UNIVERSITY OF ROCHESTER
LABORATORY FOR LASER ENERGETICS

LLE IN FOCUS

SPRING 2026 | ISSUE 9

FOCUS ON INERTIAL CONFINEMENT FUSION

SHAPING THE FUTURE OF DIRECT DRIVE

LABORATORY FOR
LASER ENERGETICS | University of Rochester

LLE IN FOCUS

Editor in Chief
Dustin H. Froula

Technical Editor
Valeri N. Goncharov

Managing Editor
Alison Arnold

Content Editor
Rosemary Shojaie

Creative Director
Jenny Hamson

Art Director
Michael J. Franchot

Editorial Photographer
Jacob I. Deats

Copy Editor
Jennifer L. Taylor

Graphic Design
Jacob I. Deats
Lamisa Fairooz
Rodi Keisidis
Heather S. Palmer

Publications
John M. Bobilin-Stahl
Jennifer C. O'Brien
Lisa A. Stappenbacher
Jessica M. Weart

Website
William Hall, Jr.

Additional image contributions by
J. Adam Fenster
Eugene Kowaluk

LABORATORY FOR LASER ENERGETICS | University of Rochester

Contact us:
www.lle.rochester.edu/about/contact/

250 E. River Rd., Rochester, NY 14623
www.lle.rochester.edu

READ LLE IN FOCUS ONLINE

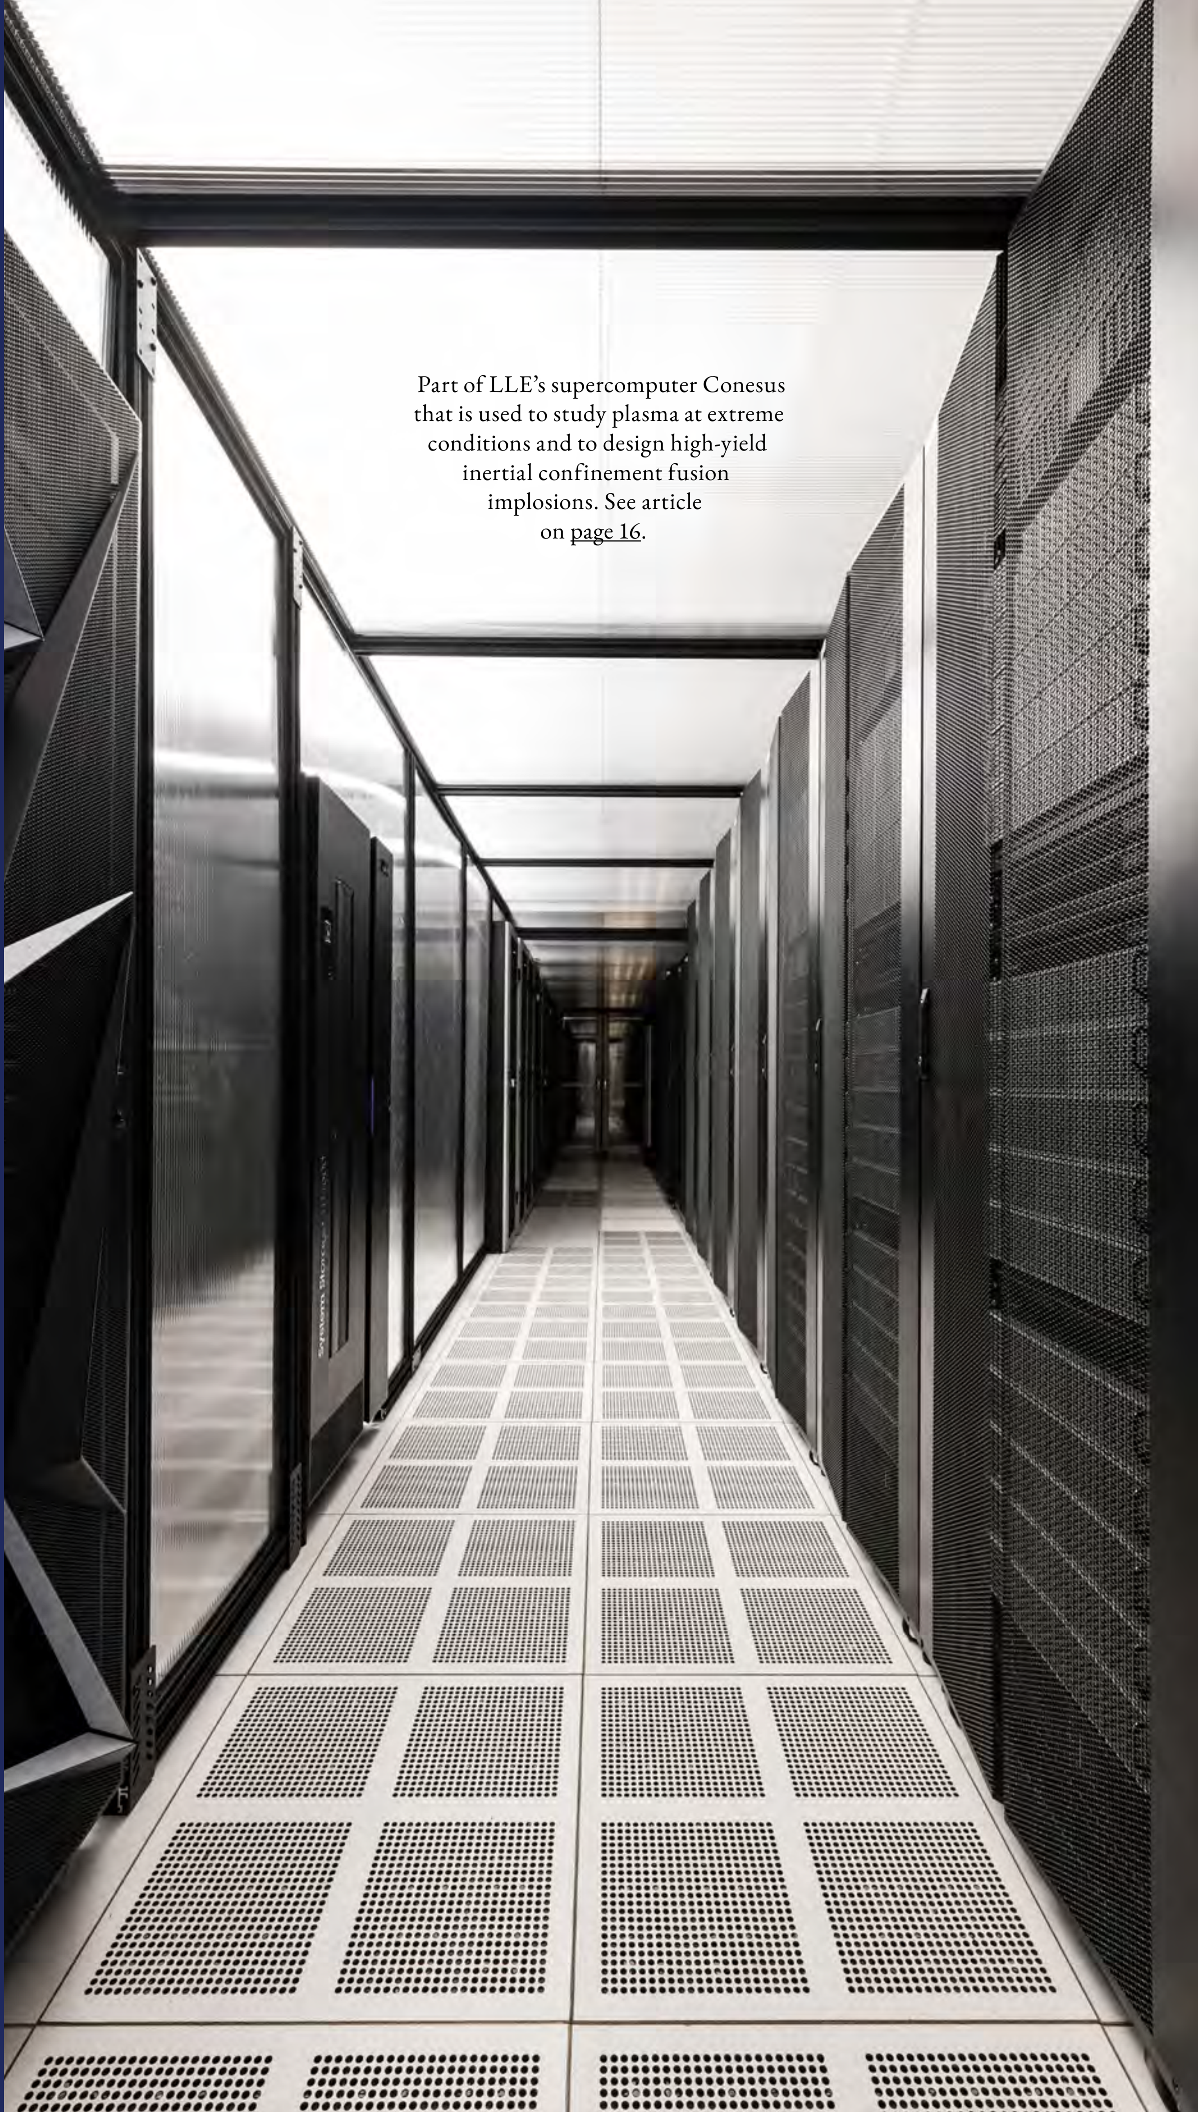

Part of LLE's supercomputer Conesus that is used to study plasma at extreme conditions and to design high-yield inertial confinement fusion implosions. See article on page 16.

LLE IN FOCUS

Focus on: ICF
Spring 2026 | Issue 9

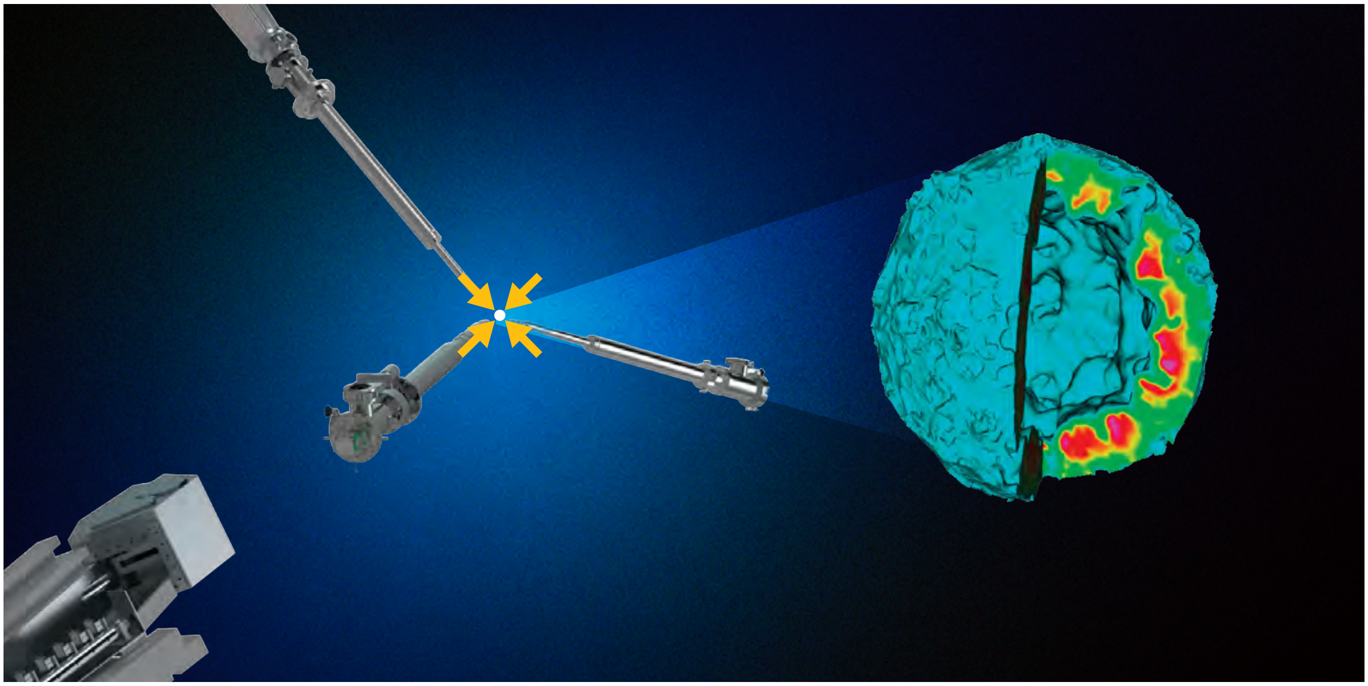

LLE develops and operates a wide range of diagnostics to study key issues of inertial confinement fusion implosions modeled using 3D hydrodynamic codes. See p. 8 for the complete article.

COLUMNS

2 From the Director
3 Highlights and Recognition
3 Publications
4 Inside the Laser Facility: Beam Smoothing for Picket-Pulse Fusion Experiments
7 US IFE Conference Focuses a National Effort

FEATURED ARTICLES

8 Fusion Science for High-Yield Applications
16 Simulation Impossible Inside the Computational Science Infrastructure Driving LLE's Research
19 The Opportunity of a Lifetime: Students Gain Rare Access to Some of the Country's Leading Science Facilities

About the cover

The cover shows the result of 3D simulation of an OMEGA cryogenic implosion at peak compression. Read more about the recent results achieved in ICF and the path to demonstrate high yields with laser direct drive on p. 8.

From the Director

DR. CHRISTOPHER DEENEY | DIRECTOR
LABORATORY FOR LASER ENERGETICS

Reading through this issue, one word kept coming back to me: scale. The scale of the science, the scale of the ambition of the questions being asked, and, not unrelated, the scale of what it takes to answer them.

The lead article on fusion science for high-yield applications puts LLE's direct-drive program in the context of a national effort that builds on ignition at the National Ignition Facility. It's a good moment to consider where direct drive and indirect drive fit and where the work is headed.

The piece on high-performance computing is, at its core, about what becomes possible when your simulation capability keeps pace with your experimental capability. The HPC Group has built an environment that lets us ask harder questions. The combination of advanced computing, exquisite experiments, and thorough diagnostics continues to advance the field.

And then there's the article about taking our undergraduates to visit national laboratories, which manages to be both a good career story and a quiet argument for why early exposure to science at scale matters. The students come back with wider eyes, and hopefully, bigger career goals.

A word about what's ahead: the next issue will be a special edition. We're marking 30 years of OMEGA, 40 years of chirped-pulse amplification, and 50 years of LLE. I find myself wanting to say something grand, but the truth is simpler: it was done by people. Hundreds of people across five decades who chose to come here and build something worth building. That's what we'll be celebrating.

I hope this issue gives you as much to think about as it gave me.

Christopher Deeney
Director, Laboratory for Laser Energetics

Highlights and Recognition

NIF & JLF User Groups Graduate Scholar Winners

Two LLE graduate students were recognized for their poster presentations at the 2026 National Ignition Facility and Jupiter Laser Facility Groups Meeting, where scientists from universities, national laboratories, and industry gathered to share share results from experiments conducted at both facilities. The meeting featured roughly 60 posters highlighting new advances in high-energy-density physics.

Ethan Smith and Hakhyeon Kim, shown above, both LLE Horton Fellows and graduate students in LLE's Experimental Group supervised by Rip Collins and Ryan Rygg, were recognized for their presentations. Smith received the Graduate Scholar Award for "Equation-of-State Measurements in Gigabar Plasmas," while Kim was named Graduate Scholar Runner-Up for, "A Free-Bound Emission Optimized Hard X-Ray Source."

Both awards highlight the strength of LLE's graduate training program and the contributions of its students to cutting-edge high-energy-density science.

LLE and Focused Energy Inc. Announce Research Collaboration

LLE and Focused Energy Inc. have established a $6.9 million partnership, the largest industrial-sponsored research agreement in LLE's history, to address key challenges in inertial fusion energy and accelerate progress toward practical fusion power. The collaboration reflects a growing model in which universities, national laboratories, and private companies work together to bridge fundamental science and commercial energy development.

LLE will apply its experimental and modeling capabilities to study laser–plasma instabilities that reduce fusion efficiency. The work builds on an existing US Department of Energy INFUSE collaboration is part of the Inertial Fusion Energy-Consortium on Laser–Plasma Interaction Research (IFE-COLoR) led by LLE.

Shown above, a conceptual rendering of Focused Energy's Fusion Pilot Plant, whose design will be informed by this work, illustrates how findings will translate laboratory results into real-world engineering conditions and advance fusion as a viable, scalable energy source.

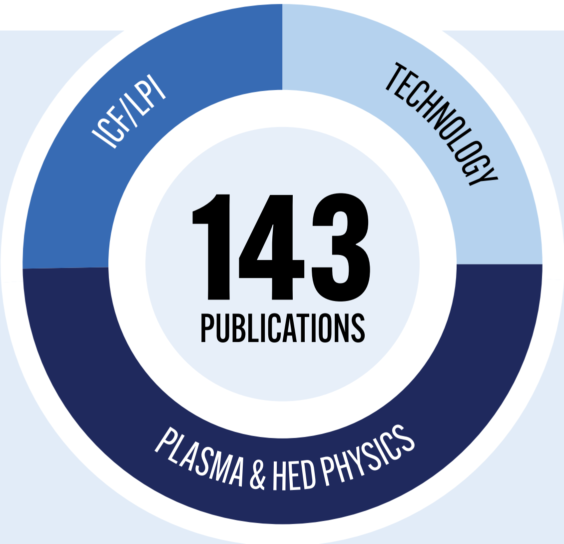

Publications

Research Highlights

LLE's science and engineering research is captured through peer-reviewed publications, which include LLE lead-authored and LLE coauthored papers. These publications reflect not only the laboratory's technical leadership but also its strong commitment to collaboration with external laboratories and academic partners. LLE averages more than 100 published articles annually across three broad areas: Technology, Plasma and High-Energy-Density (HED) Physics, and Inertial Confinement Fusion/Laser–Plasma Interactions (ICF/LPI). The graphic shows the distribution across topics for publications during the period January 1, 2025 through December 30, 2025.

Beam Smoothing for Picket-Pulse Fusion Experiments

How Spectral Dispersion and Timing Improve Irradiation Uniformity on OMEGA

The OMEGA Laser System is designed to deliver laser energy as uniformly as possible to a spherical target, using 60 beams arranged symmetrically around the target chamber. This uniformity depends not only on beam arrangement but also on the smoothness of each individual beam. The goal is to compress the spherical target evenly from all sides. If the laser light is uneven—too intense in some places and too weak in others—the target can distort as it implodes, reducing performance and making experiments harder to interpret. OMEGA uses three techniques that work together to smooth each laser beam:

- Distributed phase plates (DPPs), which split each beam into many small beamlets that overlap on the target
- Distributed polarization rotators (DPRs), which create overlapping speckle patterns with different polarizations that average together
- Smoothing by spectral dispersion (SSD), which rapidly moves the speckle pattern during the laser pulse

While all three smoothing techniques are important, SSD plays a unique role because it actively moves the laser speckle pattern during the pulse. Together, these methods ensure that the target receives a very smooth time-averaged intensity, even though the beam is not perfectly uniform at any instant. This smooth illumination is essential for symmetric inertial confinement fusion (ICF) implosions.

How SSD Works

SSD works by modulating the laser wavelength at high speed. Like light passing through a prism, each spectral component travels through the laser system at a slightly different angle. As the wavelength shifts back and forth, the beam traces a repeating pattern on the target known as a Lissajous pattern. At any one moment, the beam still has a speckle pattern, but because the pattern is moving rapidly, the target responds to a much smoother average intensity.

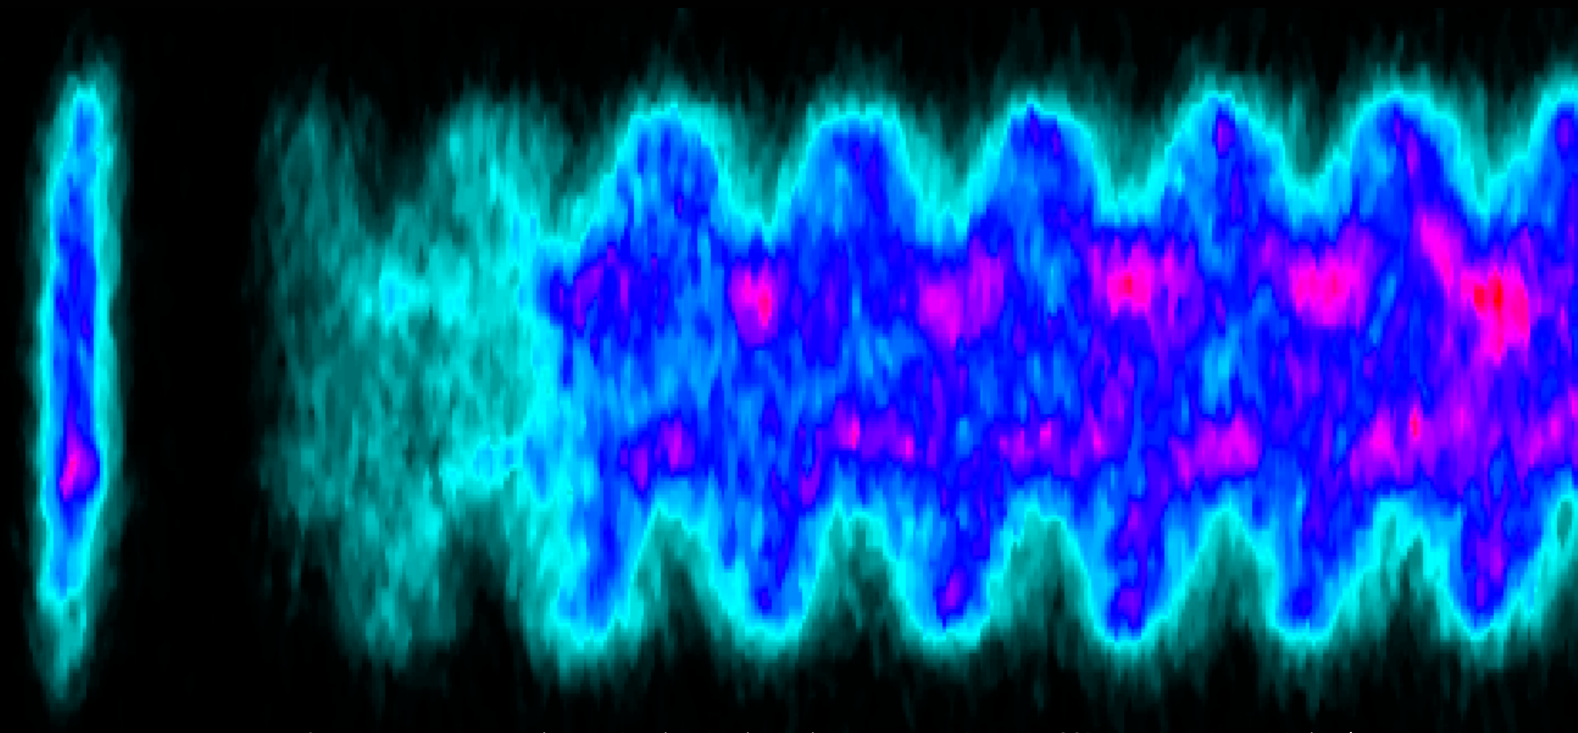

Streaked spectrometer image showing the ripple-like pattern produced by SSD wavelength modulation (wavelength versus time, with color indicating intensity). The timing of this modulation is adjusted so that the initial short "picket" pulse occurs at the zero crossing, where there is no wavelength shift, and the beam points in the correct direction.

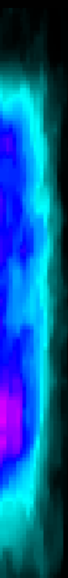

One way to understand SSD is to think about a long-exposure photograph. If the subject moves while the camera shutter is open, the camera does not record the instantaneous position of the subject; it records the average light over the entire exposure, resulting in a blurred image that represents the average motion over time. SSD on OMEGA works in a similar way. This motion-based smoothing helps ensure that the target is driven evenly during ICF implosions on OMEGA.

The Short-Pulse Challenge

Many ICF experiments use carefully shaped laser pulses that include a very short picket pulse at the beginning of the pulse. These pickets help set up the implosion before the main drive pulse arrives.

However, picket pulses can be extremely short, typically around 100 ps. In that short time, the SSD motion does not complete a full cycle. As a result, the average beam pointing during the picket can be slightly offset from the pointing during the main pulse. Even a small pointing offset during the picket can introduce large-scale asymmetries in the implosion, affecting performance and making results more difficult to interpret (Fig. 1).

SSD Sync: Optimizing SSD for Short Pulses

To solve this problem, engineers developed a technique called SSD Sync. SSD has been used on OMEGA for many years, but SSD Sync was developed to support modern pulse shapes that include very short picket pulses.

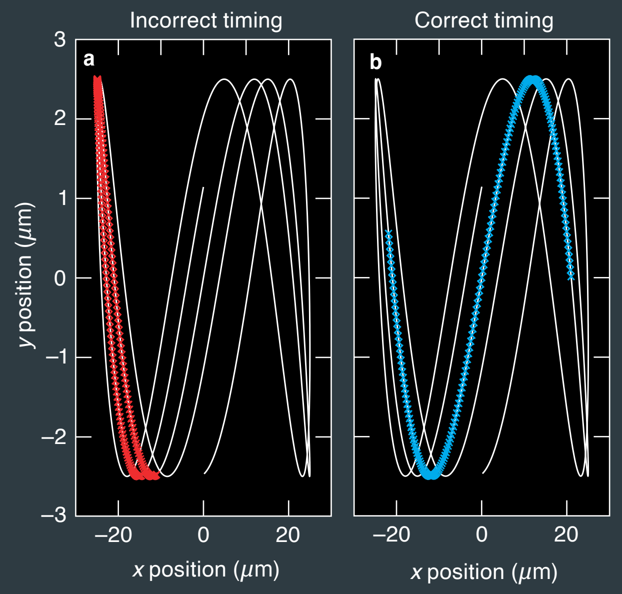

Figure 1. Effect of SSD Sync timing on beam pointing. The red and blue areas represent the beam pointing during the picket portion of the laser pulse. (a) Incorrect timing causes beam offset during the picket pulse. (b) Correct timing aligns the picket with the rf zero crossing, eliminating the offset and helping maintain symmetric implosions.

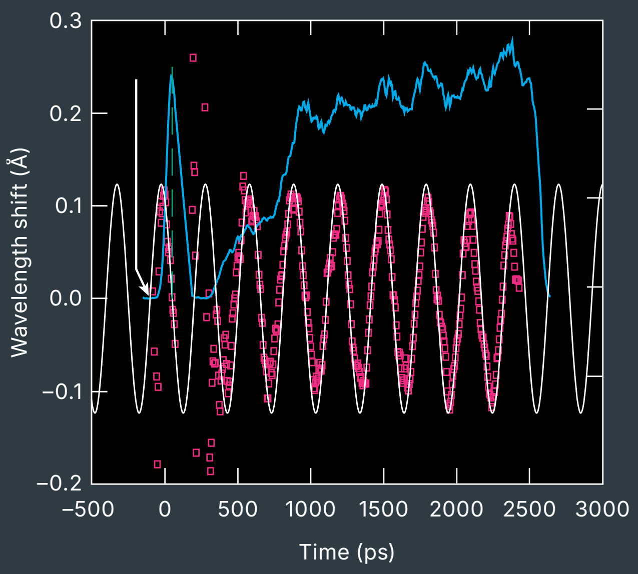

Figure 2. Streaked spectrometer measurement of SSD wavelength modulation. The red squares show the measured laser wavelength during the pulse, and the white curve is a sinusoidal fit that defines the radio-frequency (rf) phase. The blue curve is the UV pulse shape, including the short picket pulse at the beginning. Aligning the picket pulse with the zero crossing of the sinusoid minimizes beam pointing offset during the picket.

This system carefully times the picket pulse relative to the radio-frequency (rf) signal that controls the SSD wavelength modulation as measured with a streaked spectrometer (Fig. 2). There is a specific moment in the rf cycle, called the zero crossing, when there is no wavelength shift. With no wavelength shift, there is no angular shift in the beam, so it points in the same direction as the time-averaged beam.

Using a streaked spectrometer, engineers measure how the laser spectrum changes in time (Fig. 3) and adjust the phase of the rf signal so that the peak of the picket pulse occurs exactly at this zero-crossing point. This time-averaged beam pointing ensures that the beam pointing during the picket matches the pointing during the rest of the pulse [Fig. 1(b)].

Improving Consistency from Shot to Shot

SSD also introduces a small amount of unavoidable amplitude modulation as the beam moves through apertures in the laser system. By locking the rf phase, SSD Sync ensures that this modulation is the same from shot to shot. This consistency removes a variable from the experiment, making it easier for scientists to analyze implosion performance and compare results across experiments.

Why This Matters

The OMEGA laser was designed to deliver extremely uniform irradiation to fusion targets, and beam smoothing is critical to achieving this goal. Phase smoothing (DPPs), polarization smoothing (DPRs), and SSD work together to reduce laser nonuniformity and improve implosion symmetry.

SSD Sync is a refinement of this system that becomes especially important for modern pulse shapes that include very short picket pulses. By ensuring accurate beam pointing and consistent laser performance, SSD Sync helps improve experimental repeatability and the overall quality of the scientific data collected on OMEGA.

In ICF research, success depends on precision, symmetry, and repeatability. Technologies like SSD and SSD Sync operate behind the scenes, but they play a critical role in the success of ICF experiments on OMEGA.

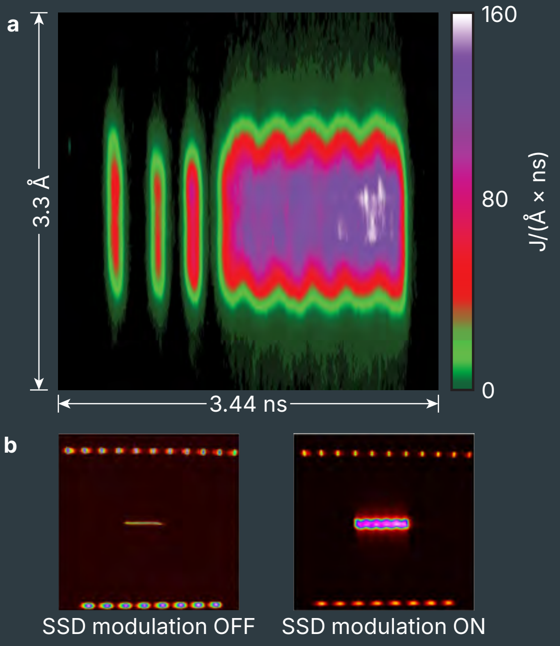

SSD modulation OFF SSD modulation ON

Figure 3. (a) Streaked spectrometer image showing the time-resolved laser spectrum. One axis represents wavelength, and the other represents time, while color indicates intensity. The oscillating pattern reveals the SSD wavelength modulation used to set the rf phase for SSD Sync. (b) With SSD off, the wavelength remains constant in time, producing a narrow, stationary trace. With SSD on, the wavelength oscillates in time, producing the ripple pattern.

Quarterly Shot Report

The Omega Laser Facility conducts experiments for research and development in support of the NNSA High-Energy-Density Program, National Laser Users' Facility, and Laboratory Basic Science, in addition to other research and educational efforts. During Q2 of FY25, the Omega Laser Facility conducted 224 target shots on OMEGA and 218 on OMEGA EP, totaling 442 target shots for 54 campaigns. OMEGA averaged 91% availability and 93% experimental effectiveness, while OMEGA EP averaged 96% availability and 94% experimental effectiveness.

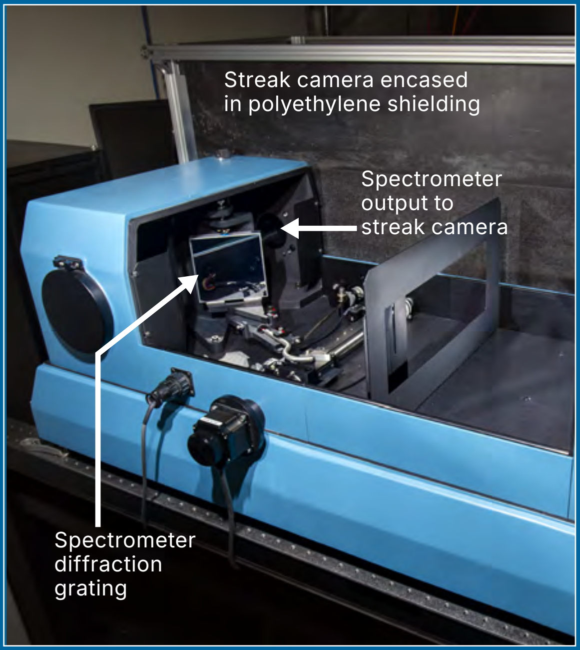

Streaked Spectrometer

To ensure that SSD Sync is working correctly, engineers need a way to observe how the laser's wavelength changes over time. This is done using a diagnostic called a streaked spectrometer, which combines two instruments into one.

A spectrometer spreads light into its component wavelengths (colors), mapping each to a different position along a line, much like a prism creating a rainbow. On its own, it provides a time-integrated measurement, showing which colors are present but not how they change during the pulse.

A streak camera, on the other hand, records how light intensity changes over time. It converts incoming light into electrons at a photocathode, then rapidly sweeps those electrons across a detector so that one dimension of the image represents time.

By combining these two instruments, the streaked spectrometer records both color and time simultaneously. The spectrometer first spreads the light into a line of wavelengths, and that line is then imaged onto the streak camera. As the signal is swept across a 2D detector, one axis represents wavelength, and the other represents time.

Each vertical slice of the image shows the instantaneous spectrum of the laser at a given moment, while each horizontal slice shows how the intensity of a particular wavelength changes over time. This allows engineers to observe the sinusoidal wavelength modulation produced by SSD and to align the timing of the picket pulse with the rf zero crossing.

In practice, the streaked spectrometer produces images like the one shown in Fig. 3, where the oscillating pattern reveals how the laser wavelength shifts during the pulse. By analyzing this pattern, the SSD phase can be adjusted to ensure proper synchronization and maintain accurate beam pointing during short picket pulses.

US IFE Conference Focuses a National Effort

More than 225 members of the US inertial fusion energy (IFE) community gathered in the Washington, DC area on March 22–27 for the second annual US IFE Conference, formerly IFE-STAR. Organized through LLE and held at the Hilton DC/Rockville hotel, the week combined a rigorous technical program with a focus on building the cross-institutional relationships the field will need as it scales.

The conference opened with US Department of Energy (DOE) leaders Jean Paul Allain, inaugural Director of the Office of Fusion, and Kramer Akli, Fusion Energy Sciences Program Manager, who outlined current US fusion strategy and mapped the evolving IFE ecosystem. Nine national laboratories contributed talks across a program spanning target physics, theory and simulation, AI and machine learning, driver development, systems engineering, and fusion materials—a breadth that reflects the field's shift toward the engineering demands of a working power plant. Universities presented topics ranging from final optics design to open-source chamber modeling to the societal dimensions of fusion deployment. A single session brought all three DOE-funded research hubs—STARFIRE, IFE-COLoR, and RISE—together for their halfway-mark reports, offering a clear picture of how the national IFE program fits together.

Private industry also presented technical sessions, with major IFE companies sharing progress on pilot-plant design, high-repetition-rate target systems, and fuel cycle development. International contributions addressed industrialization pathways and large-scale experimental infrastructure.

> ***The US IFE conference brings the entire technical community into the same room, because that cross-exposure is exactly what the field needs right now.***
>
> ***—Dustin Froula***

"Most conferences focus on one area: target physics, drivers, or materials," said Dustin Froula, principal investigator for the IFE-COLoR hub at LLE and chair of the conference organizing committee. "This one brings the entire community together. The science is no longer the only question, and the path to a power plant is starting to come into focus."

Student and early-career researchers were central to the week. A graduate student poster slam featured participants delivering two-slide presentations in three minutes before continuing their conversations at the poster sessions that followed. Across these sessions, more than 100 posters were presented, drawing contributors from across the US and international partners, covering the full spectrum of IFE research from target fabrication and plasma physics to chamber modeling and workforce development. Eleven students, from the high school to graduate level, were recognized for outstanding research contributions.

The next US IFE Conference will return to Washington, DC, in 2027.

U.S.IFE

Fusion Science for High-Yield Applications

Current Status of Inertial Confinement Fusion Research

December 5, 2022 marked a major milestone for humankind. On this day, a team led by Lawrence Livermore National Laboratory (LLNL) used the world's most powerful laser at the National Ignition Facility (NIF) to compress a mixture of deuterium and tritium (DT) to conditions in which the energy produced by fusion reactions exceeded the laser energy used in the experiment. In other words, ignition was achieved.

Fusion reactions themselves are not exotic: they power the Sun and occur during star formation throughout the universe, naturally taking place on vast scales in systems with enormous energies. On Earth, significant fusion energy release has also been observed, most notably in thermonuclear weapons. However, even these are far too energetic to be controlled in a laboratory setting.

A distinctive feature of the LLNL experiments is their extremely small scale. The team used just 0.22 mg of DT fuel contained in a target roughly the size of the head of a matchstick. With a laser input of 2.05 MJ, approximately the amount of energy required to bring two gallons of water to a boil, the experiment produced 3.1 MJ of fusion energy, satisfying the National Academy of Sciences's definition of ignition. The fuel was compressed using the inertial confinement fusion (ICF) approach, in which lasers drive the implosion of the target. You can read more about this groundbreaking experiment in the article, "Lawson Criterion for Ignition Exceeded in an Inertial Fusion Experiment," published in *Physical Review Letters* [1].

ICF shares many similarities with a conventional combustion engine. To achieve ignition in an ICF experiment, one must first compress the DT fuel—analogous to the compression of gasoline vapor by a piston in an engine. In ICF, compression is achieved either through direct laser irradiation of the target (direct drive) or by x rays (indirect drive) generated by converting laser light inside an enclosure made of metal with high x-ray production efficiency (e.g., gold). Such laser or x-ray irradiation ablates the target material, generating a rocket-like effect that drives the inward compression of the DT fuel. The fundamentals of ICF are outlined in detail in *LLE In Focus*, Issue 2.

After the fuel is compressed, it is ignited by a spark. In a combustion engine, this spark is produced by a spark plug, but in ICF, the spark is formed in a central region of high-temperature (~20 million degrees) DT plasma known as the "hot spot." These conditions are created as the surrounding shell of higher-density fuel compresses the lower-density region in the target center. For ignition to occur, a key condition known as the Lawson criterion must be satisfied—essentially, the heating rate of the plasma in the central region must exceed its cooling rate.

Reaching ignition on the NIF is a remarkable achievement because igniting the small amounts of DT fuel in a laboratory setting is extremely challenging. The success of the NIF campaign was built on more than 70 years of research conducted at numerous facilities worldwide. It also required the ingenuity and expertise of scientists and engineers at the scale of a national laboratory, working with the world's largest laser for over a decade to reach this milestone.

Why is achieving ignition in a laboratory so difficult? One reason is the extremely small scale of the igniting plasma. Although the NIF facility spans the area of three football fields, the burning plasma generated in an experiment measures only about 80 μm in diameter—roughly the width of a strand of human hair. What limits the plasma size? The short answer is the available energy.

Taking a deeper look, we can invoke the ignition criterion, which in its simplest form can be expressed in terms of the hot-spot energy and size. This is shown in Fig. 1, which

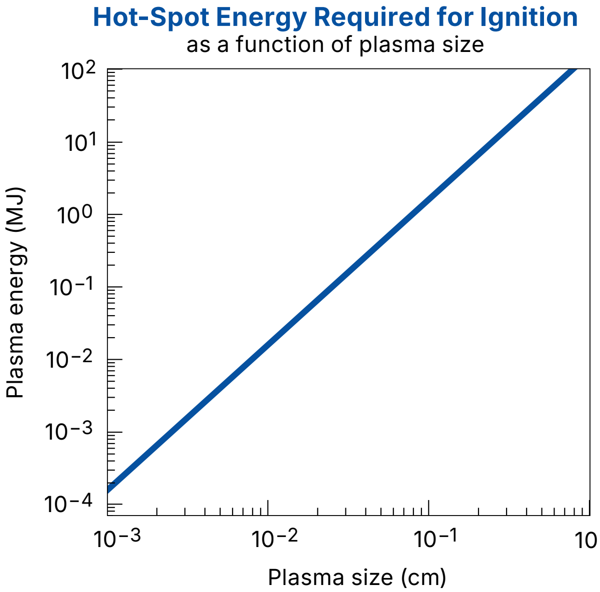

Figure 1. Hot-spot energy required for DT ignition as a function of plasma size.

suggests that for a hot spot with an energy of 2 MJ—the same as the laser energy used on the NIF—the size could be as large as 0.1 cm. Why, then, is the plasma size so much smaller in the NIF experiment?

To answer this question, it is important to understand that one of the main limitations of ICF is its low efficiency in coupling laser energy to a hot-spot plasma. Several factors contribute to this. As an example, let us consider the indirect-drive approach, where laser energy is first converted into x rays inside a gold can (hohlraum). This conversion is only about 20% efficient. The x rays then heat the target, driving mass ablation from its surface. The ablated material generates a force, analogous to rocket thrust, which is also only ~10% efficient. So, the fuel gets only a small percent of incident laser energy. Finally, at peak compression, roughly 30% of the energy coupled to the fuel reaches the hot spot. Despite the world's largest laser delivering 2 MJ of energy, the hot spot gets only ~7 kJ. Using our graph in Fig. 1, therefore, this energy will ignite a plasma only 80 μm in size.

The other challenge of ICF is that the target needs to be accelerated inward to an extreme velocity. Indeed, containing 7 kJ of energy inside a ball 80 μm in diameter is equivalent to a plasma pressure of 300 Gbar or 300 billion atmospheres, which is very similar to the conditions inside the Sun—and not surprising since fusion is the very process that powers the Sun! Squeezing energy into a tiny volume requires accelerating the targets to extremely high velocities, above 300 km/s, or nearly 1 million miles per hour. This velocity is over 30 times faster than the escape velocity of a rocket traveling to outer space.

The time scale in ICF is equally extreme. The implosion time is set by the ablation pressure generated in the ICF "rocket." The resulting ablation pressure acts on the target surface, providing the accelerating force. If the pressure is high, the target gains its energy by this force acting over a shorter distance; if it is lower, the target must be accelerated over a much longer distance.

In a typical indirect-drive implosion, the ablation pressure is of the order of 150 Mbar (150 million atmospheres, or 100 times more than what is required to form diamonds inside the Earth). This pressure accelerates the target to ~300 km/s over a distance of ~1 mm. Therefore, the corresponding implosion time—and thus the laser pulse duration—is only ~5 ns. This is an extraordinarily short time: in 5 ns, light barely travels your own height.

One of the greatest challenges the LLNL team needed to overcome in demonstrating ignition on the NIF was maintaining the target's spherical symmetry throughout the compression. Since the initial target diameter is ~1 mm, and the size of burning plasma 80 μm in diameter, the fuel must converge by a factor of 25 during an implosion. At such extreme convergence—in which the volume changes by a factor of 25^3, or 15,625—target imperfections, drive asymmetries, and the effects of target mounts are amplified by hydrodynamic instabilities such as the Rayleigh–Taylor instability (RTI). The instability growth can disrupt the shell, preventing it from reaching ignition conditions.

The Role of Direct Drive in the Quest for High Yield

Following the achievement of ignition, the key challenge becomes reaching the ultimate goal of ICF: producing high yields of the order of several hundreds of MJ. In current ignition experiments on the NIF, the relatively small hot-spot energy (~7 kJ) constrains the hot-spot diameter to about 80 μm. This results in extremely high energy densities, or pressures, within the hot spot. Such high pressures (exceeding 300 billion atmospheres) can be attained only in designs with shell velocities above 350 km/s. NIF ignition experiments reached implosion velocities of ≃390 km/s to compensate for performance losses due to shell asymmetries.

High shell velocities have two key implications for the design. First, they limit the amount of DT fuel that can be carried in the shell since heavier shells are harder to accelerate; the reduced fuel mass, in turn, constrains the achievable fusion yield. Second, high velocities compromise shell stability: a lower shell mass leads to a thinner shell that is more vulnerable to breakup from instability growth.

To address the shell stability problem, the LLNL team adopted an approach that increases shell thickness and mitigates instability growth by adding heat and raising the shell temperature. This was achieved by strengthening the first shock launched at the beginning of the implosion. In indirect-drive implosions with carbon ablators, a sufficiently strong first shock is also required to melt the carbon, thereby eliminating its granular microstructure, which would otherwise provide additional seeds for RTI growth.

The improved stability obtained through enhanced fuel heating, however, comes at the cost of reduced fuel compressibility. Higher fuel temperatures make it more difficult to compress the fuel to high densities. As a result, lower fuel densities in current ignition experiments lead to faster disassembly of the fuel after ignition conditions are reached, which result in shorter confinement times. This, in turn, reduces the burn temperature and decreases the fraction of the fuel that participates in fusion reactions.

Taken together, these limitations of current ignition experiments on the NIF point to an obvious conclusion: achieving robust ignition and higher fusion yields require coupling more energy to the fuel.

The advantages of increased fuel energy in ICF designs can be summarized as follows:

- Higher fuel energy leads to a larger "spark" region (hot spot), thereby lowering the required hot-spot pressure
- Reduced hot-spot pressure relaxes the need for extremely high shell velocities
- Reducing the velocity requirement allows for thicker shells with greater fuel mass to be used in the target design

- Thicker shells are less susceptible to the growth of nonuniformities during implosion, reducing the need for shell preheating by the initial shock
- Lower fuel temperature enhances compressibility, resulting in a higher fuel burn fraction

In short, increasing the fuel energy provides a direct pathway to robust higher yields by enabling greater fuel mass in the target and increasing the fraction of fuel that participates in fusion reactions.

To increase fuel energy, the LLNL team is currently pursuing an “extended yield capability” project, which aims to increase the laser energy from the current level of 2.2 MJ to 2.6 MJ for achieving yields of tens of megajoules. Looking further ahead, plans for a next-generation high-energy-density facility envision increasing the laser energy to as much as 10 MJ to reach the ultimate goal of hundreds of megajoules.

The direct-drive approach provides a highly attractive pathway for substantially increasing the energy delivered to the fuel. Its principal advantage is the elimination of the energetically costly conversion of laser light into x rays. As discussed earlier, this conversion consumes a lot of the incident laser energy, leaving only 20% of the laser energy to be coupled to the target. Therefore, by directly illuminating the target with laser beams, more than 90% of the laser energy can instead be coupled to the target.

OMEGA Turns 30

Sponsored by the National Nuclear Security Administration, the OMEGA Laser System has been at the heart of direct-drive inertial confinement fusion and high-energy-density physics since 1995.

Over three decades, OMEGA has supported thousands of experimental campaigns, trained hundreds of students and scientists, and generated insights that have shaped the national Inertial Confinement Fusion program.

As in the indirect-drive scheme, the absorbed laser energy in direct drive ablates the target material, generating a rocket-like effect that accelerates the target inward. Although the rocket efficiency of the laser drive is lower than that of x rays—since x rays penetrate deeper into the target—direct drive can, in principle, accelerate the shell to energies up to four times higher than those achieved with indirect drive. In a direct-drive target implosion on OMGEA driven by a 30-kJ laser pulse, approximately 1 kJ of energy is coupled to the hot spot..

Given this substantially higher coupling efficiency, a natural question arises: why has ignition not yet been demonstrated with the direct-drive approach on the NIF? The short answer is that the NIF laser is not optimized for direct-drive implosions. Experience from direct-drive experiments on OMEGA and the NIF has shown over the last decade that, to fully realize the advantages of this approach, the laser must be broadband—in other words, the laser must emit light at a wide range of wavelengths. In addition, advanced beam-smoothing technologies must be incorporated into the laser system.

The broadband requirement comes from the need to mitigate deleterious laser–plasma interaction effects. As the laser propagates through the low-density blowoff plasma produced by target ablation, it excites plasma waves—similar to the way a boat generates a wake as it moves through water. These plasma waves can then interact with laser beams arriving from different directions, scattering their energy away from the target. This phenomena is called cross-beam energy transfer, or CBET. In addition, they can accelerate plasma electrons, analogous to surfers gaining speed by riding ocean waves. These energetic electrons can travel long distances and deposit their energy in the cold fuel, increasing its temperature and thereby reducing

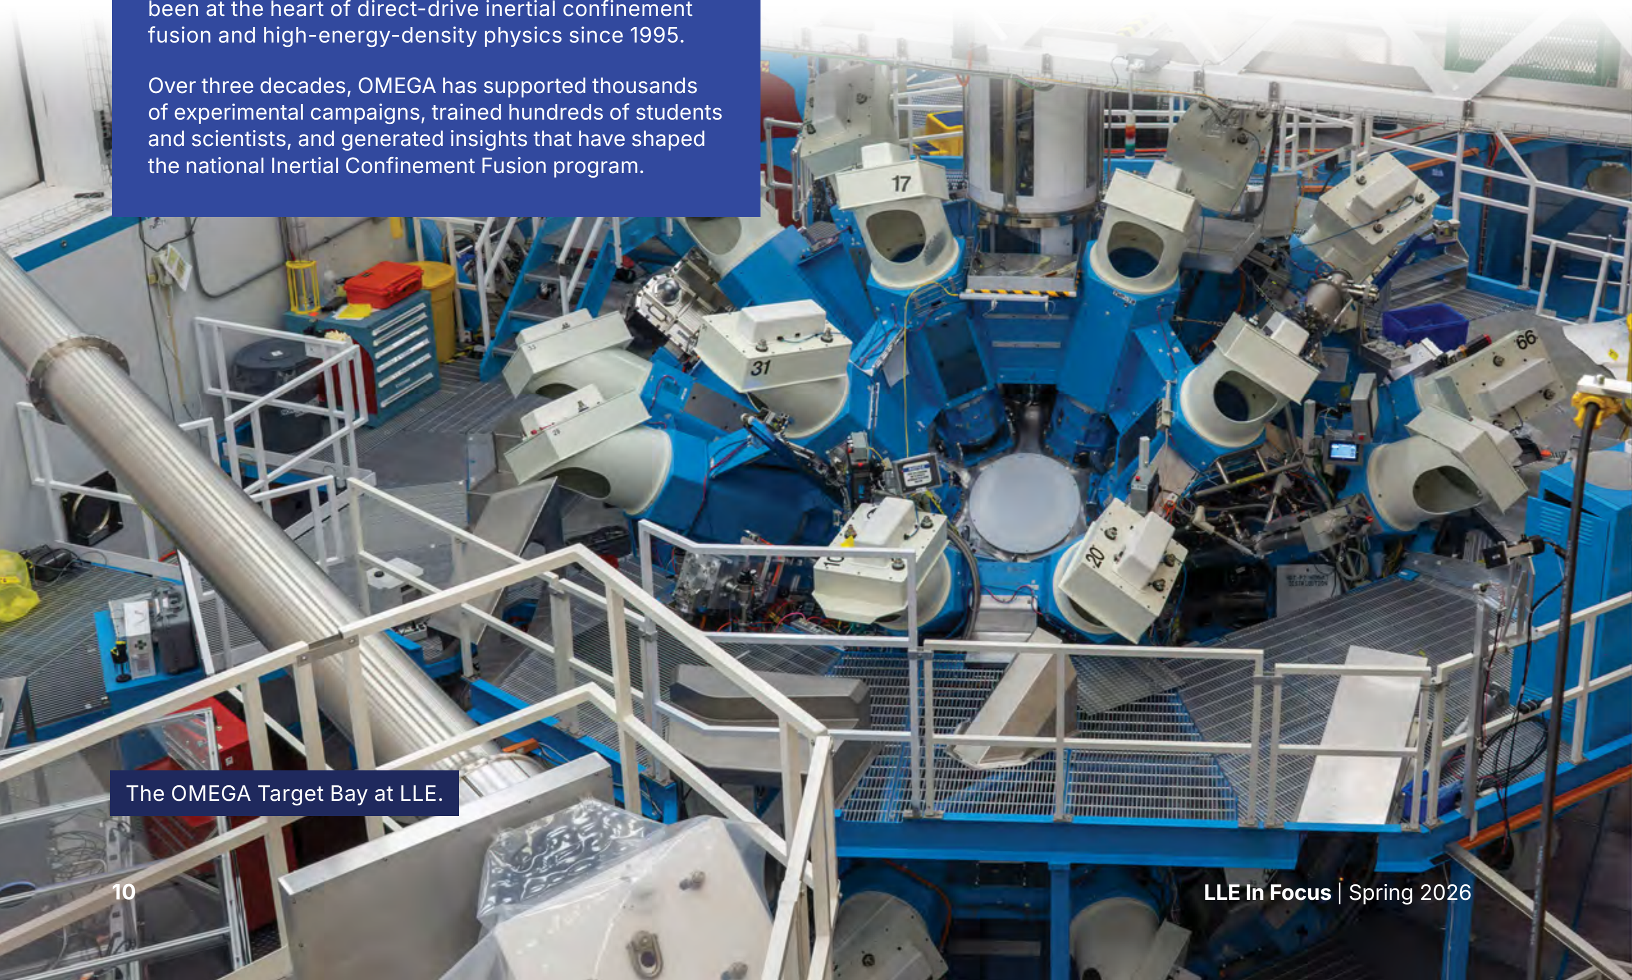

The OMEGA Target Bay at LLE.

its compressibility, as discussed earlier. More information on CBET can be found in *LLE In Focus*, Issue 3.

The excitation of plasma waves is a resonant process. If the laser operates at a single wavelength, plasma waves are driven resonantly at specific frequencies and locations within the plasma. In contrast, a broadband laser distributes its energy over multiple wavelengths, so each component drives waves at different frequencies and locations. This reduces the coherence and amplitude of the plasma waves, leading to significantly less laser scattering and weaker electron acceleration.

Beam smoothing is also critical for the direct-drive approach since the laser interacts with the target directly. Laser beams in ICF experiments consist of narrow, high-intensity speckles. This speckle pattern gets imprinted onto the target surface at the beginning of the target drive and is subsequently amplified by hydrodynamic instabilities (such as RTI) as the shell accelerates.

To reduce this imprinting, advanced beam-smoothing techniques are employed in laser systems. On the OMEGA laser, for example, smoothing by spectral dispersion (SSD) effectively moves the speckles on a time scale much shorter than the imprinting time, thereby reducing the initial nonuniformity. While a reduced version of SSD is also implemented on the NIF, it provides only a fraction of the smoothing required for laser direct-drive implosions.

Fusion Research at LLE in Support of the High-Yield Mission

Current research at LLE focuses on understanding the key physics of direct-drive implosions and defining the requirements for future laser facilities that can fully leverage the advantages of direct drive to achieve high yields, supporting both national security objectives and inertial fusion energy applications. In the next section, we highlight the key elements of fusion research by grouping them according to three main stages of an ICF implosion.

ICF Implosion Stages

A typical ICF implosion (both direct and indirect drive) proceeds through three main stages: (1) early time (shock propagation), (2) acceleration, and (3) deceleration and neutron production. To meet the requirement for a high-yield implosion, key implosion elements must be understood:

- Fuel heating
- Imprint
- Shell nonuniformity growth during acceleration and deceleration
- Laser coupling during acceleration
- Generation of and fuel heating by energetic electrons
- Fuel asymmetry and nonradial flow at peak compression
- Fuel areal density and temperature at peak compression

Most of these critical elements are common to both direct and indirect drive. In the discussion that follows, we illustrate each implosion stage, highlighting the diagnostics used to measure or infer these key quantities.

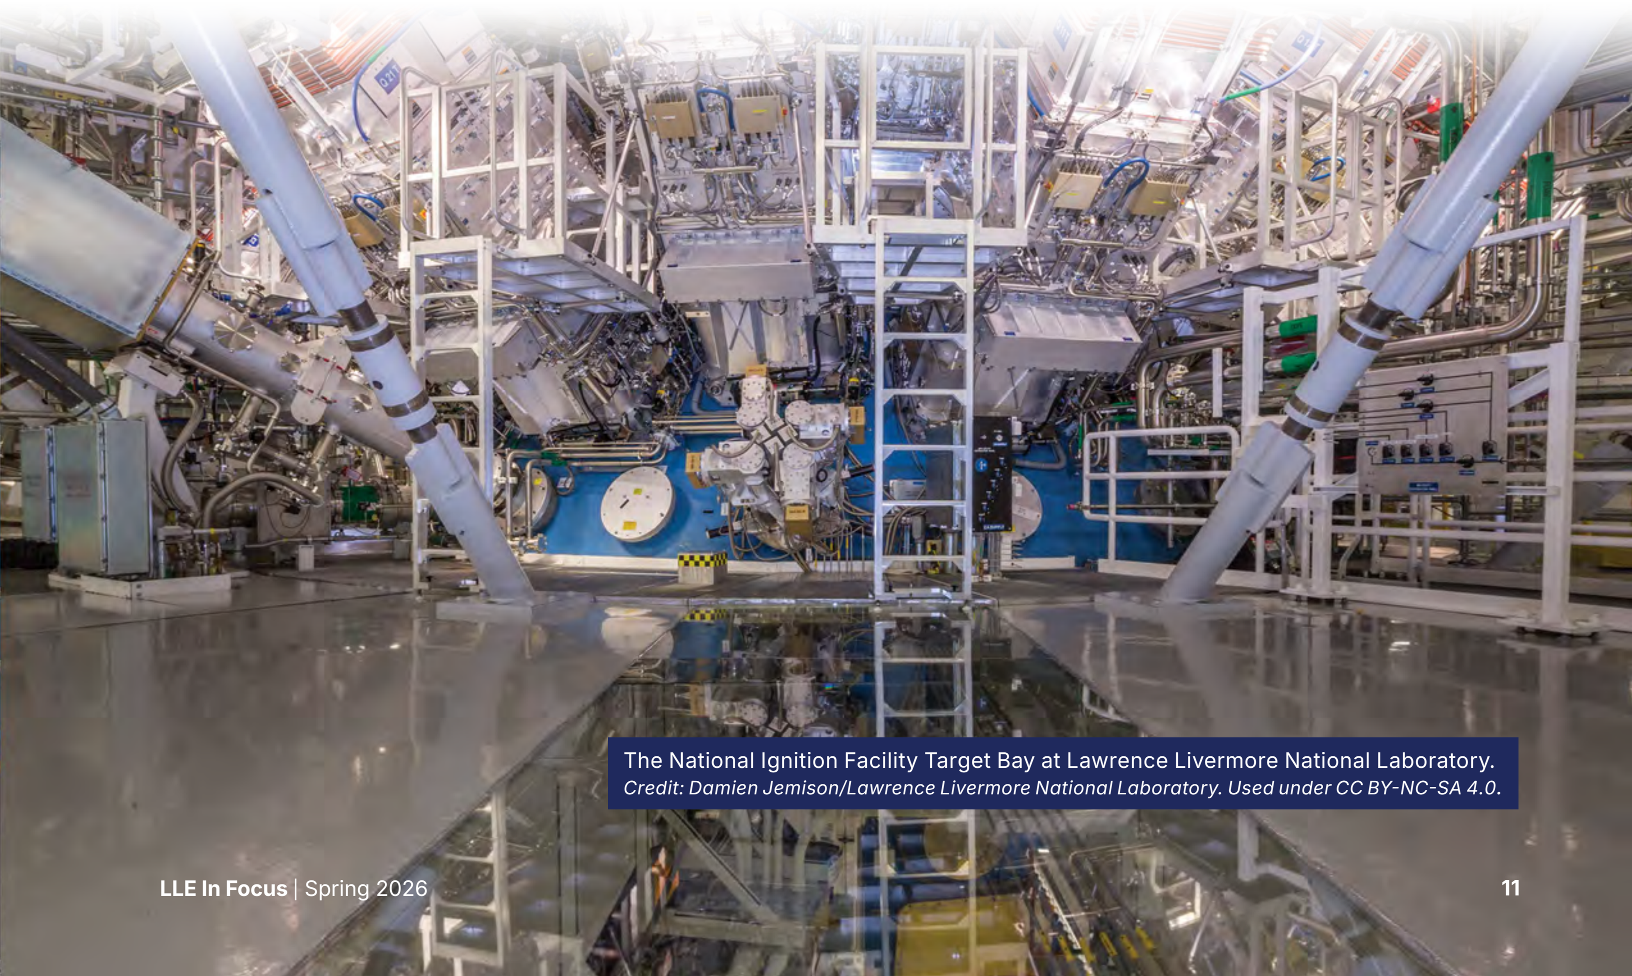

The National Ignition Facility Target Bay at Lawrence Livermore National Laboratory.
Credit: Damien Jemison/Lawrence Livermore National Laboratory. Used under CC BY-NC-SA 4.0.

STAGE 1
SHOCK PROPAGATION

At the start of an ICF implosion, a shock wave is driven into the target shell in both direct- and indirect-drive schemes. The key physics at this stage includes shock heating, which determines the fuel compressibility, and shock nonuniformities, which seed hydrodynamic instabilities.

Direct Drive

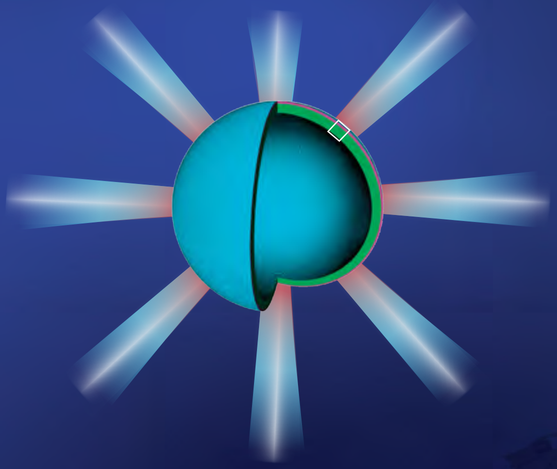

Indirect Drive

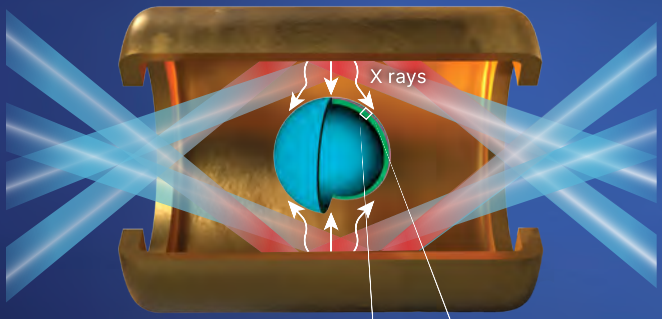

Diagnostics

LLE operates a wide range of specialized diagnostic instruments, each designed to measure specific aspects of high-energy-density physics experiments. Together, these tools provide complementary data that allow researchers to observe, analyze, and refine experimental performance with precision.

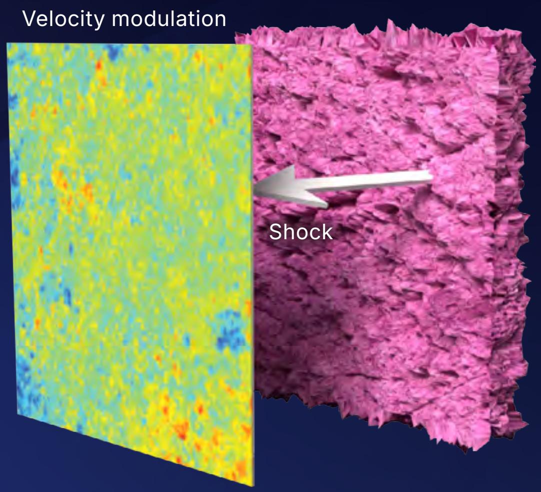

Shock Modulation Measurements

Hydrodynamic instabilities are seeded by shock modulations. In direct drive, the dominant source is laser imprint, while target imperfections—such as shell defects, ice roughness, and mounting features (fill tube, stalk, tents)—also contribute to nonuniformities in both direct and indirect drive. Shock modulations are measured using the 2D velocity interferometry system for any reflector (VISAR) or OMEGA high-resolution velocimeter (OHRV) diagnostics.

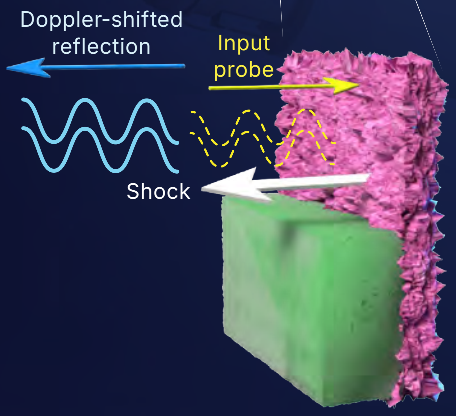

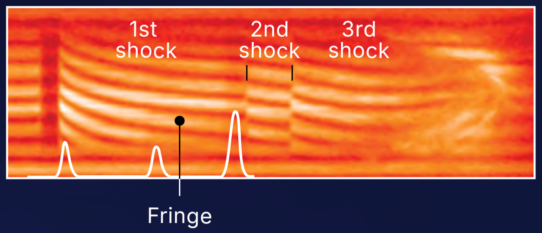

Shock Velocity Measurement

Accurately measuring the shock strength is key for setting the fuel conditions required for the high-yield designs. The shock strength is inferred by measuring shock velocity using the VISAR diagnostic. VISAR detects the Doppler shift of an optical probe beam that reflects off the shock front. The interference between the probe and reflected beam is imaged as fringes whose vertical position is proportional to the shock velocity, so changes in fringe position are proportional to the change in velocity.

STAGE 2
ACCELERATION

Shortly after shock breakout from the shell, both direct- and indirect-drive targets begin to accelerate inward. The key physics during the acceleration phase include laser coupling, the shell acceleration rate (which determines the ablation pressure), amplification of shell nonuniformities via the Rayleigh–Taylor instability, and fuel preheat caused by energetic electrons.

Direct Drive

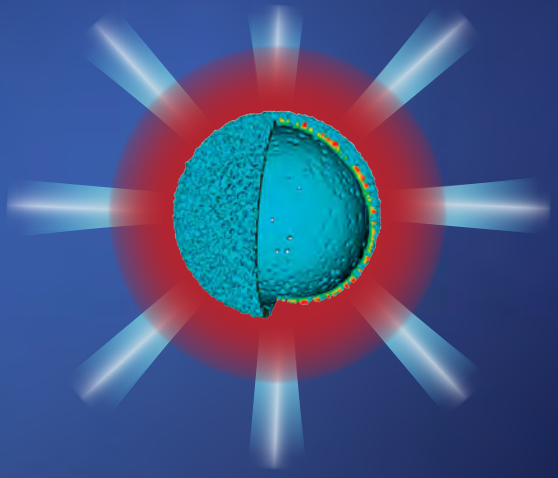

Indirect Drive

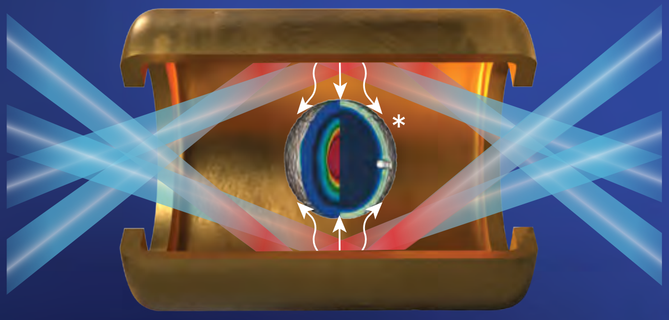

*Capsule image reprinted from D. S. Clark *et al.*, Phys. Plasmas **26**, 050601 (2019). with the permission of AIP Publishing.

Diagnostics

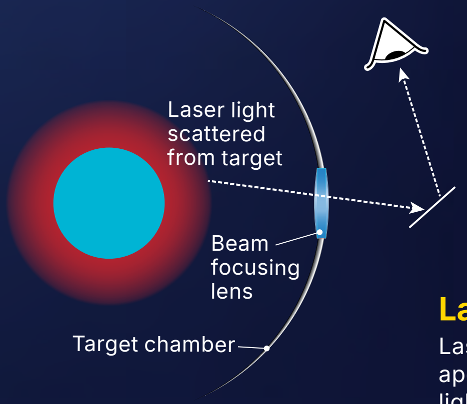

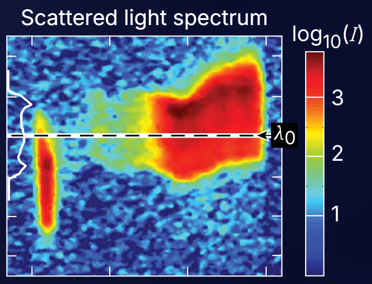

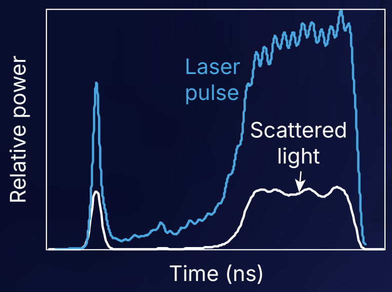

Laser Coupling Measurement

Laser coupling is inferred by measuring the scattered light by two full-aperture backscattering stations (FABS). Time-resolved scattered-light spectroscopy and time-integrated calorimetry are used to infer the absorption of light by the target.

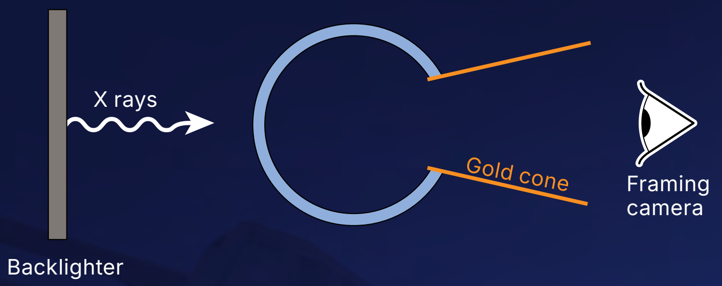

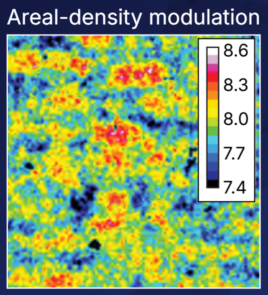

Shell Nonuniformity Growth Measurement

Nonuniformity growth can be directly measured using through-foil radiography. The shell for this measurement has an opening into which a gold cone shield is inserted. An x-ray backlighter is irradiated with a subset of the OMEGA beams, making it possible to diagnose the growth of areal-density perturbations using an x-ray framing camera.

STAGE 3

DECELERATION & NEUTRON PRODUCTION

As the shell continues to converge, the central fuel gas is progressively compressed. At a certain point in the implosion, the central pressure becomes sufficiently high to oppose the inward motion of the shell, initiating its deceleration.

With continued convergence, the central pressure, density, and temperature increase, strengthening the decelerating force. Eventually, the shell comes to a halt and subsequently begins to expand outward. The peak hot-spot temperature, reached near maximum convergence, coincides with the time of the peak fusion reaction rate.

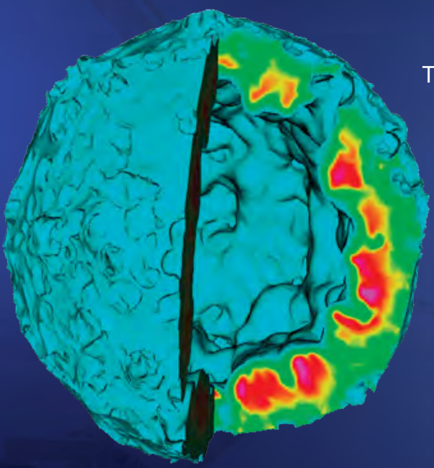

The key physics issues during this stage are common to both direct- and indirect-drive approaches and include the fuel areal density (the product of shell density and its size), as well as fuel symmetry, hot-spot size, and the extent of ablator material that becomes mixed into the hot spot.

Diagnostics

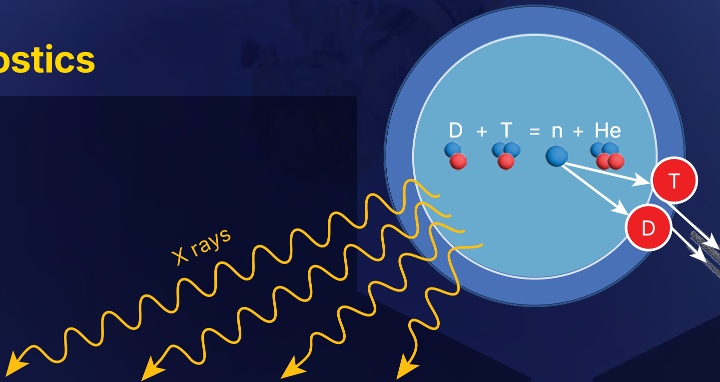

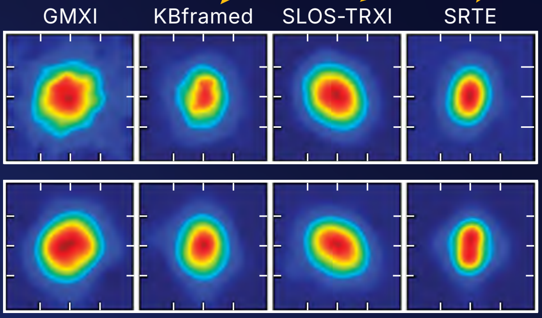

3D reconstruction using AI tools

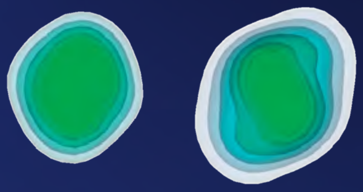

X-Ray Diagnostics

Shell and hot-spot asymmetries limit the maximum temperature and density in fuel. The hot-spot shape is inferred on OMEGA with four x-ray diagnostics positioned around the target chamber. The single-line-of-sight time-resolved x-ray imager (SLOS-TRXI), Kirkpatrick–Baez x-ray microscope (KBframed), gated monochromatic x-ray imager (GMXI), and spatially resolved electron temperature (SR-Te) imager provide the data for physics-informed machine-learning tools to reconstruct the 3D shape of the hot spot in an ICF implosion.

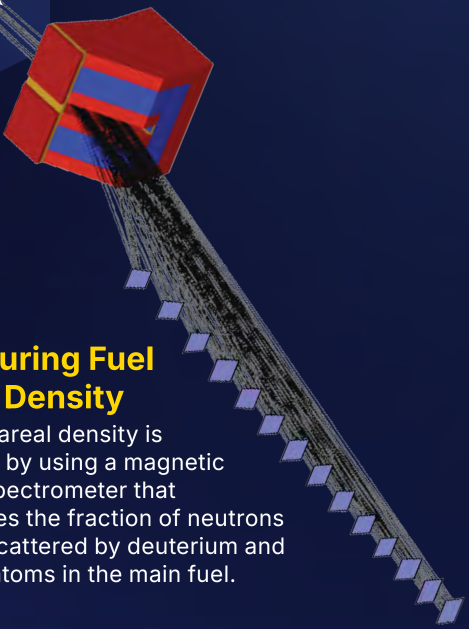

Measuring Fuel Areal Density

DT fuel areal density is inferred by using a magnetic recoil spectrometer that measures the fraction of neutrons down-scattered by deuterium and tritium atoms in the main fuel.

Using OMEGA and NIF Implosions to Validate Modeling of High-Yield Designs

The detailed measurements described on pp. 12–14 help designers improve and validate the modeling of key physics phenomena of ICF implosions. For example, in comparing the OMEGA high-resolution velocimeter (OHRV) data of shock nonuniformity with the model predictions (see p. 12), enhanced nonuniformity was observed in the data. This triggered an examination of early-time interactions of the laser light with the target material. In particular, the light penetrates and deposits its energy deep into the target until the atoms of the target materials are ionized, creating an electron barrier that reflects the light from the target's outer edge. While light travels inside the shell, it damages the shell material, introducing additional nonuniformity seeding that is observed in OHRV data. To include this effect in calculations, a model of early plasma formation has been developed and implemented in hydrodynamic codes.

The other example of how measurements feed into modeling improvements is measuring laser light absorption during shell acceleration using full-aperture backscatter (FABS) stations (see p. 13). FABS measures the color (frequency) of the laser light scattered from the target. When laser light, as with any other wave, reflects off a moving surface, its wavelength shifts due to the Doppler effect. This is similar to hearing different pitches from the sound of a race car that is moving either toward or away from you. Laser wavelength also becomes slightly longer when it reflects off an imploding target surface—that is, moving away from the source. By measuring the wavelength of the reflected laser light, a smaller-than-expected shift was observed from the imploding target. The interpretation of this deficiency in Doppler shift is that the laser light does not penetrate deeply enough and reflects earlier, prior to reaching the target surface. The light reflects from plasma waves that were launched by the laser itself due to CBET, as discussed on p. 10. This effect is responsible for 30 to 40% of light reflection from the target and a significant reduction in rocket effect and ablation pressure.

This important understanding of laser coupling limitations in direct-drive implosions led to the idea of using broadband lasers, and experiments are currently underway to understand how this technology mitigates coupling losses due to CBET. These experiments are performed using the FLUX system (Fourth-generation Laser for Ultra-broadband eXperiments). Demonstrating CBET mitigation will imply that nearly 100% of the laser energy can be coupled to the target, raising the ablation pressure to several hundred million atmospheres and making broadband technology the leading candidate for a next-generation high-energy-density implosion facility, OMEGA Next. Although much smaller than the NIF, the facility's main goal will be to demonstrate that all direct-drive deficiencies due to laser–plasma interactions can be eliminated via a spherical implosion, and that laser technology is ready for an at-scale implosion facility to ultimately deliver high fusion yields (>100 MJ) for NNSA program needs and future energy applications.

Corresponding authors:
V. N. Goncharov, I. V. Igumenshchev,
S. T. Ivancic, and M. J. Rosenberg

References

1. H. Abu-Shawareb *et al.*, Phys. Rev. Lett. 129, 075001 (2022).

Focus Points

- Achieving high yields in ICF requires coupling more energy to the fuel compared to that of current ignition experiments on the NIF.
- Direct-drive designs with broadband lasers offer a very attractive path for high-yield platforms.

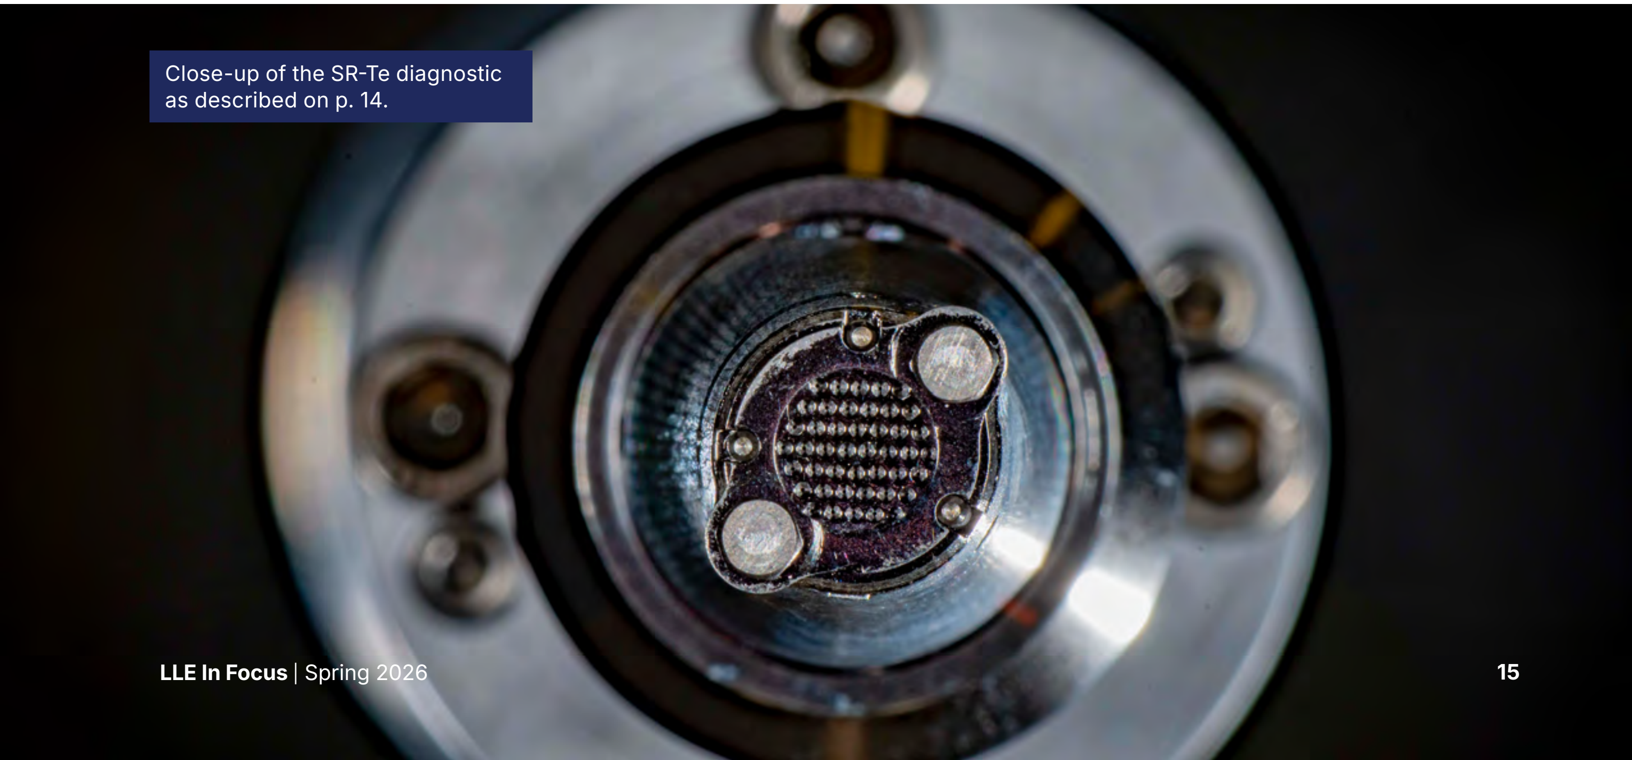

Close-up of the SR-Te diagnostic as described on p. 14.

Simulation Impossible

Inside the Computational Science Infrastructure Driving LLE's Research

Much of the most important work at LLE happens not just in a lab but in a computer, where plasma behaviors are modeled, capsule implosions are tested, and the physics of extreme conditions are interrogated. Researchers simulate experiments as close to the actual experimental conditions as possible but also have the ability to adjust parameters, explore edge cases, and probe regimes that current experiments cannot accommodate.

LLE's capacity to simulate complex physical phenomena with high fidelity is one of its more powerful research capabilities—and one that is expanding. At the center of that capability is the High-Performance Computing (HPC) Group, a team within LLE's Theory division that has built and continues to evolve a sophisticated scientific computing environment.

Simulation as a Scientific Instrument

Research at LLE has always depended on its instruments, including lasers, diagnostics, and target fabrication systems, which are each precision-engineered to reveal important information about the physics of extreme conditions. Computational simulation is an instrument in its own right and one that has for decades been central to experimental design. At LLE, simulations complement physical measurements, providing information that is not measurable in the lab. They allow researchers to see inside the physical problems and to understand not just what happened, but why, and what to explore next.

LLE's computational science infrastructure currently serves more than 150 researchers across five LLE divisions, the University of Rochester's mechanical engineering and physics and astronomy departments, as well as the Flash Center for Computational Science. Nearly half of these researchers are graduate, undergraduate, and even high school students who are developing skills in computation and computational methods, gaining through this work the ability to simulate some of the most complex physics on Earth. Training users on systems of this level of complexity is itself a significant undertaking, which is why the LLE HPC Group provides significant support in the form of documentation, hands-on training, and ongoing assistance that allows researchers at every level to work with confidence on the tools and codes they need. The range of what can now be explored computationally is growing, and LLE's HPC resources like Conesus allow researchers to answer scientific questions in ways that were out of reach a decade ago.

The DOE/NNSA has consistently been a leader in bringing high-performance computing to the mission. NNSA's Advanced Strategic Computing program, along with the ICF program, enabled LLE to acquire a new system, Conesus.

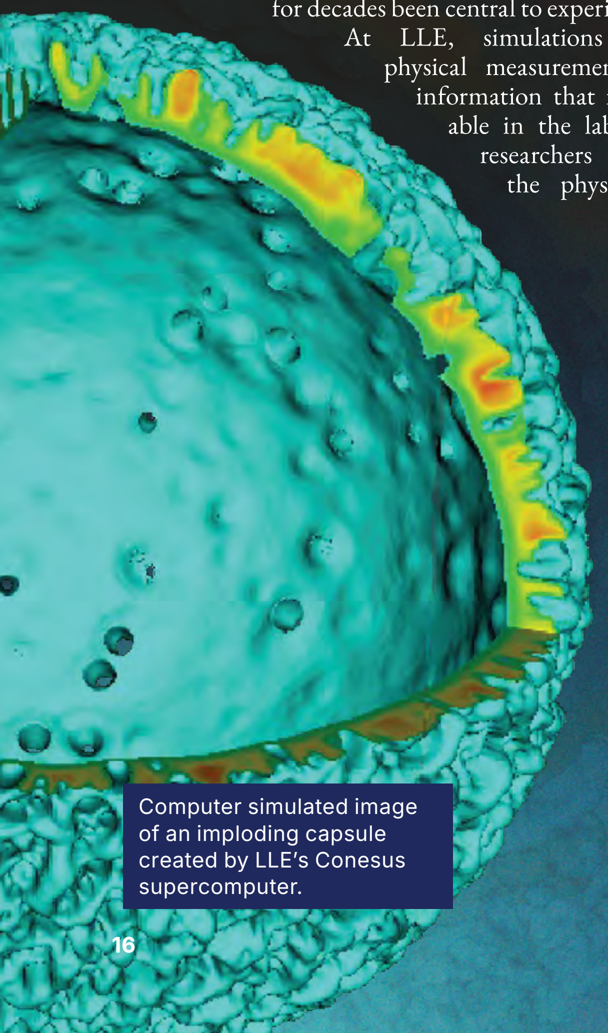

Computer simulated image of an imploding capsule created by LLE's Conesus supercomputer.

A History of HPC Computing Power at LLE by GFLOPS

(Giga floating-point operations per second, or billions of calculations per second)

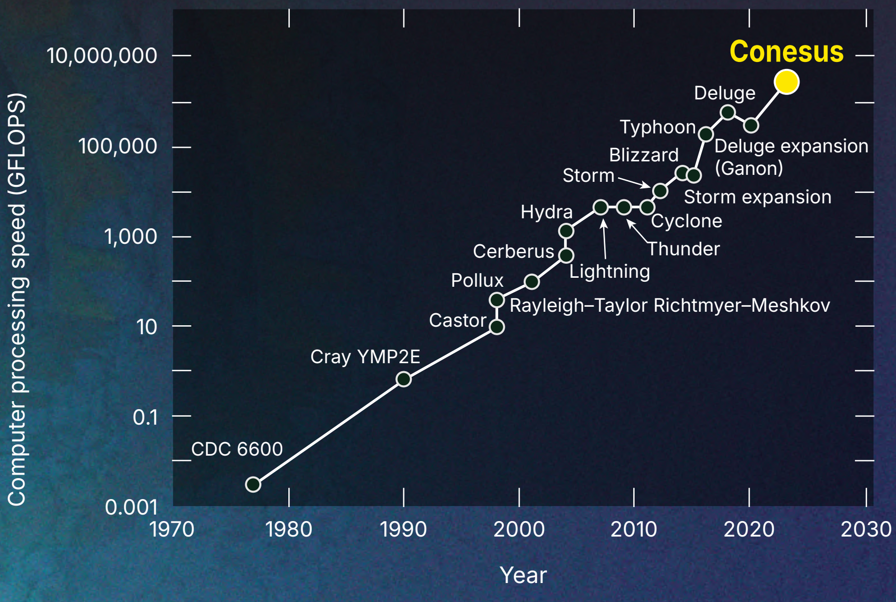

LLE's Conesus supercomputer.

The Erie storage system manages the massive datasets generated by Conesus.

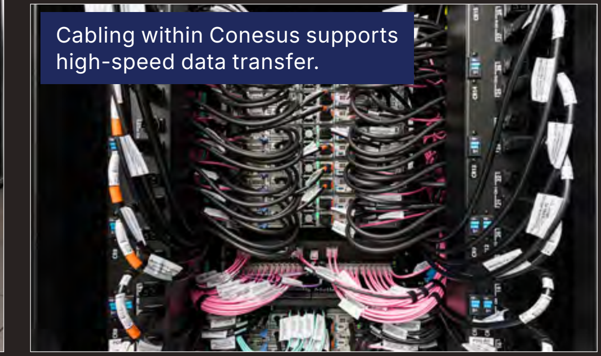

Cabling within Conesus supports high-speed data transfer.

The Scale Required to Do This Work

High-fidelity simulations demand computing infrastructure at a large scale. LLE's HPC Group currently operates approximately 800 individual computer units with 60,000 total CPU cores across four separate clusters. One cluster at LLE, Conesus, illustrates the magnitude of what this requires. At full capacity, it draws roughly 10 MWh of electricity per day—enough to power approximately 340 average American homes—and generates substantial heat that must be managed by sophisticated closed-circuit fluid-cooling systems. Despite this scale, Conesus ranked 77th on the Green500 list of the world's most energy-efficient supercomputers, a recognition that reflects how seriously LLE approaches the responsibility of operating at this size. The University of Rochester's computer center, where Conesus is housed, was purpose-built for machines like these and engineered to bear their weight, supply their power, and keep them cool.

Behind the computer processors sits an equally critical data infrastructure: petabyte-scale file servers capable of managing thousands of simultaneous read and write streams, ultrafast fiber networks connecting compute nodes, and redundant systems ensuring that when a simulation generates terabytes of data, nothing is lost. The computational instrument is only as good as its ability to capture what it produces.

Reliability as a Research Value

At the scale that LLE HPC operates, hardware and software failures are statistically inevitable. The HPC Group must not only recover from problems as they arise but also minimize their impact. The HPC infrastructure is built with redundancy throughout: multiple network paths between compute nodes, failover systems for critical services, continuous snapshots and backups of essential files, and long-term archiving for legacy data. Most issues are identified and resolved quietly, before they can affect a running simulation or delay a publication. For the researchers who depend on these systems around the clock, this invisibility is the point.

Equally invisible to most users is the software layer that makes the hardware usable. Every compute node runs an identical, carefully maintained operating system. On top of that sits a full stack of computational libraries, code compilers, debuggers, developer environments, and visualization tools—the instruments researchers use to write, test, run, and analyze their simulations. Automated job schedulers manage the flow of work across the cluster, balancing the demands of many users simultaneously while maintaining fairness and institutional priorities. Keeping all of this current, consistent, and compatible across hundreds of machines is a continuous undertaking, one that underpins every calculation that runs and allows researchers to trust that the environment will do what they expect. For science needing resources 24-7, reliability is itself a capability.

What Comes Next and Why It Matters

The computational frontier at LLE is moving fast, and the HPC Group is moving with it. The most immediate driver is scale: as simulations grow larger and more complex, the infrastructure to run and analyze them must keep pace. This means more CPUs, more memory, faster interconnects, and larger storage. At the same time, GPUs, chips originally designed for video rendering that have since become powerful tools

Members of LLE's High-Performance Computing Group, pictured from left to right: Ken Anderson, Keenan Nash, Jonathan Carroll-Nellenback, Dave Keller, and Will Ebmeyer.

CPUs and GPUs: What's the Difference?

Central processing units (CPUs) are generalists: fast, flexible processors that handle a wide variety of tasks sequentially. Graphics processing units (GPUs) were built to process thousands of simple calculations simultaneously, originally meant for rendering graphics. That same parallel architecture turns out to be exceptionally well-suited for some large-scale mathematical computations at the heart of scientific simulation.

in scientific computing, are becoming increasingly relevant for certain classes of large-scale simulation. GPUs are also highly valuable for artificial-intelligence and machine-learning (AI/ML) applications. LLE is currently procuring new, state-of-the-art GPU servers to accelerate the development of GPU-capable simulation codes and train new AI/ML models. The facilities that house these resources must also evolve to accommodate them.

Extracting meaning from the output of these simulations requires its own set of tools: from analysis software and post-processing pipelines to dedicated visualization servers capable of rendering large datasets in full 3D. These allow researchers to view and better understand the science behind what is happening in their simulations. Identifying trends from the increasingly large datasets of tomorrow will require more and more compute horsepower as well as fast storage and processing tools. What gets built next will be shaped by the scientific questions that LLE's researchers are asking today, which will in turn shape what questions become answerable tomorrow.

Expanding the Edge of What's Possible

LLE's computational science capability is distinctive, not because of any single cluster or software stack but by the combination of scale, rigor, security, and foresight that allows researchers across the institution to do work that would otherwise be out of reach—as well as the steady effort to push that boundary further.

The simulations running today are more detailed, more accurate, and more revealing than anything possible five years ago. The researchers running them—the faculty pushing the boundaries of fusion theory and the students encountering the full scale of the problem for the first time—are the reason that progress occurs. In the Meliora spirit of "Ever Better," the computational infrastructure at LLE is not something simply to be maintained and preserved. It is something to be built upon, year after year, in step with the scientists it serves.

Corresponding author: K. S. Anderson

Focus Points

- Computational modeling of fusion capsule implosions, plasma behavior, and other complex phenomena provides insights that experiments alone cannot provide.
- LLE HPC operates ~800 high-performance computers with 60,000 CPU cores and multipetabyte, high-speed storage, including Conesus, one of the world's most energy-efficient supercomputers, delivering the reliability and security required for continuous, high-impact research.
- LLE is expanding its use of GPUs, originally designed for video rendering, to support simulations, 3D visualization, and machine learning, and to advance scientific modeling and accelerate discovery.

The Opportunity of a Lifetime

Students Gain Rare Access to Some of the Country's Leading Science Facilities

For most undergraduates, national laboratories exist only as names on research papers or distant career possibilities. This past January, 27 students in LLE's Undergraduate Research Program got something rarer: direct access to three of the facilities where some of our nation's most important science and engineering happens.

The students traveled to Lawrence Livermore National Laboratory (LLNL) in California, Los Alamos National Laboratory (LANL) in New Mexico, and the Nevada National Security Site (NNSS), visiting scientists, touring restricted facilities, and connecting their own research projects to the broader national laboratory mission. For many students, the experience is eye-opening—not just for what they see, but for what it reveals about the scale and purpose of the work.

The LLE Undergraduate Research Program and experiences like these national laboratory visits are designed to address a critical gap in the US scientific workforce. Each year, billions of dollars in capability across the national complex go underused, not due to a lack of scientific ideas or PhD-level expertise, but because the nation needs more engineers and technical professionals than those currently available to execute these missions at scale. The program provides students with the opportunity to develop the practical skills required to contribute to large, complex scientific efforts while also exposing them to environments far beyond their typical experience. For many, this is their first encounter with the scale, responsibility, and impact of work conducted at the national laboratories, broadening both their perspective and their sense of what is possible.

This year, the program helped facilitate research opportunities for 15 students who previously participated in the LLE program, placing them at institutions across the country, including LLNL, LANL, SLAC National Accelerator Laboratory, and NNSS.

Lawrence Livermore and the National Ignition Facility

At LLNL, students toured the National Ignition Facility (NIF), home to a large laser system designed to compress and heat hydrogen fuel to fusion-ignition conditions. The NIF has achieved ignition—the point at which a fusion reaction produces more energy than the laser energy delivered to the target—several times in recent years. For students working on laser and plasma physics experiments at LLE, the visit illustrated how their research connects to work being carried out at larger facilities.

"I learned a lot about the laser technology that has been developed for successful ignition shots," said Mason Weiss, one of the participating students. "I also greatly appreciated the chance to talk to some researchers about the work culture and opportunities available at LLNL."

LLE undergraduate students at the National Ignition Facility during their visit to Lawrence Livermore National Laboratory.

Los Alamos and the History of Nuclear Research

Los Alamos has played a central role in nuclear science since the Manhattan Project and today remains an active research institution working across weapons, energy, and basic science. Students learned about the laboratory's history and how current programs there relate to research underway at LLE, along with practical information about summer internship pathways.

The Nevada National Security Site and Stockpile Stewardship

The Nevada National Security Site, where the United States conducted nuclear tests until 1992, now supports stockpile stewardship—the ongoing scientific effort to certify that the nation's nuclear weapons remain safe and reliable without underground testing. The visit gave students a clearer picture of LLE's collaborative role in that mission and the kinds of research careers it supports.

Sandia National Laboratories and Pulsed-Power Science

In previous years, students also visited Sandia National Laboratories (SNL) in Albuquerque, NM. At SNL, a leader in national security science and engineering with world-class pulsed-power facilities used to study materials and plasmas under extreme conditions, students make connections between the work at SNL and research at LLE, and are introduced to a wide variety of internship and career opportunities.

Building Connections for the Next Generation

The national laboratory trip has been part of LLE's undergraduate program for several years. Over the past three years, nearly 60 students have visited LLNL, LANL, NNSS, and SNL. Next year's trip will include Savannah River National Laboratory in South Carolina.

The visits are intended to help students understand the national laboratory system, make professional connections, and consider career paths they might not otherwise encounter. Several students from previous trips have gone on to summer internships at the facilities they visited. The experience also gives students a clearer sense of how large-scale research environments operate and where their own interests might fit within them.

"The trip opens their eyes to opportunities in science, technology, and engineering outside of their limited experiences," said Dustin Froula, Director of LLE's Plasma & Ultrafast Laser Science & Engineering Division. "Students are already taking advantage of the contacts they've made to do summer internships at the national labs, which will help line up future jobs at these institutions once they graduate."

For students at an early stage of their careers, the exposure can be clarifying. "Seeing these facilities in person, and talking to the people who work at them inspires me to keep pursuing work that is both interesting and has a positive impact on our communities and country," said undergraduate Ethan Mentzer. "The lab trip helped me gain lots of insight into how I might structure my future." Froula offered a simpler summary: "The experience was extraordinary—the kind that imprints itself for life."

Students on the viewing platform of the Sedan Crater at the Nevada National Security Site. The largest man-made crater in the US, it was formed by a 104-kton underground nuclear test on July 6, 1962. The blast displaced 12 million tons of earth, creating a hole ~12,800 ft. wide and 320 ft. deep.

The Principal Underground Laboratory for Subcritical Experimentation (PULSE) complex at the Nevada National Security Site.

Students tour Los Alamos National Laboratory.

This material is based upon work supported by award numbers Department of Energy (DOE) National Nuclear Security Administration DE-NA0004144; DE-SC0024863, DE-SC0025573, DE-SC0024415, DE-SC0025517, DE-SC0021057 and DE-SC0023783; US National Science Foundation 2329970; Department of Defense; Air Force Office of Scientific Research FA9550-24- 1-0160 and FA9550-25-1-0241; Empire State Development Corporation C230113; Lawrence Livermore National Laboratory B671119; Savannah River National Laboratory Subaward 800006058; Lawrence Berkeley National Laboratory Subcontract 7742075; Commissariat à l'énergie atomique 3957; Johns Hopkins University Subaward 2004758202; European Research Executive Agency Project 101095207; Princeton University SUB0000841; Duke University DE-SC0022573; Sydor Technologies; Seurat Technologies; and Focused Energy. Sponsor information is subject to ongoing changes and reflects details available as of the publication date.