To: Members of the National Academy of Sciences Committee on the Prospects for Inertial Confinement Fusion Energy Systems, and the Panel on Fusion Target Physics

From: Dr. Stephen E. Bodner, retired, former head of the laser fusion program at the Naval Research Laboratory

Date: December 9, 2011 (revised)

Introduction

When your committee was created, at the request of former Under Secretary Steven Koonin, it was told to *assume* that the NIF (National Ignition Facility) would reach ignition, and was asked for recommendations on how to develop fusion power after success on the NIF. That assumption has become increasingly tenuous. This memo outlines a path to fusion power that does not rely upon the NIF.

Over the past year, Dr. Koonin periodically reviewed the progress towards ignition at the NIF. In his November 8, 2011 memo,[1] he listed some of the remaining problems in the program, and he then noted that:

> *Surprises encountered on the path to ignition make it impossible to predict confidently the rate of progress on those issues of greatest concern to the NIC [national ignition campaign] and so ignition by the end of FY-12 is not assured. It would be prudent therefore to devote some effort to understanding what might be the criteria for, and nature of, a "Plan B" post-FY12.*

At the November 2011 Annual Meeting of the American Physical Society's Plasma Physics Division, there were many discussions in the hallways about "Plan B." For some, it should be a continuation of the schedule-driven approach for another year or two, but using different indirect-drive target designs that could be quickly developed and tested. For others, Plan B would be a multiyear science-based program, first determining the causes of the problems, and then designing different versions of the indirect-drive concept that would have a better chance of succeeding.

If ignition with the NIF is close to success, but for some unknown reason has not quite reached this goal, then of course the government should adopt either or both of the

---

[1] http://fire.pppl.gov/NIF_NIC_rev4_Koonin_2011.pdf

above approaches. However in this memo I will try to show that there is now experimental evidence that the NIF program is very far from success; in fact the indirect-drive approach to ignition is almost certain to fail. Switching to a science-based program will only delay the admission of failure. We all know that in science sometimes there are breakthroughs, and apparently impossible problems are solved, but I do not see any such breakthrough on the horizon.

For some others, Plan B would be some new direction for the NIF: either a shift primarily to non-ignition nuclear weapons research ("high energy density physics"), or a change to a totally different type of target design, such as direct-drive. When the NIF contract was first signed, there was an agreement that the NIF "would not preclude" using direct-drive targets. In reality, it didn't happen. The NIF has insufficient beam smoothing and insufficient laser bandwidth for a proper direct-drive test. Also, the chamber portholes that would be needed for direct-drive were covered up with concrete shielding or allocated to essential chamber functions. Unofficial and rumored estimates from LLNL say that the conversion to symmetric illumination for direct-drive would cost over $300 million and take at least two years. Since the paying customer is the weapons program, it won't happen.

There is a Plan B by the University of Rochester's Laboratory for Laser Energetics (LLE) to test the direct-drive concept using the available chamber portholes, along with an upgrade of the laser beam smoothing and laser bandwidth. Their "polar-drive" approach is potentially feasible, but there are uncertainties and it is premature to evaluate their chance of success. LLE has proposed a program to address these uncertainties.

The NIF program has been heavily marketing their indirect-drive target concept as a means of developing commercial electric power. As this memo will show, not only is ignition highly unlikely using indirect-drive, it would be virtually impossible for their approach to have enough performance for a power plant, even if the target worked as proposed. Simply too much energy is lost in the conversion from laser energy to x-rays, and the target energy gain would be too low. Not to mention the complication of recycling 580 tons of target debris each year.

It now appears likely that the direct-drive fusion concept, developed primarily at LLE and NRL with NNSA funding, could achieve both ignition and high energy gain, *if one used an appropriate laser system*. No one in the direct-drive program would have voluntarily chosen to build a NIF-type laser to test their target designs. The direct-drive concept also has a reasonable chance of meeting all the other requirements for a fusion power plant, including high enough target gains, laser durability and efficiency, total cost of electricity, durable chamber, no ultrahigh vacuums or gigantic magnetic fields, modular plant size well below a gigawatt-electric, minimal debris, etc.

Most people who have seen the direct-drive story understand its inherent attractiveness for fusion energy. To further develop this concept, independent of the NIF, one would use a staged approach, with a set of intermediate appropriations and intermediate goals. Each evaluated before proceeding to the next stage of increased funding. It would be a managed-risk strategy, not a "success oriented" strategy. As I will try to show, it is possible to obtain confidence in the eventual success of laser fusion at modest cost, without first igniting a target. Fusion power can be developed in a reasonable time without another commitment of massive funding.

This memo has the following sections:

- A section on the recent problems with the NIF ignition campaign, as reported at the November Plasma Physics Meeting. This supplements the memo of former Under Secretary Koonin.

- A more general analysis of the scientific and management errors that led to the failure of the NIF. This will provide us with a set of "lessons learned" that should aid in the development of a better fusion energy program.

- A summary of the direct-drive fusion energy concept.

- An overview of laser beam smoothing.

- A review of the LLE plan to improve the optical smoothing on the NIF, to attempt ignition using direct-drive targets.

- An outline of a managed risk strategy for the development of direct-drive laser fusion for a commercial power plant.

My review was undertaken with partial support from the Natural Resources Defense Council (NRDC), and some of the information and conclusions in this memo will likely appear in modified form in subsequent reports of the Council.

Within this memo is some information that has previously been presented to your committee. It is included here again for completeness, so that what I write will be understandable to a wider audience.

Status of National Ignition Campaign

There has been remarkable progress in diagnosing the NIF implosion experiments. The measurements are truly impressive, with some first-rate scientists in the program. There is an overall parameter that measures progress towards ignition, called the "ITF," that has been increasing and is now only a factor of ten away from ignition success. Thus there seems to be almost a disconnect between major progress as measured by the fusion parameter ITF, and the concerns by many scientists that the ignition campaign will ultimately fail. Why the disconnect?

1. The CH ablator is excessively preheated, for reasons unknown. This preheating reduces its density, which then reduces its ability to compress the DT fuel. As Dr. Koonin noted: "*the inflight ablator thickness exceeds predictions by a factor of 2 or so*."
2. The CH ablator is also moving inward at too low of a velocity, and for an unknown reason it slows down prematurely on its flight inward. The drawing to the right plots the center-of-mass (CoM) velocity versus radius. [2] The ablator moves to the left, to smaller radius. The solid lines are the code predictions. The circles and triangles are the experimental data. The red triangles and red curves are for their most recent design, using a silicon dopant in the ablator. Note how in all cases the shell is at too low a velocity, and then prematurely slows down. This prevents sufficient compression and heating of the DT central hot spot that is supposed to ignite.

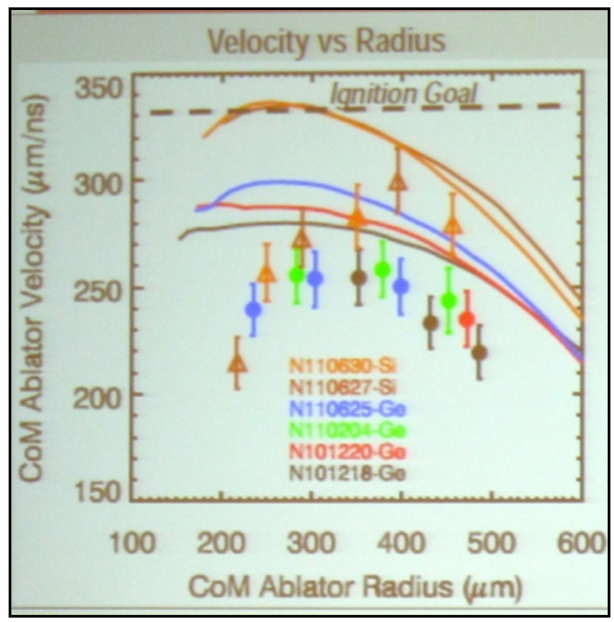

D. Callahan

   The ignition parameter ITF that NIF scientists use varies as V^8, where V is the velocity. To claim that they are within 10% of success, they apparently used the peak value of one red triangle. The energy needed for ignition [3] varies as V^{-6}. That means they are off by a factor of two in required energy delivered to the capsule. Somehow, lots of the laser energy is being diverted and not used to drive the capsule shell inward. Unless that missing energy can be recovered, they will fail.

---

[2] The NIF figures shown here, with the author's name, were presented at the recent Plasma Physics Meeting.

[3] "The Physics of Inertial Fusion" S. Atzeni & J. Meyer-ter-Vehn, chapter 5

3. Looking through the entrance holes at the capsule, measurements show an apparent stronger emission at 900 eV than predicted by their computer model. The reason for this discrepancy is not known.

   All the evidence, (too thick an ablator, too slow an implosion velocity, too early a slowdown, more intense emission from the capsule exterior), seems to point in the same direction. For some reason, a significant fraction of the laser energy is being used the wrong way -- to overheat both the body and the surface of the ablator -- instead of being used to burn off the outside of the ablator and push it inwards like a spherical rocket.

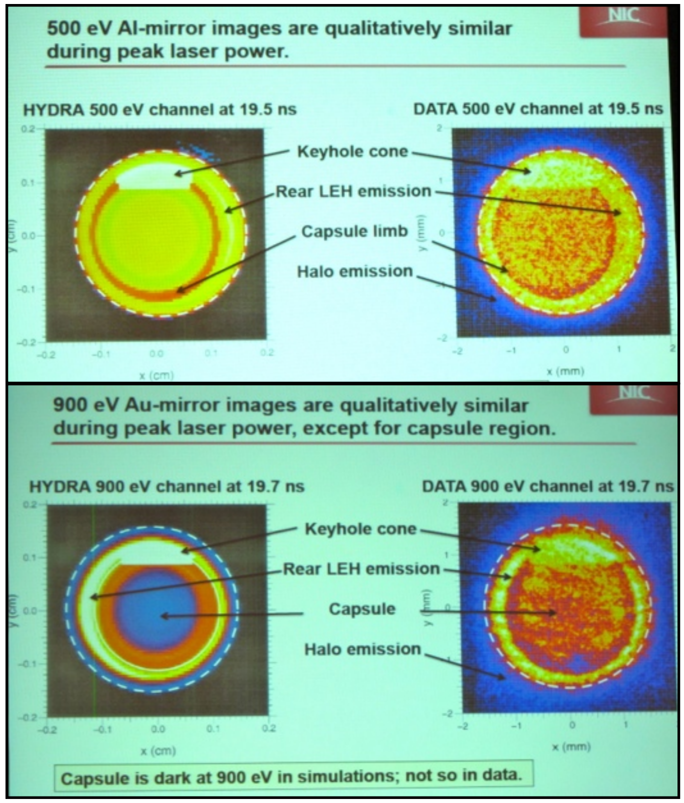

N. Meezan

4. There has been an attempt to fit the implosion data (shock timing, velocities, etc.) with the target design code by artificially adjusting the laser deposition energy. In a nice presentation, it was shown that to obtain agreement they had to assume that somehow the laser power in each of the steps is reduced by 15% to 50%. More evidence that there is some unknown loss of laser energy.

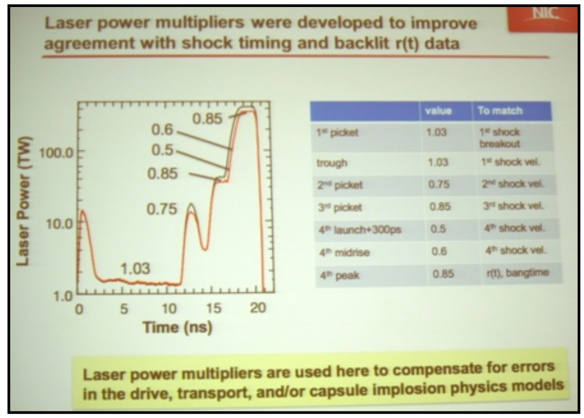

O. Jones

5. If one fits the experimental data on the hot spot ignitor region with a simple analytical model, called the "Isobaric model fit", and then compares these numbers with a computer model that has been "adjusted" to fit the shock times, etc, then the experimental hot spot mass is still about a factor of 8 below the adjusted model, and the pressure is about a factor of 3 below the model. Another indicator of energy losses and preheating.

	DT: N110620 - CHGe		
Hot spot property	Isobaric model fit	Semi-empirical model	SemiEmpir/ Isobaric model
Burn wt. density (g/cc)	26±13	100	~4
Mass (μg)	2.8±1.4	24	~8
Pressure (Gbar)	80±40	270	~3
Yield (kJ)	1.5	27	18
Ion Temperature (keV)	4.4	3.9	$(3.9/4.4)^{4.7} = 0.6$
Waist X-ray P0 (μm)	24	24	1

O. Jones

6. There is now a measurement of the self-generated magnetic fields near the entrance holes, using the Omega laser. About a million gauss. Similar fields are expected on the NIF. What the magnetic fields might be *inside* the hohlraum target, and their impact, is still unknown. The implications were not explored in this presentation.

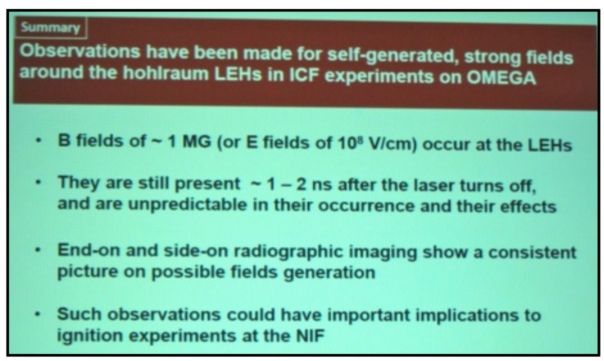

C. K. Li

Recently, about 17% of the incident laser energy has been scattering back out of the hohlraum by the stimulated Raman backscatter instability, equal to about 220 kilojoules. The reflected energy percentage seems to keep creeping up. It was 12% last year, and less than 10% the year before. However this 17% is included in their latest computer modeling, so that doesn't explain the mystery. This Raman instability produces up to 80 kilojoules of electrons at a temperature of about 18 keV; x-rays produced by those electrons could be one of the causes of the mysterious ablator preheat.

Perhaps there is also another Raman instability within the hohlraum that produces scattered laser light that does not leave the hohlraum. This unnoticed laser-plasma instability could also be removing some of the laser energy and placing some of it on or in the CH ablator, either through direct electron preheating, or indirectly by creating high energy x-rays on the hohlraum wall that then preheat the ablator. Such an effect was suggested in another talk at the meeting, by L. Yin and colleagues.

Alternatively, as mentioned above, perhaps some of the incident laser light is scattered by plasma acoustic waves that are near the entrance holes, and that light directly heats the ablator.

There were other interesting problems presented at the meeting that I am not reviewing here. There is ablator/fuel mixing that may or may not be excessive; the analysis they presented seems to me incomplete. There have been attempts to tune the various laser beams to simultaneously control all types of asymmetries during the implosion. They have made progress, but so far they have not succeeded, probably because the hohlraum wall was designed to be too close to the ablator.

Fundamental Reasons for the NIF Failure; and Lessons Learned

1. The NIF laser used a fundamentally new optical design, compared to their previous laser called Nova. The reason was to reduce costs. Ha ha. To test this new NIF laser design, Livermore Lab first built a prototype laser called Beamlet. However Beamlet was not run long enough or at a high enough performance level, to fully address all the problems. Imagine if there had been a requirement that Beamlet operate for a while. Probably the optical damage problem would have been discovered before they began building the NIF. They would have also discovered that the laser was limited to a bandwidth of about 50 GHz at high energies with beam smoothing, versus the 270 GHz in their design specification. Perhaps the NIF project never would have proceeded. Perhaps they would have developed enough justification to eventually proceed. They could have then sought competitive bids for a fixed price contract to build the 192 parallel beam facility.

   <u>Lesson learned</u>: Require that new laser technologies be tested to the extent possible before committing to the full facility. A fusion laser has numerous parallel beams; one can validate the technology by building only a few of the beams.

2. The ignition target design has significantly larger dimensions than their previous targets, and some of the laser-plasma interactions are sensitive to geometrical size and shape. Imagine if the first few NIF beams had been used for extended

experiments, (they briefly were, until 2005, when all testing was strangely cancelled). It is likely that they would have discovered the importance of including dielectronic recombination in their computer modeling. [4] Maybe they would have discovered how really terrible the laser-plasma instabilities could be, since their design is way, way above the instability thresholds, and they could have either stopped or redesigned the laser to pursue direct-drive.

Lesson learned: Try to verify the target physics at each stage of program development. This may appear to slow the program and increase total costs, but sometimes it will improve the program and save money.

3. Through my long involvement in the laser fusion program, from 1970 to 1999, I saw that the computer models would be carefully matched to some experiment, then the experimental parameters would be changed in some way, and the predictions would fail, sometimes disastrously. The physics equations would then be modified to fit the new data. The failure would happen again. This process of prediction and failure happened repeatedly. Gradually, the modeling has become more accurate and more predictive. However no matter how many experiments are performed and calibrated against the computations, there is always the risk of failure when the parameters are extended into a new regime. The failures arise for two basic reasons. First, the physical equations are always approximate. They are not sufficiently complete, and never will be. Second, the equations are solved on a numerical grid, and the methods of solution on this grid are always approximate. Any good modeler of complex hydrodynamics knows these two basic limitations, and should warn others to take the predictions with a grain of salt.

   Lesson learned: Recognize that computer modeling is essential, given the complexity of a fusion target. One cannot understand what is happening without them. Use these models to provide insight and judgement, and to provide a prediction of what is likely to happen. But never place too much confidence in their predictions. Always be wary, and always have a backup plan.

4. An old friend recently reminded me of the early days of the laser fusion program. Edward Teller, who had helped initiate the program, was informed that there was “real” plasma physics in this concept. He then scowled, and finally asserted: “It will never work.”

---

[4] For an explanation of the importance of dielectronic recombination, see my report from last year: http://docs.nrdc.org/nuclear/files/nuc_11010601a.pdf

Lesson learned: Try to design targets that avoid laser-plasma instabilities, by staying below their thresholds. If you can't completely avoid all "real" plasma physics, then try to design close to the instability thresholds.

5. One of the basic rules of engineering design is to put in a safety factor, to allow for the unknowns. Because of its extreme inefficiency, the NIF ignition target had to be designed with a very small safety factor. Then, when things went wrong, there was no recourse. I once estimated how much laser energy would be reasonably required to reach ignition using the inefficient indirect-drive concept. I estimated 50 megajoules.

   Lesson learned: Design both the fusion target and the laser with significant safety factors. Then, if there are some surprises, one can recover.

6. Anyone involved in research knows that written proposals, and regular reports to the sponsor, and meetings with the sponsor, are not a useful way of maximizing the value of a research program. The one technique that works best is competition. With a competing lab, scientists will always perform at their best. In the laser fusion program, Livermore Lab has historically always dominated the program. The pity is that opportunity for competition has always been there with excellent programs at the other NNSA-funded labs (NRL, Rochester, Los Alamos, Sandia).

   Lesson learned: The best way of enhancing the path to fusion energy would be through setting up competing efforts. Then Congress and DOE can relax, confident that the scientists will be working as hard and as creatively as is possible. Duplication in research does not waste money; it maximizes output.

7. I have been thinking back about all of those research reports, and journal articles, and review articles, even books, by the NIF scientists, that confidently described how to design an ignition target. Page upon page of experimental data and equations and computational studies. Even plans for a fusion power plant. Most of it now seen to be fundamentally wrong. And no oversight from the rest of Livermore Lab to bring them back to their senses.

   Lesson learned: I have a set of questions that disturb me. It is not part of the NAS review, but it is bothersome. Are we really going to rely on NIF-trained scientists to maintain our nuclear weapons stockpile in future decades? Can we really trust any part of this lab to maintain our nuclear weapons? Aren't they going to oversell the NIF, since it is such a significant fraction of their lab budget? In the past, when the nuclear weapons were still being designed, the competition between the two design labs was essential. Again, competition is beneficial in research. Do we still need and want two labs just for stewardship and maintenance, given the sad story of the NIF?

Direct-Drive Laser Fusion

This figure shows the design of the optimum direct-drive target, for single demo shots, and for a fusion power plant, consistent with our knowledge of the physics and engineering, using either a KrF gas laser or a frequency-converted diode-pumped solid-state laser (DPSSL). [5]

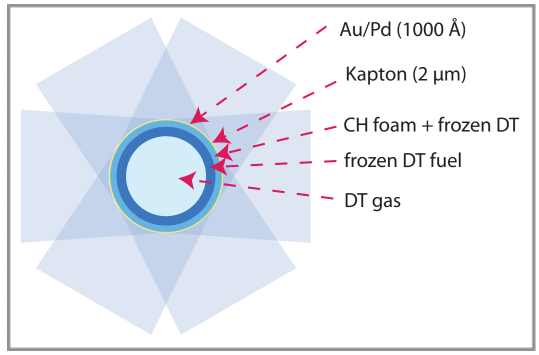

The frozen DT fuel is surrounded by an ablator consisting of a low-density CH foam wicked with DT. Surrounding this ablator is a thin coating of kapton, acting as a seal, then a very thin overlay of gold and palladium, about 1000 Å thick (0.00013 grams), that serves as thermal protection during the acceleration of the target into the middle of the hot reactor chamber. That is the whole shebang.

Knowing that they might be stupid, the designers kept it simple. The Au/Pd coating produces a short flash of x-rays when it is heated, but the CH foam is sufficient to protect the fuel from this early radiation flash. Adding DT to the CH foam improved the ablator's rocket efficiency. The CH foam inhibits the formation of inhomogeneous ice crystals in the frozen DT. Even if the laser burned through the ablator into the main DT fuel, it would not lead to preheat of the remaining fuel.

The cost of each target has been estimated at 17 cents in mass production, which would be about 6% of the electrical energy value of the explosion. If the shot rate is five pps, it would produce total debris, in D, T, H, C, and Au/Pd, of 44 lb/year. (*For indirect drive, with D, T, H, C, plus 1.3 gm of Pb per shot, the total debris would be 580 tons/year.*)

To get a feeling for the energy efficiency and robustness of direct-drive, consider the graph on the next page of calculated target gain versus incident laser energy (target gain is defined as the thermonuclear energy produced divided by the total laser energy that is incident on the target). There are three curves for the direct-drive target. These three curves use different laser pulse shapes to control the characteristics of the target's implosion.

For most of the fusion program history, a target gain of 100 to 150 has been the estimate for the minimum needed for an economically attractive power plant; because of the inefficiencies of the lasers. I believe that this is still a lower bound. The NIF program has

---

[5] A. J. Schmitt, Phys. Plasmas 17, 042701 (2010)

a major sales campaign for a power plant that would use the "indirect-drive DPSSL" concept. Their target gain = 60 design uses optimistic scaling assumptions from their target gain = 1 NIF ignition design. It also requires what I have been told is a surreal estimate of the total DPSSL efficiency.

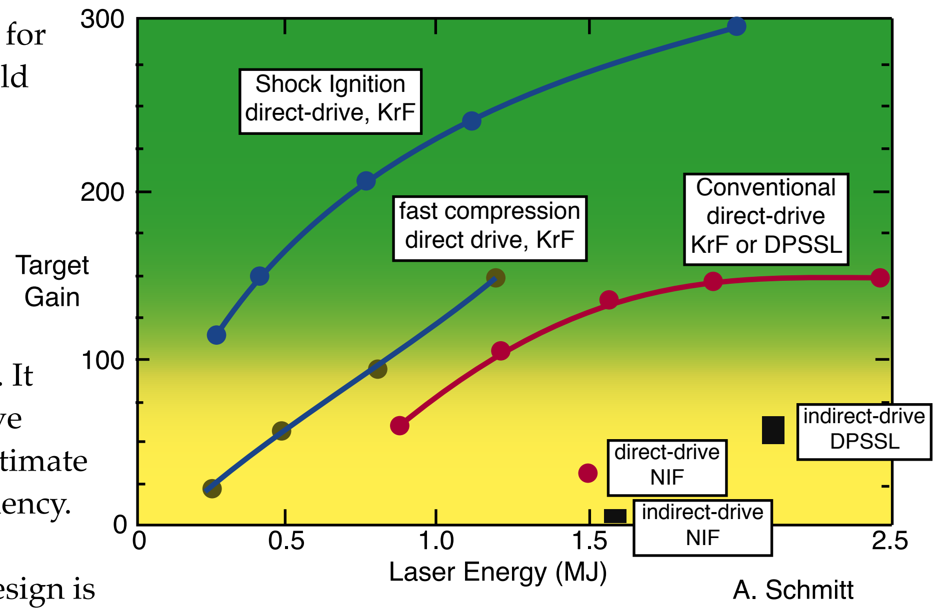

The "shock ignition" design is the most attractive, because it produces the highest energy gain with the smallest laser energy. [6] If it works, one could build a power reactor with, for example, a 0.7 megajoule laser that produces up to 200 MWe net output. A rather compact and low-cost and attractive modular power plant. If this target does not perform as predicted, there is still some room below the three curves.

What about "real" plasma physics? [7] None of these target designs is susceptible to the Stimulated Raman Backscatter instability, the most troublesome of the NIF instabilities. The laser intensity is too low, and the radial plasma distance too short, to induce this instability. The only two instabilities that are of concern for direct-drive are the "two-plasmon decay" instability, and the "cross-beam transfer" mode (both also occur with indirect-drive).

The threshold of the two-plasmon decay mode has been carefully measured experimentally, and it agrees with theory. Some of the direct-drive target designs have a laser intensity above threshold for this instability; but only by a factor of about two. For all laser-plasma instabilities, including this one, its detailed behavior when *above* threshold is complex, and it is still impossible to accurately predict for the plasma parameters of a fusion target. Careful measurements will be needed with a larger laser, and with plasma parameters that are closer to an actual fusion target.

---

[6] The shock ignition design was first proposed by LLE scientists. The version with KrF was designed by A. J. Schmitt of NRL.

[7] Writing of this Section was greatly aided by reading the Sept 20-21, 2011 presentations by A. J. Schmitt of NRL and D. H. Foula of LLE to the NAS target panel

For the cross-beam energy transfer, there is no sharp intensity threshold. The cross-beam transfer however uses an acoustic wave; this acoustic wave grows on the picosecond time scale. If the laser bandwidth exceeds 1000 GHz, it may help quench the acoustic wave, and thus help quench the cross-beam transfer. This hypothesis seems reasonable, but has not yet been tested.

In summary, the best strategy for direct-drive has been to design with as low a laser intensity as is consistent with hydrodynamic stability, and to try to design with some slack, to allow for surprises. Even in the worst case, the targets have a laser intensity that is only a few times above the threshold of the two-plasmon decay mode, assuming good laser beam smoothing.

Using the shortest possible laser wavelength is also helpful. Theory shows that a shorter wavelength directly raises the laser intensity instability threshold. A shorter wavelength also shifts the absorption region to higher plasma densities. This raises the rocket efficiency and the target gain. This also means that less laser intensity is required to produce a given pressure, thus further reducing the risk of plasma instabilities. A KrF laser, with a wavelength of ¼ micron, is significantly superior to a solid state glass laser with a frequency-converted wavelength of ⅓ micron, assuming other important target parameters are held constant.

Optical Beam Smoothing

The most important breakthrough in direct-drive was the invention of laser beam smoothing. This converted the direct-drive fusion concept from something apparently impossible, because of nonuniform laser illumination, to the lowest risk approach to laser fusion. Smoothing is needed for direct-drive because fusion lasers have inherent optical distortions that would produce excessive intensity nonuniformities when illuminated directly on the target. Smoothing is accomplished by trading off the focal spot size for a more uniform profile.

There are two smoothing techniques, called ISI and SSD. These are different ways of accomplishing the same goal. The ISI technique was invented at NRL in 1982 for a glass laser, and then redesigned and simplified for a KrF glass laser in 1987. The ISI method is well matched to a gas laser, but not to a frequency-converted glass laser. In 1988 Rochester scientists responded with the invention of the SSD smoothing technique, which is well matched to glass lasers.

ISI is very simple to implement with a gas laser such as KrF; it merely relays, from the oscillator to the target, a statistically random speckle pattern; when time-averaged it produces a smooth image. SSD is more technologically complex, with frequency modulators and phase plates. There are subtle differences in the final result, and the smoothing of ISI is somewhat superior. The more important difference may be the rate of smoothing, which is determined by the laser's bandwidth. For KrF the bandwidth is about 3000 GHz, for glass lasers about 1000 GHz. Thus the KrF laser smooths three times faster, and it achieves better smoothing for a given smoothing time.

There are a number of uses of optical smoothing:

1. At the beginning of the laser pulse, smoothing with a bandwidth of 1000 GHz or more prevents the initial "imprinting" of pressure nonuniformities on the target. These pressure nonuniformities would be a source for deleterious hydrodynamic instabilities. Later in the laser pulse, after the plasma corona is formed, it is predicted that the corona can itself provide some of this temporal smoothing, and the broad laser bandwidth may or may not be needed.
2. Smoothing can provide a very smooth controllable pressure profile at the longer transverse wavelengths around the target (modes less than 30) where the plasma corona cannot provide sufficient smoothing. This is important all during the inward acceleration. The smoothing is sufficient to implode a direct-drive target with high convergence without shell breakup or excessive distortion. It has been estimated that the residual laser nonuniformities during the acceleration inward should be less than 1% peak/valley, to provide a comfortable safety margin for the implosion.
3. Nonuniform laser illumination can self-focus inside a plasma, because the plasma response provides a positive lens for the laser light nonuniformities. The increase in laser intensity can then accentuate the laser-plasma instabilities; a very dangerous problem. One of the most important discoveries by NRL was that the rapid intensity variations associated with optical smoothing prevents this filamentation. Estimates are still uncertain, but probably a bandwidth of about 100 GHz suffices to prevent this filamentation.
4. With large enough bandwidth, 1000 GHz or more, smoothing may also help limit the cross-beam transfer instability, because the acoustic wave in this mode cannot respond rapidly enough to the changing laser intensities. This fourth justification has not yet been experimentally verified.
5. Ideally, the focal spot size should decrease during the implosion, to better match the size of the imploding capsule. This "zooming" saves significant laser energy and significantly raises the target gain. Zooming also reduces the risk of the cross-beam transfer mechanism, because it reduces the spatial extent of beam overlap. With ISI, this can be easily accomplished at the low-energy front end of the laser. Thus all beams can be zoomed during the implosion, in a series of steps. This has already

been demonstrated on the Nike KrF laser. While zooming is feasible on a glass laser, it is far more cumbersome. Different beams, each with their own set of focal spots, have to be turned on and off during the implosion.

Which laser has the more advanced technology? KrF technology had been considered more challenging than glass lasers. Those challenges have been overcome in the NRL program. Let's compare the energy on target produced by the largest U.S. amplifiers, using each lab's claim of their best performance; no editing. Building many beams in parallel is just more money spent, not evidence of a more advanced technology, so lets compare the energy available from the final laser amplifier.

Name	Location	Type	UV Energy	# of final amplifiers	UV Energy/ amplifier
NIF	LLNL	Glass	1,600 kJ	192	8,300 J/amplifier
Omega	Rochester	Glass	40 kJ	60	670 J/amplifier
Nike	NRL	KrF	4 kJ	1	4,000 J/amplifier

What about the future application to a fusion power plant? Over the past decade Congress funded the High Average Power Laser program (HAPL), managed by NRL. This national program included more than 30 institutions, including DOE national labs, universities, small businesses, and NRL. The program experimentally evaluated lasers, final optics, target fabrication, target injection, target engagement, chamber technologies, and auxiliary systems. The values sought were simplicity, durability, cost, and ability to test on a small scale. The emphasis was on experimental validation. Within the laser category, each year NRL voluntarily provided *exactly* equal amounts of DOE funding to itself and to Livermore. Here are the laser achievements.

Name	Location	Type	Energy/ pulse	Rep rate	Run time between failures
Electra	NRL	KrF	300 J	2.5 - 5 Hz	5 – 10 hrs
Mercury	LLNL	DPSSL	50 J	10 Hz	0.5 – 2 hrs

Electra also ran at 700 J, but for not as long. A technological breakthrough had been demonstrated that would have resulted in further significant improvements on "run time" for Electra, and would have been implemented if funding had not ended.

The HAPL program only partly evaluated the basic challenges of building a fusion power plant. However no failure mode was found in any area, and there was substantial progress and possible paths found around the remaining problems.

If one's goal is to achieve a single-pulse evaluation of ignition and high gain using a direct-drive target, and to minimize the risk, then the government should put ***all*** its money into KrF, because of the four advantages of laser frequency, laser bandwidth, zooming, and better smoothing. There is no counter-justification for a glass laser. However, if the goal is a fusion power system that provides electricity at acceptable cost, then it is too early to know which laser is best. Perhaps both lasers could achieve sufficient target performance, and there could easily be some discovery in the future of a cost or damage problem with a KrF system, such that a DPSSL would become the best choice. We don't know yet, and it is prudent to pursue both options.

Testing Direct-Drive Targets with the NIF

The LLE Omega laser uses "2D-SSD," which means that the interference speckles are smoothed in both transverse directions. The NIF laser has only 1D-SSD, which means that the speckles are only smoothed in one direction.

The Omega laser has a bandwidth of about 1000 GHz, obtained with two sets of crystals. At high power levels the NIF can only operate with about 50 GHz. I was told that if they operated the NIF at higher bandwidths, simultaneously with high power levels and long pulse durations, it would damage the laser.

The Omega has symmetric illumination of the spherical target, from 60 directions. The NIF laser would aim the beams from a set of portholes clustered near the North and South poles; the same ports that were used for indirect-drive. This is called "polar-drive." As noted above, the option of symmetric illumination on the NIF was made impractical because of various decisions during the construction phase.

However the NIF exists now, and if it is not shut down, then it may be available part-time for other uses. So, the LLE scientists have proposed ways of solving the above problems. First they have noted that at early times, during the low-power "foot" of the laser pulse, when the imprinting on the target is most dangerous, the NIF can operate up to about 500 GHz. LLE has proposed to use 1D-SSD at about 500GHz during the foot of the pulse. They also propose to use several simultaneous SSD modulation frequencies. Their calculations indicate that with these two changes, the NIF should have acceptably low imprinting of perturbations during the pulse foot.

In the high power portion of the pulse, LLE has proposed to just reduce the bandwidth, since their calculations predict that the NIF laser beam quality, with just phase plates, would be sufficient to keep the long-wavelength perturbations at a low enough level that the target can be imploded to ignition and gain.

They also think that the filamentation instability might not be dangerous for this type of target design, and in any case they think the two-plasmon decay instability would probably not be driven to higher and more dangerous levels by any filamentation.

Understand that the LLE scientists would not have chosen this version of the NIF, if they had a choice. They would be more cautious. But the NIF is what it is. The LLE scientists plan to first test the above scenario using their own few-beam version of the NIF. Their laser will be modified to match the above conditions, and then used to accelerate a flat foil target. Most likely, other tests would have to be performed later using NIF beams with more total energy on the foil, perhaps 100 kJ.

To deal with the non-symmetric polar drive of the laser beams, and the different refraction of different laser beams by the coronal plasma, they would adjust the power levels between the various laser beams; and for some beams they would change the shape of the focal profile from a circle to an ellipse. Their calculations indicate that this would provide sufficiently uniform illumination.

The LLE scientists have a reasonable basis for their approach, and I can find no flaw in their analysis. They plan to test the basic physics and underlying assumptions every way they can. It helps that NRL has agreed to work with LLE on this. It not only brings an independent assessment, but it adds a bit of that much needed competition.

However the above approach seems to be violating many of the "lessons learned" from the failure of the NIF indirect-drive program. Instead of beginning with simple and symmetric illumination, the illumination is polar and initially nonuniform, requiring careful retuning to achieve symmetry. It is premature to trust the assumptions in their computer modeling of the laser illumination, the target implosion, the light refraction, and the laser-plasma instability behavior. This all has been forced on the community by the previous bad decisions when building the NIF. It appears to be a difficult decision whether or not to proceed with a direct-drive test using the NIF. My recommendation is contained at the end of this memo.

Strategy to Develop a Fusion Power System

As mentioned above, the NIF laser was originally supposed to be brought up in stages, with preliminary target experiments accompanying the laser construction. These experiments were to have been science-driven, not schedule-driven. For some inexplicable reason, that staging was cancelled in 2005. Now consider a different scenario. Assume they had kept the original staging scenario, and assume that the problems were somehow solvable. I assert that even before the NIF laser was finished, fusion experts could have been reasonably certain whether ignition would succeed.

We could not be absolutely certain, because not all the physics phenomenon are perfectly scalable in size, and there would have been no fuel ignition, but we could determine if the odds were good. Perhaps you can now see where I am headed.

Certainly a fusion demonstration is ultimately necessary, to remove remaining doubts by the experts, and to convince the outsider who doesn't understand all the scientific details. Fusion is difficult because some of the physics can't be scaled down too far in size. If one takes a direct-drive fusion target design and reduces its dimensions by, for example, a factor of five, then the laser energy required to implode would be reduced by a factor of $5^3 = 125$. A fusion target designed for one megajoule would then only require an 8 kilojoule laser. It wouldn't achieve ignition of course, because there would not be enough DT fuel. However the scaling also does not work properly because the smaller target would have a thinner plasma blowoff. The laser-plasma interaction with this thinner plasma would be fundamentally different. Smaller targets are helpful to evaluate some of the implosion physics, but they have limitations.

However one can also construct planar foil whose thickness is closer to that of a fusion target. Then the plasma blowoff is optimized for the study of the laser--plasma instabilities above threshold, but one loses the spherical convergence. This planar target is useful because most fusion scientists believe that the riskiest part of an implosion occurs early in time, when the laser first compresses and partly accelerates the shell inward. There are limitations with foils, because the edge effects restrict the distance the foil can be accelerated. With an intermediate-size laser, using large enough targets, or using other simplified geometries, such as a cone, one should be able to properly evaluate the laser-plasma coupling that is not accessible using intermediate-size spheres.

By combining these experiments, spheres and foils and cones, one can investigate most of the physics of fusion with an intermediate-size laser, somewhat smaller than required for ignition and high gain. By "somewhat smaller" I mean approximately 100 kJ. These various tests would not be perfect, but they would be close enough to be very helpful to the experts. If there is a remaining flaw in the direct-drive design, it would probably be found at this energy level. If the direct-drive design is sound, it would probably also be determined.

Some of my friends have been skeptical of fusion energy, and would not change their minds even with a demonstration of high energy gain. Their concerns are elsewhere. They question whether any concept as complex as fusion, which requires billions of dollars and decades to evaluate, would ever become simple enough, and robust enough, and low-cost enough, to be placed in the power grid. Some worry that the minimum size of the facility would be too expensive, with too great a risk of losing an entire investment due to some high-technology failure. A 100 MJ thermonuclear yield has the explosive energy of 50 pounds of TNT -- a large explosion to be repetitively contained in a target chamber. They worry whether the helium nuclei from the burn, buried in the chamber wall, would eventually cause swelling bubbles and premature mechanical failure of the wall material. They worry whether not only the NIF, but all fusion lasers, would have excessive optical damage from the high-intensity laser light and from the explosions. Etcetera.

By the late 1990s the direct-drive target concept had become increasingly attractive. Along with solving the problems of laser beam non-uniformity and hydrodynamic instability, the first of the "conventional" high gain target designs had been developed, with the possibility of a target gain above 100. Thus it seemed time to begin addressing those other nasty problems that concerned my friends. If they were correct, then evaluating these other problems would be the quickest and cheapest way to kill the whole program. My colleagues at NRL and I thus designed the "High Average Power Laser" (HAPL) program, which attempted to study all the above issues, and more. Congress chose to fund this program over the strong objections of NNSA. I retired just after the HAPL program began, and this program was managed by Dr. John Sethian of NRL. It is now generally agreed that HAPL was an extremely successful and cost-efficient program. No failure mode was found, and solutions were partially or fully demonstrated to all the concerns of my skeptical friends. [8] I think they are starting to waver a bit.

It is now time for the next step. The following figure summarizes my suggested plan for direct-drive laser fusion energy. It is a variant on a recent proposal from NRL. [9]

---

[8] J. D. Sethian et al. IEE Trans. Plasma Science, Vol 38, pgs 690-703 (2010);

[9] S. P. Obenschain et al. Fusion Science & Tech., Vol 56, pg 594, (Aug 2009)

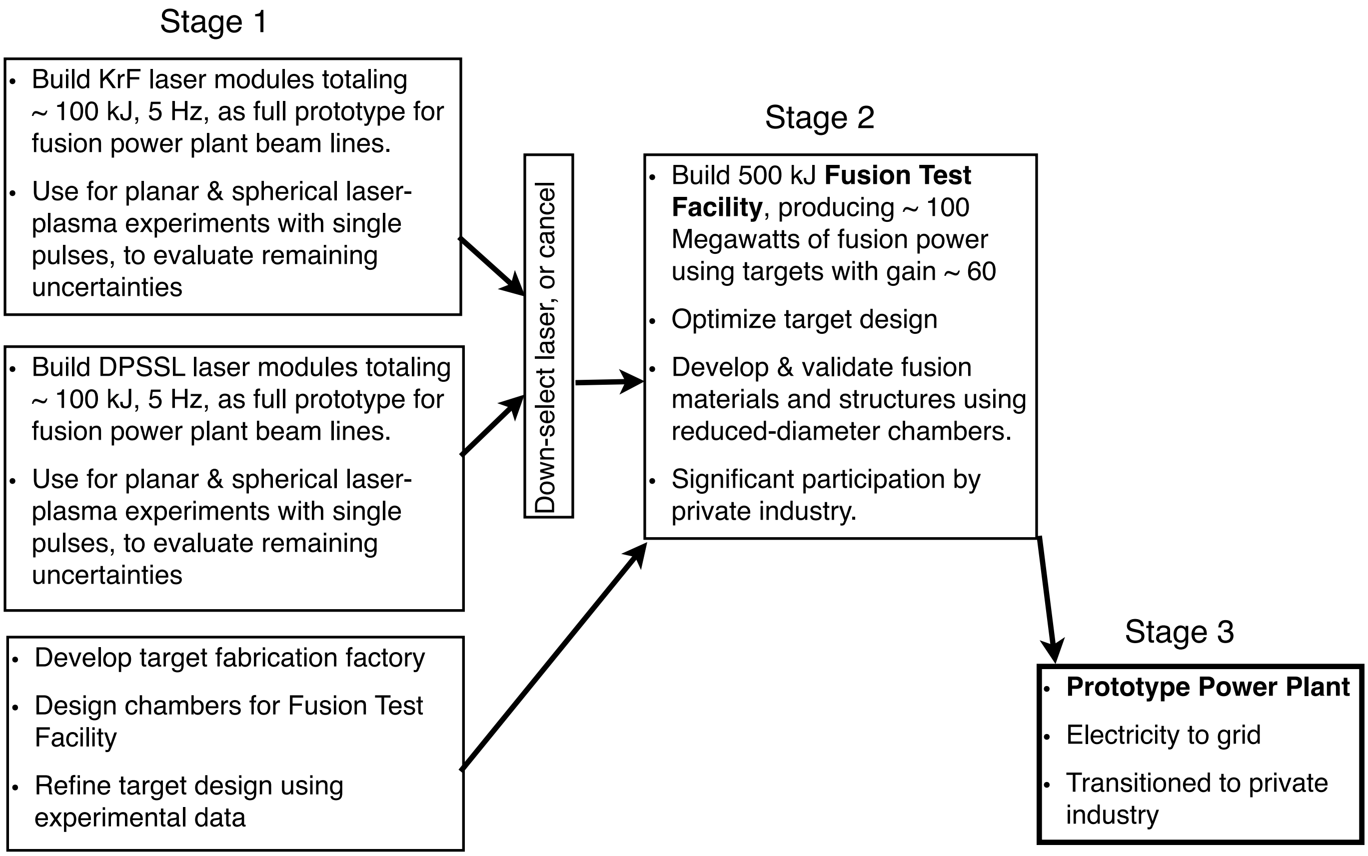

In Stage 1, two competing organizations would obtain exactly equal funding to develop laser modules using a KrF laser and a DPSSL (diode-pumped-solid-state laser). These lasers would have to meet all the requirements for a power plant, with durability, rep rate, beam quality, pulse shaping, etc. Their energies would total about 100 kJ on a target. The lasers would then be used with planar, conical, and spherical targets in a single-pulse mode to evaluate the remaining uncertainties in the laser-plasma interaction, etc. Most likely, NRL would team up with one of the DOE laboratories for the KrF development. Most likely, Livermore would lead the DPSSL development. There would then be a down selection to the best laser. At the same time, other groups would build target factories for Stage 2, and still other groups would design target chambers to use in Stage 2. The schedule for completion of Stage 1 would be determined by the competition. If one laser group falls behind and the other laser group is fully successful, it will be time to down select. If both groups stop making significant progress, or if the target evaluations do not indicate that fusion would be successful, the entire program should be cancelled. If the program proceeds, the experts will now be reasonably confident that fusion would be successful.

Stage 2 would be a Fusion Test Facility, built with significant participation by industry. From the target gain curve on page 13, a 500 kJ laser should be able to produce an energy gain of 60 or more. One could build a laser with less laser energy, but that would not have a sufficient safety factor. First of course one would validate the target performance using single pulses. Then, at 5 pps, the facility would produce about 100 MW of fusion energy. The facility would be used to evaluate chamber wall materials, heat removal, etc. Sufficiently small diameter chambers could be used so that the flux of neutrons, helium, etc., would match that of a power plant.

When this Fusion Test Facility has completed its testing, industry would then build Prototype Power Plants at full explosive yield, and connect to the grid.

I also favor a continued evaluation of the potential of heavy ion accelerators for fusion. Why? Because it is not a laser. Also because it can have a much higher electrical efficiency, and therefore is consistent with lower target energy gains. As I understand it, the heavy-ion community is still exploring possible fusion target designs to put at the end of the accelerator. They also do not yet have an attractive chamber design concept. Finally, they need a development plan, like the above plan for lasers, in which the accelerator performance and the target performance can be evaluated without first committing billions of dollars. The problem is that a heavy ion accelerator is a serial device, not a parallel device like a laser. If one builds only part of the accelerator, then the ion energy would be lower, and its deposition in a target would be different. These may not be insuperable problems, but they need to be solved. This program is not yet ready for major funding. I do not favor holding back the development of a laser fusion power plant to await possible and uncertain development of a complete heavy ion fusion concept. However if lasers stumble, then hopefully a heavy ion fusion concept would be ready to take the lead.

Now let us return to an evaluation of the LLE proposal to modify the NIF to test direct-drive. There are three possible outcomes: (1) the preliminary and final tests are fully successful, and the target gain validates the computer modeling; (2) the preliminary tests up through 100 kJ are a disaster, and no polar drive ignition test is attempted; (3) the preliminary tests are complicated -- not disastrous, but not successful either.

With outcome (1), there should be an immediate start on Stage 1 of the above scenario, to determine if rep rated lasers can meet all of the needed specifications. But one could leave out the testing of laser-target physics, and probably reduce the size of the initial laser development from 100 kJ to about 25 kJ.

With outcomes (2) and (3), one would not be able to determine if the problems were fundamental to the direct-drive concept, or whether they were due to the limitations of the NIF. Perhaps the KrF laser, with its many advantages, would have worked.

There would be no way of determining the answer. Thus, one would still have to proceed with the full Stage 1, to determine if direct-drive laser fusion is viable.

In summary, looking at the program from a scientific viewpoint, not a political viewpoint, the Stage 1 program should proceed regardless of whether there is any further evaluation of direct-drive on the NIF. If the NIF test proceeded and was successful, it would only shift away some of the costs of Stage 1 to the NIF program. If the NIF test failed, it would only delay the Stage 1 development.

From a political viewpoint, proceeding with a NIF direct-drive test is a gamble with unknown odds. If it comes up heads, money would be more likely to be appropriated for Stage 1. If it comes up tails, money would be less likely. Said another way, proceeding with the NIF direct-drive test is equivalent to asking the government to provide hundreds of millions of dollars for the sole purpose of then using the data to convince the government to provide more hundreds of millions. I think it would be better to just shut down the NIF, shut down Nike, shut down Omega, and begin immediately with Stage 1.

The only counter-argument I can think of would be a new claim that ignition using direct-drive is an important component of the nuclear weapons program. I would then wonder if this were actually the Sunk-Cost Trap; the subconscious need of the laser fusion community to justify its past decisions and expenditures on the NIF.

"The inferno of the living is not something that will be; if there is one, it is what is already here, the inferno where we live every day, that we form by being together. There are two ways to escape suffering it. The first is easy for many: accept the inferno and become such a part of it that you can no longer see it. The second is risky and demands constant vigilance and apprehension: seek and learn to recognize who and what, in the midst of inferno, are not inferno, then make them endure, give them space."

Italo Calvino, final paragraph from his book "Invisible Cities"