An Assessment of the National Ignition Facility

This analysis was prepared at the request of the
Natural Resources Defense Council

by
Dr. Stephen E. Bodner

December 15, 2010

Preface

In October 2010, the Natural Resources Defense Council asked Dr. Stephen Bodner to attend the 2010 Annual Meeting of the American Physical Society, Division of Plasma Physics, and provide a written assessment updating NRDC on the technical status and prospects of achieving inertial confinement fusion in the National Ignition Facility (NIF).

Dr. Bodner received his Ph.D. in physics from Princeton University. From 1964 to 1974 He worked at the Department of Energy's Lawrence Livermore National Laboratory in Livermore, CA, first in the magnetic fusion program and then in the laser fusion program. In 1974 he joined the laser fusion program at the Naval Research Laboratory (NRL), and in 1975 was appointed the head of that program. The NRL laser fusion program, funded by DOE, then grew into a large team effort in all aspects of laser fusion: fusion target design, laser design and construction, and laser-target experiments and theory. He is the author or co-author of numerous papers in referred journals, and has been a Fellow of the American Physical Society since 1980. He retired from NRL in 1999, and was a member of the congressionally-mandated 2005 review of the NIF ignition target design.

At NRDC's request, Dr. Bodner has sought to orient his report for readers who have only general knowledge of the NIF target concept, and some general knowledge of science. However, like any specialized scientific field, his discussion necessarily involves a bit of jargon, so there is a glossary at the end of this article.

His report raises questions that highlight the importance of the National Research Council's just launched review of the prospects for Inertial Fusion Energy, and the technical status of alternative approaches that may be less costly and more successful than the path pursued by the current National Ignition Campaign focused on the NIF.

Christopher Paine, Director
NRDC Nuclear Program
December 15, 2010

Background

The NIF scientists are to be commended for their success in building the world's largest laser, and for performing experiments in a scientifically fascinating but challenging environment. However there are physics problems, unique to their approach to fusion, that are becoming increasingly obvious impediments to achieving ignition. These problems are the topic of this assessment.

The first NIF target experiments were reported at the 2009 Annual Meeting. At this year's meeting, the NIF scientists provided an updated analysis of their 2009 experiments, plus initial results from a few laser shots taken so far in 2010. I attended most of the NIF oral presentations, and had a chance to renew many old friendships.

As a reminder, here is a drawing of the current version of the NIF ignition target design. The target consists of a spherical capsule located in the middle of a cylindrical can. The can is called a hohlraum. The capsule is a frozen layer of DT fuel surrounded by a plastic CH ablator doped with germanium.

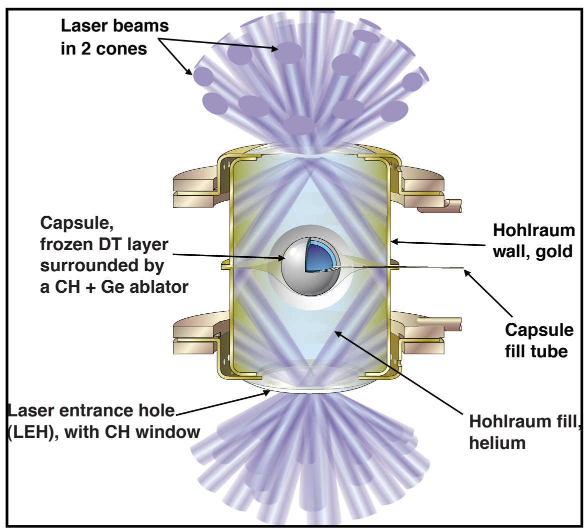

Updated from a figure provided by NIF program

The inner wall of the hohlraum is lined with gold. The region between the can and the capsule is filled with helium gas. At the top and bottom of the hohlraum there are laser entrance holes that are initially covered with a thin layer of plastic (to confine the helium).

The NIF laser consists of 192 beams. Half are focused through each entrance hole. The laser beams illuminate and heat the gold liner, producing a hot plasma. The hot plasma generates x-rays, and these x-rays then heat the outside of the capsule, which drives it inward like a spherical rocket. Most of the DT fuel is kept cold and compressed to high density, while a central portion of the DT heats to thermonuclear temperatures and ignites. This spark plug then propagates into the surrounding cold compressed DT fuel, producing net energy gain.

Introduction

In large organizations of any type, sometimes, but not always, there is a dichotomy between what the senior management says and what is seen at the worker level. The management will say: "Yes, there were some problems before, but everything is under control now, and success seems assured." Meanwhile the workers are facing big problems with no clear solution. Sometimes one should believe the management. They have a better overview of the situation, better judgment, and more experience. Things do work themselves out sometimes, in spite of apparent disasters. At other times the management is arrogant, or disconnected from reality, or too deeply committed to be able to recognize defeat, or too desirous of keeping the money flowing in from investors and sponsors, and the program eventually fails. At the Chicago meeting I saw this management reassurance and worker uncertainty, and the question is, should one buy the optimistic story?

Claim of successful NIF implosions has been overstated

In the Fall of 2009 Livermore Lab began attempts at target implosions using the NIF facility. The first implosions had more pressure on the poles of the capsule than the equator, and the implosions pancaked. However by use of a novel "plasma optical mixing" near the entrance holes of the hohlraum, they were able to shift some of the laser energy from one set of laser beams to another. This rebalanced the pressure around the capsule and produced a roughly symmetric implosion, with total light reflection in the range of 10–15% (*their previous claims of lower reflection have been retracted*).

In the Fall of 2010 the NIF energy was successfully increased to 1.3 MJ, and the hohlraum wall generated an x-ray radiation temperature of 300 eV. This is the laser energy and temperature that they plan to use for their attempts at ignition.

One of the 2009 implosion experiments produced a symmetric-looking implosion when viewed from the equator. [1] However all of the other laser shots, including ones that they claimed were symmetric, show a very blotchy appearance which indicates to me that the capsule is breaking apart during the implosion. What they claimed was a success looks to me like a failure.

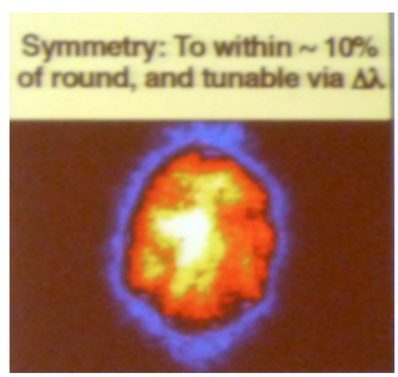

E. Moses, MR1.00001

[1] Most of the figures in this report are photographs from the November, 2010 Chicago meeting. Each is labeled with the name of the presenter, and the talk number.

Their claim of 10–15% total reflection also understates the problem of light scattering. Two-thirds of the NIF laser beams have negligible reflection; for the other third the reflections are in the range of 25–40%, or even higher. This is a big problem.

A surprising failure in the computer modeling

During the 2009 experiments there were various other puzzling results. Besides the unexpected initial pancake shape, they found that the radiation temperature and the backscatter spectrum differed from the modeling; the implosion velocity was too low; the ablated mass was too low; the stagnation pressure was too low; there was low neutron yield; there was unexplainable ablator/fuel mix, etc. They shut the laser down in December 2009, and gradually over the next several months put together a partial picture of what had actually been happening. They described their new understanding at the 2010 Chicago meeting. This was ***not*** some subtle physics effect of interest only to experts. It changed their whole view of how their targets were performing, and of the risks of achieving ignition.

They found that there had been fundamental problems in their computer modeling of atomic physics and of x-ray generation. In technical language, they had not included an effect called "dielectronic recombination". This is standard, well-established physics. As the NIST web site notes: *"Both radiative and the dielectronic recombination are important capture processes which play a dominant role in determining the charge state balance of highly ionized astrophysical and laboratory plasmas."* The target designers just left out this important effect. Also, they had approximated the x-ray spectrum in the hohlraum with only 10 energy groups, far below standard practice.[2]

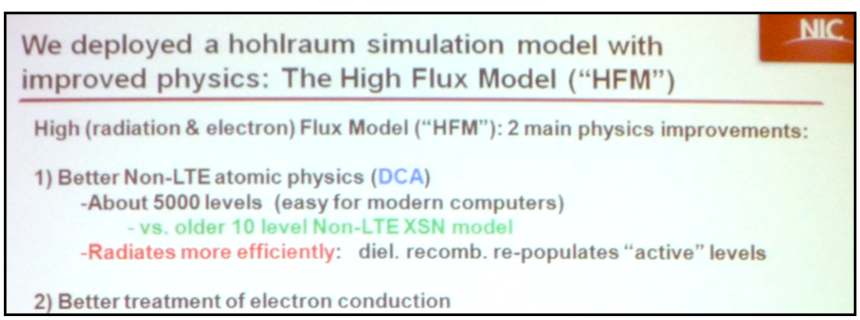

M.D. Rosen, NO5.00003

---

[2] We know that when people have to admit a mistake, they sometimes mumble in a low voice. That is apparent in this figure, which abbreviates the change as "diel. recomb." and makes the phrase subsidiary. However in later discussions with the presenter, and with other NIF scientists, they readily admitted that most of the change in the hohlraum calculations was due to the addition of dielectronic recombination. In their previous modeling they also had to fiddle with the electron flux limiter equation, to match experiments. With their improved atomic physics, they no longer needed that fiddling. Returning to a near-standard conduction model, they state: "2) Better treatment of electron conduction"

Putting dielectronic recombination into their modeling of target performance, increasing the number of energy groups from 10 to a more reasonable 5000, and accounting for the reflection from the innermost laser beams which had not been measured, they then obtained better agreement between modeling and the experiment.

This change in the computer modeling then changed all the important physical parameters. For example, the predicted electron temperature inside the hohlraum went from 4.4 keV to 2.6 keV, a reduction of 40%. This lower plasma temperature enhanced a plasma instability called Stimulated Raman Scatter (SRS) that reflects laser light. This helped explained why, for the "inner" laser beams, the light reflection is now typically 25 - 40%.

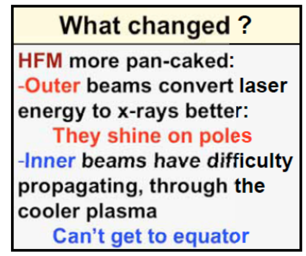

M.D. Rosen, NO5.00003

Finally, and this is important, because of changes in temperature and density gradients, the x-ray emission produced by the gold plasma nearer the entrance holes of the hohlraum was now much stronger than the x–ray emission from near the hohlraum waist. The enhanced x-ray emission near the poles of the capsule, combined with the reduced x-ray emission near the equator, (due to both the enhanced SRS and the modified density gradients), led to the pancaking of the target. One of the implications of these changes is seen in the following figure.

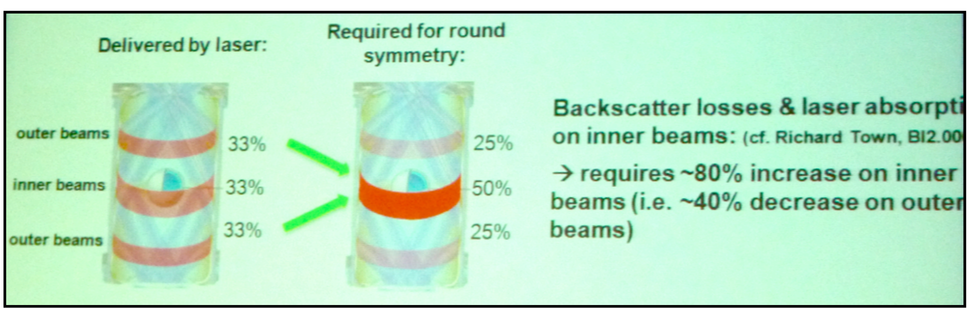

P. Michel, PO6.00015

The NIF laser was designed to deliver 33% of its energy in each of three rings; two outer rings, and a much wider inner ring. They discovered that they needed about 50%, not 33%, of the beam energy *delivered* to the inner ring. Livermore plans to solve their problem by shifting the laser beam energy from the outer to the inner beams using "plasma optical mixing" near the laser entrance holes. The chance of this working sufficiently for target ignition is discussed below.

Surprising level of light reflection, and high-energy electron beams

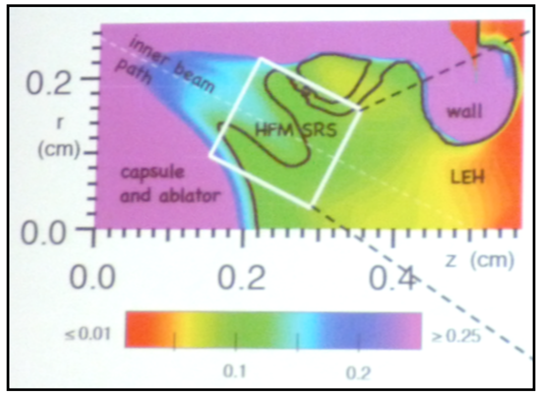

D. E. Hinkel, BI2.00002

Because of the reduced hohlraum temperature, there was a change in both the physical location of the SRS instability and its amplitude. Using their updated computer design code, this figure shows where they think the major part of the SRS is now located within the hohlraum. It is inside the white box near the region labeled "HFM SRS"; now far from the hohlraum wall.

The SRS also produces a forward-aiming high-energy electron beam. The original expectation was for a low level of SRS in each laser beam, produced mostly near the gold wall. The electrons would first hit the wall. Most would be absorbed there, and only a small fraction would scatter onto the spherical target. But now that the high-energy electron beam is produced farther away from the wall, there is a risk that a significantly larger fraction will directly hit the capsule, preheating the DT fuel, and also non-uniformly heating the capsule ablator. And those electrons that do not hit the capsule will now hit other parts of the wall, non-uniformly, and unpredictably. Preheating of the cold DT fuel will reduce its final compression, and perhaps prevent ignition. The nonuniform heating of the ablator will lead to nonuniform ablation, making it perhaps impossible to symmetrically compress all parts of the capsule.

To evaluate this preheat, they presented[3] at the Chicago meeting a new computational model in which the electrons would now be generated within the low-density plasma. However this model made the simplifying assumptions that everything was spherically symmetric, that these high-energy electrons would be created isotropically, and that they would propagate according to a collision-dominated "diffusion model". During the question period, I pointed out to the speaker that, in the low-density portion of plasma, the distance to a collision for these high-energy electrons was long, not short. Therefore their use of a diffusion model greatly underestimated the risk of preheat of the capsule. He responded that the hot electrons will be isotropized by the self-generated magnetic fields produced within the hohlraum.

He is partly correct, if these magnetic fields are oriented so that they prevent flow

[3] J. D. Salmonson, PO6.00010, Chicago DPP mtg

directly towards the capsule [4]. But if the magnetic fields are oriented the other way, then they could preferentially *collimate* the high-energy electrons *towards* the capsule. Of course, the actual geometry is not spherically symmetric, and the electron beams are more collimated than isotropic. Also, it is very difficult to predict or measure where these magnetic fields will be located. No one knows their direction, their amplitude, or how they change during the laser pulse. All we know is that there are huge, nonparallel, temperature and density gradients [5]. These gradients produce a loop current that will then produce huge magnetic fields, probably in the millions of gauss; the largest magnetic fields ever produced. The hohlraum design computer codes do not yet include self-generated magnetic fields. And their new spherical diffusion computer model is not, in any case, an appropriate way of studying high-energy electron preheat.

I hope they are not also using a diffusion approximation for the x-ray transport within the low-density portion of the hohlraum.

What was the justification for ignoring well-known atomic physics?

As noted above, there are many serious ramifications brought on by their previous neglect of dielectronic recombination, and their use of only 10 energy groups. The new calculations reveal increased risk of nonuniform wall heating, of implosion asymmetry, of light reflection, of fuel preheat, etc. What I have not been able to figure out is why they previously ignored this physics. Scientists working in laser fusion research have long understood the importance of including dielectronic recombination. At NRL, the scientists included dielectronic recombination as a matter of course when we developed our target design code in the early 1990s. Other labs did as well. Livermore Lab had a version by 1999.

At the Chicago meeting I asked how this error could have been made. I was told that it would have been computationally intensive to include dielectronic recombination. True, but other labs have included this effect. However let's assume they are correct, and it was previously too computationally intensive for their specific type of target design. Then why their early confidence in ignition? Why build a multibillion dollar laser, locking into what turned out to be the wrong laser and target geometry, when you *know* that you have left out some key atomic physics?

It is also worth noting that this major error could have been discovered much earlier: through asking for help; or through an independent, detailed, and extensive scientific

---

[4] They would still need a completely different diffusion coefficient.

[5] The system is not even non-axisymmetric, because there are a finite number of laser beams.

review process; or through experiments with their previous laser, Beamlet, which they shut down and disassembled prematurely. They did not need a megajoule of laser energy from the NIF to find this basic problem with their target design code.

Plasma Optical Mixing as a Cure for Implosion Asymmetry

Let's now consider their plasma optical mixing, which is supposed to fix the asymmetry problem, if not their other problems. The figure on the right shows the geometry of how the lasers shine on the target. There are actually four cones of laser light, in two groups of two. Measured from the axis, the two "inner" cones are at 23.5° and 30°. The two "outer" cones are at 44.5° and 50°. In total, there are 96 beams in each entrance hole.

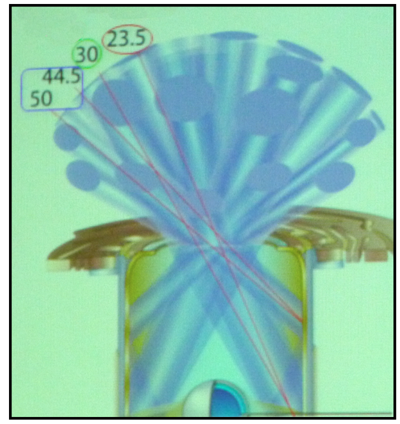

P. Michel, PO6.00015

When any two laser beams cross at any angle within a plasma, the interference between the laser beams will produce an ion acoustic wave. The ion acoustic wave then induces a transfer of energy from the higher frequency laser beam to the lower frequency laser beam. The transfer is strongest when the laser frequency difference matches the (Doppler shifted) ion acoustic frequency. When the two laser beams initially have comparable amplitude, this transfer can best be thought of as a type of "optical mixing", not as an instability. As can be seen, this is a geometrically complex overlap. By carefully choosing the frequency shift between the outer and inner laser cones, and adding a second frequency shift between the two inner cones, they hope to maximize their energy transfer and still minimize the laser reflection.

Let's look at one of the experimental plots presented at the Chicago meeting. The horizontal axis is the laser energy contained in 4 beams (a quad) *after the beam transfer*. That is, they used a theoretical model to calculate and predict the energy transfer, and then added this transferred energy to the initial energy in the beams, to produce the values on the horizontal axis. Most of the variation along the horizontal axis is due to different predicted energy transfers in different experiments, not to actual changes in energy from the laser.

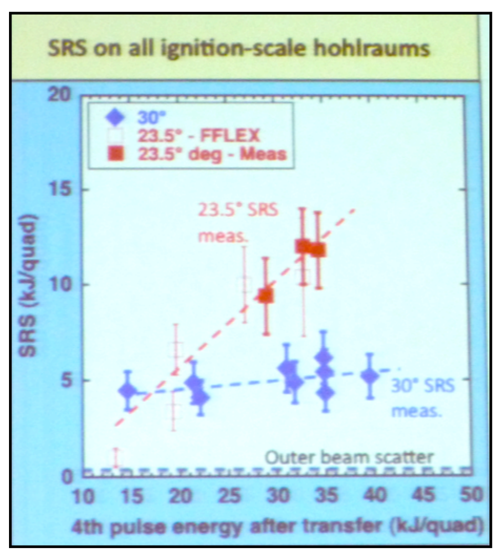

J. Moody, GO5.00001

Look at the data for the 30˚ cone, shown in blue. The SRS reflection is *nearly constant* as they try to increase the energy in this cone. Why? More energy transferred to the incoming 30˚ cone should lead to more energy in the reflected 30˚ cone; but it doesn't. On the other hand, the reflection from the 23.5˚ cone does increase. Something much more complex is going on than what they have modeled. A likely scenario is that the energy did transfer, to some extent, into the 23.5˚ cone, but it did not transfer into the 30˚ cone, or it went somewhere else in the hohlraum. Something is wrong with their theory.

What exactly is wrong? I don't know, and they don't say, but all models are approximate, and there are many possibilities. I did notice that their calculation of the ion acoustic frequency ignores the potentially huge self-generated magnetic fields. These magnetic fields could significantly change the effective coupling frequency, thereby driving the energy transfer out of resonance [6].

Why is all this detail important, as long as they have shown that their technique can improve the implosion symmetry? Because achieving ignition requires better symmetry control than their current experiments, and that symmetry will not happen unless they can fully control the laser energy deposition. One has to prevent not only pancake and sausage-shaped implosions, but also cloverleafs, tulips, bananas, etc.[7]

You don't have to take my word for it; for the past 20 years the Livermore scientists have been regularly explaining that they would achieve ignition by carefully tuning, versus time, the laser energy balance between the inner and outer laser cones. Their assumption was that there would *not* be any unknown spillage, or uncontrolled feedback, or other lack of control of the laser light.

Sort of like tuning your car for maximum fuel economy when there is a leak in the gas tank.

---

[6] NIF measurements near the entrance hole found T_e near 2.0 keV, and n_e near 0.05 n_c. If B is 2 megagauss, then the Alfvén velocity exceeds the ion acoustic velocity.

[7] To achieve ignition they need to simultaneously control the applied pressure asymmetries in both the Legendre P_2 (to less than 1%) and the P_4 and P_6 (to less than 0.5%). Tuning to minimize one mode can enhance another. Tuning symmetry in several modes may now be more challenging.

High energy electron beams can produce three types of implosion failure

To get a feeling for the importance of high-energy electron preheat consider this figure that shows how they now predict the laser energy will be allocated.

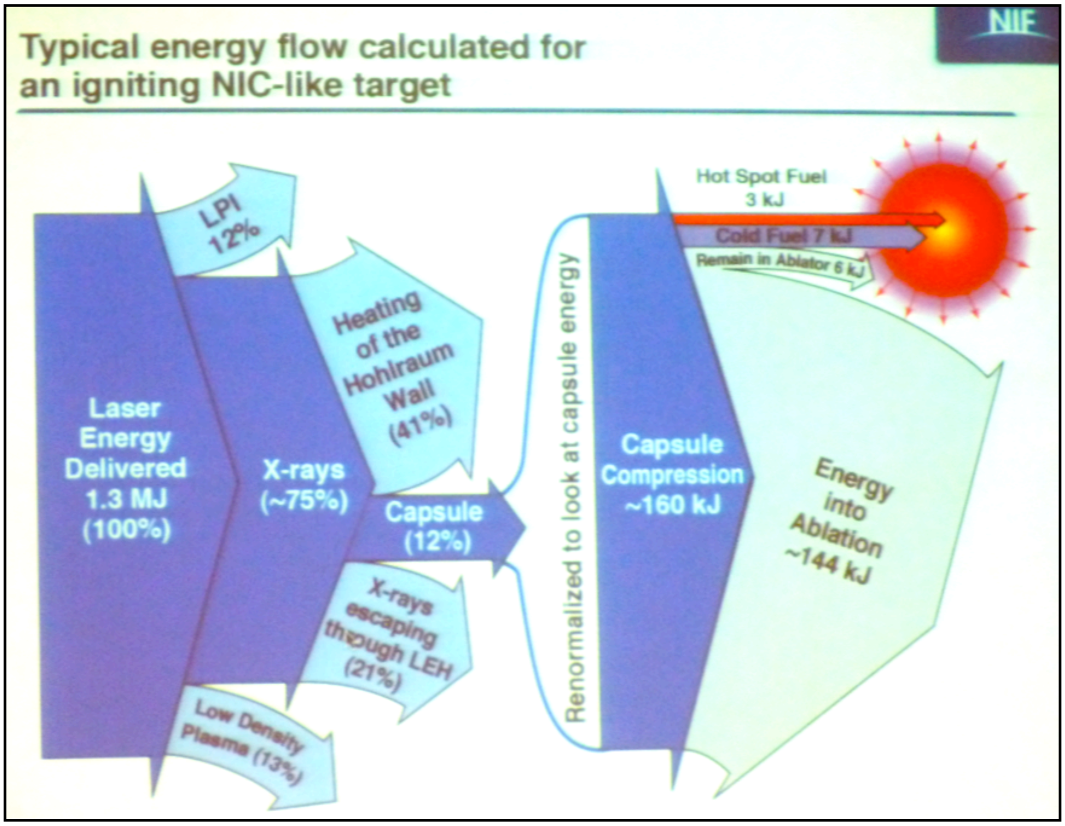

E. Moses, MR1.00001

In this example, they have assumed 1.3 MJ of incident laser light. Of that energy only 160 kJ is deposited in the capsule; and of that energy, most is lost in the rocket ablation process. Only 3 kJ is deposited in the hot spot fuel, and 7 kJ is used to compress the cold fuel, for a total of 10 kJ. Thus only 0.8% of the 1.3 MJ of laser energy goes into the fuel. Ignoring the gross inefficiency of their indirect-drive target concept, they have to get it just right.

Any significant fuel preheat can prevent compression and ignition. The NIF computer modeling predicts that ignition will fail if the DT absorbs more than 1 kJ of high-energy electrons during the high power portion of the laser pulse, or more than 10 J (*not a typo*) during the "foot" of the laser pulse. That's 0.08% and 0.0008%, respectively, of the total laser energy.

Are these numbers achievable? The laser-plasma interaction (LPI), scatters about 12% (156 kJ) of the laser energy back out the entrance holes. This is mostly SRS. Accompanying the SRS is the high-energy electron beam, not mentioned in their graph. Typically, these high-energy electrons will have 60% as much energy as the reflected light.[8] So, we can expect about 100 kJ of high-energy electrons, in some energy distribution. This energy will end up in three places:

1. Part will deposit on the gold wall, but in a nonuniform way that depends on both its initial aiming and on the self-generated magnetic fields. The deposited energy will heat the gold, and produce some x-rays. However since the deposition locations are unpredictable and uncontrollable, this will add to the uncontrollable x-ray illumination asymmetry on the capsule implosion.
2. Part will deposit within the capsule ablator, but again *non-uniformly*, leading to nonuniform shock speeds and nonuniform ablation rates, which could also destroy the implosion.
3. The remainder, with electron energies above about 170 keV, will penetrate the capsule, penetrate the ablator, and preheat the cold DT fuel. If that penetration portion is more than 1 kJ, out of the 100 kJ of high-energy electrons, the implosion will fail.[9]

Their severe laser damage problem may, or may not, have been solved

There was a NNSA/National Ignition Campaign meeting held during the week of this Chicago conference. I did not attend, but I was told that with about 15 people in the room, the NIF manager finally admitted under questioning that in last year's experiments every laser shot, even the low energy laser shots, had produced significant damage to the final optics. That may explain why they unexpectedly shut the NIF laser for about 8 months this year, after a very short operating period last year, and replaced a large fraction of the final optics. The people in this meeting were assured that this optical damage problem has now been fixed. The NIF laser team deserves praise for this important accomplishment. However it would seem prudent to wait for a formal written explanation of why the laser was so heavily damaged last year, and how they fixed it, along with detailed evidence that the fix worked this time.

---

[8] L. Divol, NO5.00013

[9] They already have indirect and uncertain estimates of fast electron generation. They announced that they intended to directly measure the preheat in the capsule immediately after the Chicago meeting. Since they have ablator/fuel mixing, this measurement may be difficult.

Earlier attempts to control laser-plasma instabilities

The reader probably knows that there are two categories of laser fusion targets: direct-drive, where the laser directly illuminates and heats the capsule ablator; and indirect-drive where the laser light shines into a hohlraum, produces thermal x-rays from the wall of the hohlraum, and these x-rays then heat the capsule ablator.

Back in the 1970s, laser fusion appeared impossible because of the large levels of laser-plasma instabilities, such as SRS. Then there were two major advances in the 1980s that changed everything. One advance was the use of the highest possible laser frequency, in the ultraviolet; either a frequency-tripled glass laser or a KrF gas laser. This frequency increase dramatically raised the threshold of the laser-plasma instabilities. The other advance was a set of inventions that provided ways of producing ultra-smooth laser beams, so that one could seriously think of directly and uniformly illuminating the capsule, instead of indirectly via x-rays. As a bonus, these smoothing techniques also dramatically raised the effective threshold of the laser-plasma instabilities.

The experimental results were so good that one could now design direct-drive fusion targets that operated near the threshold of the various laser-plasma instabilities, with enough energy gain for an economical pure-fusion reactor. Fluid instabilities would still need to be controlled, but there would no longer be a need to worry about the laser-plasma instabilities. Delightful!

All of these advances were invented, developed, tested, and modeled at the various laboratories that work on direct-drive laser fusion. None of the advances originated from the NIF program or its predecessor. These advances are also not obviously applicable to the indirect-drive target concept. The problem is that to fit all of the laser beams into the small entrance holes of a hohlraum, one has to use much higher laser intensities -- way above the threshold of the laser-plasma instabilities. This higher intensity light also propagates through longer distances in the plasma. With higher intensities and longer path lengths, there is simply no way to use these techniques to get down near the instability threshold.

The Livermore scientists insisted that they too could use the inventions that were developed for direct-drive, and they worked on this until the mid 1990s. They of course understood that they could not get their ignition design to near the instability thresholds, but they thought they could keep their instabilities weak enough to obtain ignition. They eventually claimed success, although some disagreed. Congress and DOE then gave billions to Livermore to build the NIF. We were told the NIF was necessary to attract the best possible scientists to Livermore, so that they could be trained as our next generation of nuclear weapons experts, in the name of national security.

Well, the laser has finally been completed, many years late, many billions over budget, and has not yet reached its full performance requirements. What do we find? Huge levels of the SRS laser-plasma instability, with about 25 - 40% reflection from the inner laser beam cones [10], plus unwanted energy in the directed beams of high-energy electrons. So the direct-drive inventions did not prove sufficient for indirect-drive.

Until the laser was completed, Livermore scientists planned to use a very wide laser bandwidth, which is part of the beam smoothing technique, to help control their laser-plasma instabilities. It was even in the NIF specifications. However when they found that the NIF laser could not produce the widest bandwidths, they dropped this approach and it is no longer mentioned. I did not think a wide bandwidth would fix their laser-plasma instabilities, but I was curious to see a test.

Shifting their laser to green light is not a solution to their problems

With their poor results, it is worth asking why exactly did they shift the laser wavelength to its third harmonic, in the ultraviolet? Since they can't control the laser-plasma instabilities, which are at very high levels in the inner beams, why not just use the second harmonic, in the green? With green light, they will probably have a much reduced level of optical damage to the laser, and they can then produce much more laser energy, perhaps 3 MJ or more. They will certainly have intense laser-plasma instabilities, but they never got rid of them with the ultraviolet. So why not shift to the green?

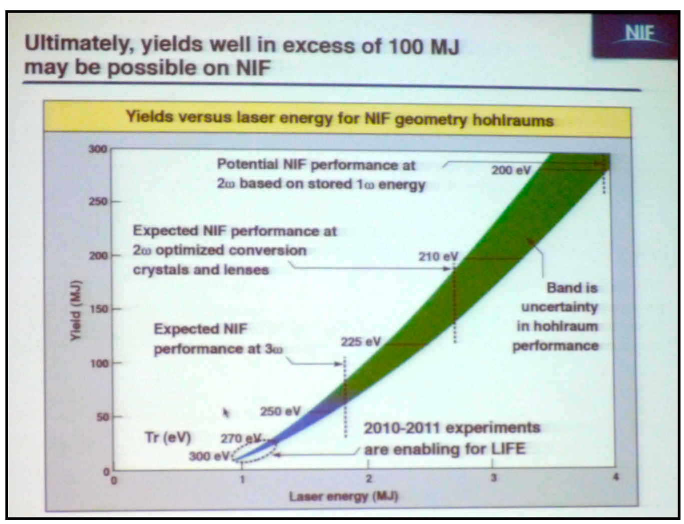

E. Moses, MR1.00001

It turns out that they are considering just that. This figure was also presented at the Chicago meeting. Laser energy is plotted on the horizontal axis. Notice that just a little bit of the curve, near the bottom, is colored blue/violet. Most of the curve is green. I was told that Livermore has been internally discussing what they will do if their

[10] N. B. Meezan et. al., Phys. Plasmas, Vol. 17, 109901 (2010)

current approach fails. They are seriously considering converting the NIF to the green.

Going to green laser light is, to put it bluntly, a dumb idea. The target performance will worsen because all the instabilities are worse at a lower laser frequency.

Summary

When the NIF program was being planned, the target designers understood that there were uncertainties in their computer modeling; perhaps extra backscatter, perhaps some unexpected wall motion, perhaps errors in some physical parameters. So they requested a laser with 1.8 MJ of energy, when only 1.5 MJ would be fired onto the target. The extra 20% of energy would be held in reserve. The 1.5 MJ ignition design even included 10% reflection from each of the laser beams, so that the target actually absorbed only 1.35 MJ. If there was more than 10% reflection in any of the laser beams, they could simply use the reserve power in those beams to compensate for the reflection.

The NIF laser has so far only produced 1.3 MJ, so there will not be any reserve energy unless they have indeed solved the optical damage problems. They even had to scale down the size of the ignition target. Extra laser energy would help, but not enough to compensate for both 25 - 40% backscatter and the reduced x-ray emission from the equatorial region of the hohlraum.

Their design concept also assumes that the laser beams will deposit their energy close to where they are aimed. As explained above, the plasma optical mixing is not working as expected, and there is the likelihood that high-energy electron beams are spraying energy onto parts of the hohlraum wall, not in the nice isotropic way that they are still assuming. And they have underestimated the potential of high-energy electron deposition into the ablator and into the fuel.

In addition, the capsule and the surrounding plasma are encased by a gold hohlraum that limits diagnostic access. Some quantitative and qualitative data is available from x-ray emission and light backscatter, but they can not directly and quantitatively measure the local wall emissivity, the local density and temperature profiles, the magnetic fields, the internal plasma velocities, light scatter and refraction into other angles, etc. And there may be still more problems to uncover in their computer modeling, since the implosion velocity is still below predictions, the ablator is still mixing with the fuel more than predicted, etc.

Imagine buying some very expensive, custom-made computer gaming unit, which regrettably arrives years late, and with all kinds of charges on the invoice that weren't there when you placed the order. Unfortunately, the unit has both hardware and software problems. When levers are moved, unexpected things start changing, contrary to what was specified in the manual. Also, a portion of the screen is blank, so part of the action is unobservable. Also, it needs lots of routine but expensive repairs. Unfortunately, there is no warranty and no refund. The manufacturer tells you not to worry — just send more money for an upgraded unit that will be sent to you whenever it's ready. The manufacturer swears he is making progress, and is optimistic that the problems will someday be fixed.

Glossary

Term	Explanation
Alfvén wave	A low frequency ion wave in a plasma with a magnetic field. The magnetic field acts like rubber band and provides a restoring force.
Beam smoothing	A combination of spatial and temporal averaging that trades off some of the focusability of laser light to produce highly-uniform focused laser light, without hot spots. Rapid temporal averaging requires a wide laser bandwidth.
Capsule	The spherical target with an inner layer of DT fuel. It is located in the middle of the hohlraum. See drawing on page 3.
Dielectronic recombination	This is a process in which an electron is captured by an ion without radiation, with the extra energy absorbed instead by another electron already in orbit around the ion.
Diffusion	A spatial spreading of particles (such as electrons) that is dominated by random collisions. The textbook example is the spread of heat.
Doppler shift	Waves traveling through a plasma have an apparent frequency shift when the plasma is moving.
Energy groups	In a computer program, the continuous energy spectrum is approximated by a finite number of groups.
Ion acoustic wave	A low frequency ion sound wave in a plasma.
Isotropic	Uniform in all directions.
keV	Kilo-electron-volts. Thousands of electron-volts. A unit of energy for particles, It can also be unit of temperature, for radiation or for particles.
Legendre polynomial	In linear systems, waves are mathematically characterized by sines and cosines. On a sphere, the analogous waves are characterized by Legendre polynomials.
MJ	Megajoules. Millions of joules. A unit of energy.
NIST	National Institute of Standards and Technology
SRS	Stimulated Raman Scatter. A laser wave decays into a backward traveling laser wave and a forward electron plasma wave. The electron plasma wave then produces a forward high-energy electron beam.
Thermal x-rays	A broad spectrum of radiation from a hot body is called thermal. For room temperatures it is mostly in the infrared. At fusion temperatures it is x-rays.