $\$10^5\ \mathrm{m}^{-1}$ it would cost $\$10^9$. More detailed estimates put the cost closer to $\$2 \times 10^9$, approximately the same as an RF system and a factor of four too high for a 1 GW_e power plant. There are three ways to reduce cost:

(i) reduce the target requirements,

(ii) reduce the number of metres,

(iii) reduce the cost per metre.

There are potentially several ways to reduce the target requirements. One way, the fast ignitor, has received considerable attention. While the fast ignitor may reduce target requirements, there are, in the opinion of this author, formidable difficulties in using it for commercial power production. Fast ignition is important for laser drivers because without it, the predicted ηG product is marginal. Because of the high efficiency of accelerators, 15–40% in most studies, the higher gain promised by fast ignition is not essential. The principal advantage for ions, if fast ignition can be made to work, is lower driver energy.

Even if one considers standard targets that rely on implosion for ignition as well as compression, improvements in gain are possible. This statement is particularly true of ion targets. It is well known that directly driven targets, ignoring questions of instability, can be far more efficient than X-ray targets. Strictly speaking, the classification of targets into directly driven targets and X-ray targets is not correct for ions. There is a continuum between the two extremes. Fluid instabilities and beam imprinting may prevent the use of targets that are completely directly driven. Nevertheless, it appears that significant improvements in efficiency are possible by moving toward the directly driven end of the continuum.

Unfortunately, there is little experimental information on directly driven ion targets (none at high beam intensities). Evidently, if a new accelerator is built to test accelerator and beam physics issues; it would be extremely useful if it could also address the physics of ion direct drive. Regarding direct drive, it is noteworthy that most directly driven targets require a large number of beams to produce illumination that is nearly spherically symmetric. This beam geometry complicates target chamber design and may even preclude the use of thick liquid jets to protect the wall from neutrons. Because ions deposit their energy over a relatively large distance, it might be possible to design directly driven targets that can accept one-sided or two-sided illumination.

Consider the second method of cost reduction; reducing accelerator length. Two options are: (1) increase the gradient to a value greater than 1 MV m^{-1} or (2) reduce the acceleration voltage (the integral of the accelerating gradient along the machine). Since vacuum breakdown limits exceed 10 MV m^{-1}, it may be possible to exceed 1 MV m^{-1}.

There are several ways to reduce the total acceleration voltage. One can simply reduce the ion kinetic energy, leaving the ion mass high—an atomic mass of approximately 200. From a target standpoint, this option is actually advantageous. Target energy gain increases with decreasing kinetic energy (actually decreasing ion range) until the ion range drops below approximately 0.02–0.03 g cm^{-2}. This ion range corresponds to an ion kinetic energy of 2–3 GeV for atomic mass 200. Therefore, for singly charged ions, one requires a total acceleration voltage of only 2–3 GV rather than 10 GV. If one wishes to use even lower voltages, one must either increase the ion charge state or reduce the ion mass. Unfortunately, lower acceleration voltage always leads to increased space-charge forces in the accelerator, in the focusing system, and in the target chamber. One has always had the option of increasing the number of beams or the beam size to handle the space-charge forces associated with lower acceleration voltages, but both of these methods entail additional cost. The optimal operating point will be a compromise between the acceleration voltage and the cost of the beam lines. Note that the new target designs by Tabak and Callahan-Miller do require voltages less than 10 GV.

The space-charge problems are most severe during final focusing onto the target. The obvious way to overcome the space-charge constraints without increasing the cost of the beam transport system is to neutralize the beams with electrons or with a plasma. Consequently, several years ago, we expanded our research efforts in this area. This research has been quite successful. Numerical simulations show that a small amount of plasma in the target chamber can effectively neutralize ballistically focused beams as they propagate to the target (Callahan 1996). In addition, there are more speculative methods of neutralization that employ current-carrying plasma channels to guide the beams to the target. The channel methods have several attractive features. The chamber pressure can be relatively high (of the order of 1 Torr) so concerns about recondensation and chamber vacuum pumping are minimized. Moreover, the channels have a diameter that is comparable to the target diameter so the beam ports through the chamber wall are small. These small ports greatly mitigate concerns about neutrons, radiation, and target debris. Experiments on channel transport are currently underway at Berkeley and at the Naval Research Laboratory. We are hopeful that at least one of the neutralization techniques will be successful.

Next consider the third method to reduce costs, namely, reducing the cost per metre. We may be able to reduce the cost per metre by reducing the transverse dimensions of the accelerator and also by improvements in technology. In order to reduce the transverse dimensions, we must increase the average current density in the accelerator. Here we define average current density as the total current carried by all the beams divided by the area of the hole through the induction core. Increasing the current density helps in three ways. It reduces the amount of ferromagnetic material that is required for the core. This reduction leads to a reduction in the pulser energy needed to drive the core. A reduction in pulser energy leads to increased accelerator efficiency. We are now studying two options to increase the current density. The first method is to reduce beam size. In a quadrupole (alternating gradient) transport channel, the current that can be transported increases linearly with the magnetic field at the edge of the beam and with channel radius. On the other hand, the area occupied by the channel is proportional to the square of the radius, so a larger number of small beams, at a fixed magnetic field, leads to increased beam current density. Unfortunately, there are limits to this strategy. Inevitable misalignments mean that the beam is not always precisely centred in the transport channel. Clearly the channel must be larger than the sum of the beam radius and the alignment error or the beam will hit the vacuum pipe. The precision of industrially produced parts has improved in recent decades. If this trend continues, it will lead to smaller accelerators.

Ampere's law also places a lower limit on channel size. The magnetic field in a quadrupole is proportional to the current in the conductors and inversely proportional to the quadrupole radius. If one reduces all tranverse dimensions of a quadrupole by the same factor, the cross-section of the conductor decreases as the factor squared while the channel size decreases linearly. There is a limit on the current density in the conductor so quadrupoles scale poorly to very small sizes. Progress in the capacity of superconductors to carry high current densities will enable the use of smaller quadrupoles. However, there may be some economic penalty associated with a larger number of smaller beams, even if the entire beam array is smaller, simply because more parts are required. Minimizing this penalty will require the development of automated fabrication systems.

We have begun working with manufacturers to determine if cost reductions are possible for the major components of the accelerator. The answer appears to be yes if we purchase the large quantities needed for power plants. For example, we have previously used a cost of \$10 J^{-1} for capacitors. The manufacturers inform us that it may be possible to reduce this cost by a factor of five to ten. Similarly, it appears possible to reduce the cost of thyratrons (often used as switches in induction accelerators) by a factor of several to perhaps \$1 to \$2 per megawatt.

Table 1 shows the effects of some of the methods of cost reduction discussed in this section. Option A is a 'standard' 10 GV accelerator. Option B uses lower target requirements, lower acceleration voltage, and less expensive components to reduce length, mass, and cost. It must be emphasized that the numbers given in table 1 are

Table 1. *A comparison of two driver options*

(Option A is a 10 GV accelerator. Option B uses lower target requirements, lower accelerator voltage, and less expensive components to reduce cost.)

	option A	option B
acceleration voltage (GV)	10	1.3
total beam energy (MJ)	6.9	3.3
ion mass (amu)	207	84
ion charge	+1	+1
maximum gradient (MV m^{-1})	1.0	1.5
cost of thyratrons (\$ MW^{-1})	2.00	1.33
cost of capacitors (\$ J^{-1})	10	2
mass of ferromagnetic cores (Gg)	45	13
length of accelerator (km)	10.1	0.96
cost of accelerator (G\$)	2.0	0.56

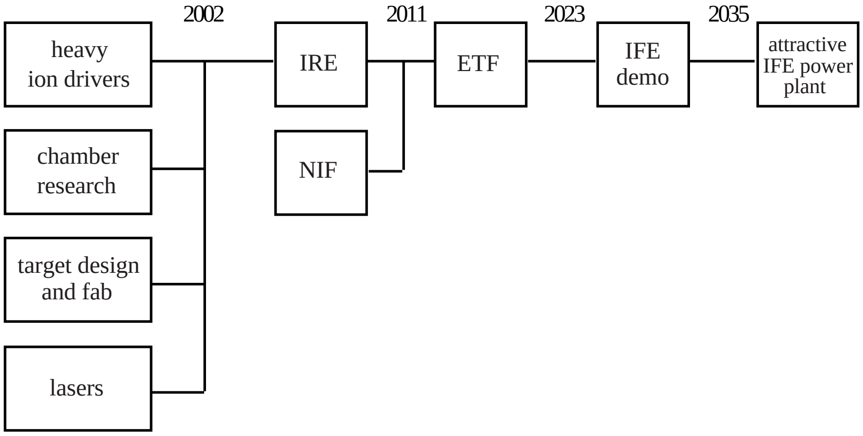

Figure 7. The proposed plan for the development of inertial fusion energy in the United States. See text.

based on simplified computer models (Meier 1999) rather than detailed design. The uncertainties in absolute cost are large but the relative costs are meaningful. To the extent that the cost models are correct, Option B has a cost approaching our goal of \$5 × 10^8. It must also be emphasized that Option B is less conservative in terms of both physics and technology than Option A.

6. Research plans

Previous sections have indicated that heavy-ion fusion is promising. It appears possible to achieve high values of ηG and significant cost reduction appears possible. Nevertheless, larger-scale beam experiments must be performed and less expensive technologies must be demonstrated. The inertial fusion energy community in the United States has developed a comprehensive plan that will address these important issues and ultimately lead to an attractive power plant. Figure 7 illustrates the main features of this plan. The philosophy in formulating this plan was to work backward from the end-point. The final goal is to build the Demonstration Power Plant at a cost of less than \$$3\times10^9$. The Engineering Test Facility (ETF) is a prerequisite. The Engineering Test Facility will be designed to produce large levels of fusion power, but it will not necessarily have the balance of plant (turbines, etc.) required for the Demonstration Power Plant. Based on recent experience with ITER and other large projects, a cost goal of \$$2\times10^9$ has been adopted. It does not appear possible to build a full-scale fusion driver until some of the critical issues relating to cost and focusing have been resolved. Also, since significant improvements in target performance appear possible, using direct drive for example, it is prudent to build a smaller facility to test these target concepts before investing in a full-scale driver. This smaller facility is called the Integrated Research Experiment (IRE). The IRE, together with the National Ignition Facility (NIF), will provide the basis for the construction of the ETF.

Although this paper has emphasized heavy-ion fusion, the proposed development plan also includes lasers. If the proposed plans are accepted by the government, there will be a four year programme to develop both lasers and heavy-ion accelerators for fusion. There will also be a greatly expanded programme in target design, target fabrication, and chamber research. The accelerator programme will have five principal elements.

1. Completion of the small experiments shown in figure 5.
2. End-to-end numerical simulation.
3. Construction of a multi-beam injector with full-scale beams. We currently have an injector with one full-scale beam.
4. A beam transport experiment with one or more full-scale beams. This experiment will be similar to MBE-4, but the beam current will be increased from approximately 10 mA to 1 A.
5. Technology development to reduce costs. The development will focus on ferromagnetic materials, multi-beam quadrupole arrays, insulators, and pulsers.

The estimated cost of the entire heavy-ion fusion programme (including target design, target fabrication, and chamber research) is, for the next four years, less than \$20 M per year. For comparison, the magnetic fusion programme in the United States currently costs more than \$200 M per year.

7. Conclusions

In summary, inertial fusion has several attractive features. Liquid wall protection leads to long chamber life, good economics, and good environmental characteristics. Heavy-ion accelerators are well matched to the engineering requirements of fusion power production. Progress leading to good economics at modest plant capacities appears possible. Finally, because much of the target physics will be done at the National Ignition Facility for defence applications, the incremental cost of a heavy-ion fusion programme is relatively low.

8. Disclaimer

This document was prepared as an account of work sponsored by the United States Government. Neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favouring by the United States Government or any agency thereof, or the Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California and shall not be used for advertising or product support purposes.

This work was supported by the Director, Office of Energy Research (Office of Fusion Energy Science), US Department of Energy, under Contract no. DE-AC03-76SF00098.

References

Badger, B. (and 41 others) 1985 HIBALL-II: an improved conceptual heavy-ion beam driven fusion reactor study. Kernforschungzentrum Karlsruhe Report, KfK-3840 (July).

Barnard, J. J., Fessenden, T. J. & Lee, E. P. (eds) 1996 *Proc. Seventh Int. Symp. on Heavy-ion Inertial Fusion, Fusion Engineering and Design*, pp. 32–33.

Callahan, D. A. 1996 Chamber propagation physics for heavy-ion fusion. *Fusion Engng Design* **32–33**, 441.

Callahan-Miller, D. A. & Tabak, M. 1999 *Nucl. Fusion.* (In the press.)

Barletta, W., Klanner, R., Knoll, G. F., Siegbahn, K. & Ypsilantis, T. 1998 *Nucl. Instrum. Meth. Phys. Res.* A. (In the press.)

Hewett, D. W., Kruer, W. L. & Bangerter, R. O. 1991 Corona plasma instabilities in heavy-ion fusion targets. *Nucl. Fusion* **31**, 431.

Hogan, W. J. (ed.) 1995 *Energy from inertial fusion.* Vienna: IAEA.

Lawson, J. D. 1988 *The physics of charged-particle beams.* Oxford University Press.

Lee, J. D. 1994 Waste disposal assessment of HYLIFE-II structure. *Fusion Technol.* **26**, 74.

Lindl, J. 1995 Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain. *Phys. Plasmas* **2**, 3933–4024.

Meier, W. R. 1999 Systems modeling and analysis of heavy-ion drivers for inertial fusion energy. Lawrence Livermore National Laboratory Report UCRL-JC-130954, 28 May 1998. *Fusion Technol.* (In the press.)

Meier, W. R. (and 27 others) 1992 OSIRIS and SOMBRERO inertial fusion power plant designs. US Department of Energy Report, DOE/ER/541000-1.

Moir, R. W. (and 16 others) 1994 HYLIFE-II: a molten-salt inertial fusion energy power plant design—final report. *Fusion Technol.* **25**, 5.

Sahin, S., Moir, R. W., Sahinasian, A. & Sahin, H. M. 1996 Radiation damage in liquid-protected first-wall materials for IFE reactors. *Fusion Technol.* **30**, 1027–1045.

Discussion

E. A. Little (*University of Wales, Swansea, UK*). With regard to chamber lifetime and the use of thick fluid walls as coolant, etc., it would be possible to incorporate a moderator into the fluid and thereby degrade the neutron spectrum down from 14 MeV to energies of less than 2 MeV, typical of fast/thermal reactor systems. This would increase the vessel lifetime and also permit better choice of materials and design since the behaviour of the materials (e.g. steels) is well known at these energy levels, whereas the behaviour of materials under 14 MeV neutrons is less well known.

D. C. Robinson (*UKAEA Fusion, Culham Science Centre, Abingdon, UK*). What are the relative merits of the European and US approaches to the heavy-ion beam accelerator?

R. O. Bangerter. The approaches are complementary. In general, we believe that the induction (US) approach has advantages, because of the limited number of beam manipulations, in achieving the required phase-space density. The RF (European) approach is based on more mature technology.

P. Vandenplas (*Laboratory for Plasma Physics, Royal Military Academy, Brussels, Belgium*). How constant does the electric charge Z of the ions remain during the acceleration, and what influence does this have on the length of the pulse?

R. O. Bangerter. The charge remains constant during acceleration (+1 in most designs). The vacuum is good enough to prevent ionization. The beam density, at least in the induction approach, is low enough and the time is short enough to prevent significant beam–beam ionization or charge exchange. The beam–beam effects are a concern in the RF systems.

K. Lackner (*Tokamak Physics Division, Garching, Germany*). I would like to stress the large qualitative difference between the reactor studies for alternative magnetic-confinement concepts and inertial-confinement-based systems on one hand, and of ITER on the other hand. Practically all the critical issues of the ITER design, raised also at this meeting—halo-currents and runaways during disruptions, divertor heat loads, compatibility of good confinement with densities around or above the Greenwald density—have emerged as such only during these very detailed design studies. Other, in principle elegant ideas, like direct cooling by liquid breeding materials, had to be dropped once one considered them with the degree of realism associated with an actual design effort. Reactor studies for alternative magnetic confinement and inertial confinement systems are obviously far away from this level of maturity, and one should not compare speculative design advantages of them with the down-to-earth difficulties of a tokamak-based reactor.

R. O. Bangerter. I agree that the ITER design work is very impressive. It is a very important contribution to fusion science and engineering. I strongly disagree with the implication that we should not look for better solutions. The search for better solutions inevitably requires a comparison with known solutions.

A. Gibson (*Bluebonnets, West End Cholsey, Oxfordshire, UK*). Can Dr Bangerter comment on the difficulties of securing sufficiently homogeneous target illumination with heavy-ion drive compared to those of laser drive, both for direct and indirect drive.

R. O. Bangerter. We believe that two-sided illumination is adequate for heavy-ion indirect drive. The required pointing accuracy is a few hundred microns. Laser indirect drive does not appear to have adequate gain for commercial power production. Laser direct drive requires quasi-spherical illumination using many beams. The required smoothness of the beams and the pointing accuracy are far more stringent for direct drive.

C. Gormezano (*JET Joint Undertaking, UK*). Dr Kilkenny has mentioned that the NIF will be achieved with $G = 10$ with 1.8 MJ. In Dr Bangerter's presentation, he has given values of $G = 85$ with 3 MJ input energy. Is the difference due to direct-drive techniques against indirect-drive techniques?

R. O. Bangerter. The target design that I showed gives $G > 130$ from 3.1 MJ of input energy. Both this target and the NIF target are indirectly driven. Between 1.8 and 3.1 MJ, G increases rapidly with increasing energy. More importantly, the target with the higher gain was closely coupled in the sense that it used a small hohlraum. Also the target with the higher gain is driven by ions while the NIF target is driven by a laser. Finally, the ion target design is preliminary and somewhat speculative.

D. R. Sweetman (*The Old Priory, Wallingford, Oxfordshire, UK*). How serious is the problem of dealing with the debris resulting from the micro explosions? There is only of order 100 ms to clear the chamber and some of the debris may be in the form of a gas.

R. O. Bangerter. Calculations and preliminary experiments are encouraging, but much more work is needed. Typical molecular or sound speeds in the chamber are of the order of several hundred metres per second. The chamber size is several metres, so a hydrodynamic timescale is approximately 10 ms, which is short compared to the 100–200 ms between shots.

D. C. Robinson. What is the present position on the interaction between the heavy-ion beam, the target and the plasma, and what is being done in the short, medium and longer term in this area to validate calculations?

R. O. Bangerter. There has been a theoretical and numerical effort in many countries, spanning two decades, to address these issues. There is general agreement that the theoretical uncertainties are small. There are now experimental data from Sandia and NRL in the US and France and Germany in Europe. The proposed IRE mentioned in my paper is expected to validate the conclusions.

G. H. Wolf (*Institute for Plasma Physics, Jülich, Germany*). Dr Bangerter mentioned that the target physics is already well understood. Does this mean that, when modelling the processes leading to ignition, in particular concerning the turbulent 3D processes, empirical scaling assumptions that have to be compared with experimental results are not needed?

R. O. Bangerter. The target physics appears to be reasonably well understood. The understanding is the result of years of comparison between experiments, theory, and numerical simulation. Some processes such as the laser–plasma interaction are still difficult to model based on first principles.

S. Zweben (*PPPL, Princeton, USA*). In tokamak physics we try to check the theory and modelling against experiments, and to identify dimensionless parameters which characterize this physics and how it scales to ignition. Are there any dimensionless parameters in IFE target systems for which the physics has not been checked at the appropriate dimensionless parameters in present experiments? Or will not be checked in the NIF?

R. O. Bangerter. In IFE we also try to check against experiments. As in MFE, there are a number of parameters, not always dimensionless, that we can compare to full-scale ignition experiments. These comparisons have been made. The NIF will address nearly all important capsule physics issues for indirect drive with ions or lasers and for direct drive with lasers.