LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-JRNL-402612

The Mercury Laser System-A scaleable average-power laser for fusion and beyond

C.A. Ebbers, E.I. Moses

March 28, 2008

Photonics Spectra

Disclaimer

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

The Mercury Laser System—A scaleable average-power laser for fusion and beyond
Lawrence Livermore National Laboratory

By Chris Ebbers and Edward Moses

Introduction

Nestled in a valley between the whitecaps of the Pacific and the snowcapped crests of the Sierra Nevada, Lawrence Livermore National Laboratory (LLNL) is home to the nearly complete National Ignition Facility (NIF). The purpose of NIF is to create a miniature star—on demand. An enormous amount of laser light energy (1.8 MJ in a pulse that is 20 ns in duration) will be focused into a small gold cylinder approximately the size of a pencil eraser. Centered in the gold cylinder (or hohlraum) will be a nearly perfect sphere filled with a complex mixture of hydrogen gas isotopes that is similar to the atmosphere of our Sun. During experiments, the laser light will hit the inside of the gold cylinder, heating the metal until it emits X-rays (similar to how your electric stove coil emits visible red light when heated). The X-rays will be used to compress the hydrogen-like gas with such pressure that the gas atoms will combine or "fuse" together, producing the next heavier element (helium) and releasing energy in the form of energetic particles. 2010 will mark the first credible attempt at this world-changing event: the achievement of fusion energy "break-even" on Earth using NIF, the world's largest laser!

NIF is anticipated to eventually perform this immense technological accomplishment once per week, with the capability of firing up to six shots per day—eliminating the need for continued underground testing of our nation's nuclear stockpile, in addition to opening up new realms of science. But what about the day *after* NIF achieves ignition? Although NIF will achieve fusion energy break-even and gain, the facility is not designed to harness the enormous potential of fusion for energy generation. A fusion power plant, as opposed to a world-class engineering research facility, would require that the laser deliver drive pulses nearly 100,000 times more frequently—a rate closer to *10 shots per second* as opposed to *several shots per day*.

Mercury—a laser system designed for high repetition rate and scalability

NIF is a world-class facility for both science and engineering. The headquarters of laser facility expertise, it coordinates the cleaning, inspection, and installation of over one million small, medium, and large optics. Riding on the coat tails of this unique project are a number of smaller laser systems, including the Mercury laser facility. Mercury (named after the Roman God, not the element) is not nearly as energetic as NIF; it only has the potential to generate roughly 1/300th as much energy as a single beamline of its behemoth cousin. The Mercury laser is currently ramping up its output and is now operating at 60% of its 100-J design capability, in contrast to the 18,750-J (at 1 μm) output of one of NIF's 192 beams. Of what use is this relatively small laser system? The Mercury laser is designed to demonstrate the feasibility, reliability, and scalability of a 10-Hz, 100-J-class, diode-pumped laser system. Mercury is a prototype laser capable of scaling in aperture to a NIF-like beamline, while still running at a repetition rate 100,000 times greater.

NIF and Mercury—side by side

In Table 1, we compare the components of one beamline of the National Ignition Facility to the single beamline of the Mercury laser system. An 19 kJ beamline of NIF has an area of roughly 38 cm x 38 cm, or approximately 13 J of energy for every square centimeter of aperture. NIF uses 16 slabs of neodymium-doped phosphate glass (Nd:glass)—each slab is 45 cm x 78 cm x 4.2 cm thick. The slabs are oriented at Brewster's angle to minimize transmission losses, so the effective path length is 5 cm per slab. As shown in Figure 1, this NIF beamline is seeded by a precision single-mode fiber laser, tuned to the laser gain peak at 1053 nm. Large flashlamps store energy in the 16 slabs of laser glass, providing the 19 kJ of laser energy for each beamline. In the main amplifier, 11 of the slabs are multi-passed four times, while the five remaining slabs comprising the power amplifier are double passed. The output of each beamline is frequency doubled and tripled (generating green and ultraviolet light from the incident infrared). The NIF system will meet its beam quality specifications by radiative and convective cooling alone. The overall efficiency of the laser, from wall plug to laser light on target, is approximately 0.7%. The NIF building, with a footprint equal to three football fields, is essentially built with the same stringent stability as an optical laser table.

In contrast, the Mercury laser system occupies a room the size of a handball court. Mercury's architecture is similar, in some respects, to that of a single NIF beam. A precision fiber-based laser serves as the front end. The gain medium is multi-passed four times. Image relaying and spatial filtering are used for beam transport. A frequency conversion module sits at the end of the laser system. In contrast, the aperture of the Mercury laser system is 3 cm x 5 cm. Despite a smaller aperture, Mercury's designed operating fluence (light energy per square centimeter) has an average similar in magnitude to that of NIF: 9 J/cm^2 at 1047 nm for Mercury versus 13 J/cm^2 at 1053 nm for NIF. It is in the details that the differences between the two laser systems become apparent.

Eliminating sources of heat

By changing the laser pumping mechanism from flashlamps to laser diodes, the electrical efficiency can improve by more than an order of magnitude. Flashlamps emit white light in a manner similar to a camera flash. Laser diodes, similar in structure to the ubiquitous red laser pointer, efficiently convert electricity into single-color directional light. NIF uses flashlamps, which currently are the only economical way to produce megajoules of stored energy. The Mercury laser system uses diode lasers, the pumping mechanism of the future, to provide the stored energy. The two methods are compared in Figure 2.

Mercury has two laser amplifiers. Each amplifier assembly is pumped by four 100,000-W laser diode arrays to provide the necessary gain for a maximum output capability of 100 J. The diode light is first angularly redistributed by hollow concentrating optics to a 3x5 cm aperture in which the angular divergence of the beam is made nearly identical in both the x and y axes. The diode light is then guided by a rectangular reflective cavity called a "homogenizer", which maintains the angular distribution while smoothing the spatial pump profile. The face-cooled helium gas amplifiers efficiently remove approximately 3 W/cm^2 of heat with minimal thermal wavefront distortions. Due to the current high cost

of diodes, pulsed high-energy lasers are not generally diode pumped. Diode lasers are currently expensive, priced at roughly $1.60 per peak Watt of output power, not including the cost of the power supply or electrical pulse generators. However, it is expected that as demand grows, diode prices will drop from dollars per Watt to pennies per Watt.

The design of a complex laser facility is driven by multiple, sometimes conflicting, issues. We concentrate on three issues that have the most direct impact on efficiency and cost regarding the selection of a laser gain medium. These issues are: quantum defect, upper lasing level lifetime, and emission cross section. Essentially, these can be translated to the amount of energy (heat) left in the laser gain medium after efficient lasing, the number of diode pump lasers required to effectively provide the gain, and the amount of laser light (fluence) needed to circulate through the laser gain medium to efficiently extract the stored energy. Heat deposition, fracture toughness, and thermo-mechanical effects in the gain medium play a secondary role in the design of an energetic laser system. The design of the gas cooling system for the amplifier slabs must account for these material parameters. In principle, the gain medium of a high-average power laser can be nearly any glass, crystal, or transparent ceramic laser host.

The Mercury laser gain medium is crystalline strontium fluoro-apatite doped with ytterbium (Yb:S-FAP). The more well-known crystal from the apatite family, calcium hydroxy phosphate, is the material of teeth and bones. Strontium can readily substitute for calcium, while fluorine readily substitutes for the hydroxyl ion (fluorine substitution is the same mechanism by which teeth are made more resistant to decay). The ytterbium ion provides the laser gain. The ytterbium dopant absorbs diode light near 900 nm and emits near 1047 nm. The mere fact that a 900 nm photon is absorbed and produces a 1047 nm photon implies that even if the laser is running at its highest efficiency, at least 15% of the energy deposited into the Yb:S-FAP host is converted to heat that must be removed. Another laser gain medium, Nd:phosphate glass, absorbs diode light at 808 nm and emits laser light at 1053 nm, leaving 24% of the incident diode pump light converted to heat. Eventually, this heat must be removed from the laser material, whether it is crystalline, glass, or a transparent ceramic.

The Mercury laser system diodes are pulsed on for 900 millionths of a second, at a rate of 10 pulses per second. Thus, the diode lasers are only “on” for 1% of the time, and off 99% of the time, for a duty cycle of 1%. (In contrast, a continuous diode laser such as a laser pointer is “on” for the length of time that you hold the button—effectively a duty factor of 100%.) The duty cycle is limited by the upper laser level lifetime: for Yb:S-FAP, the upper-level lifetime is 1,100 millionths of a second (1,100 μs). In comparison, the common laser crystal, Nd:YAG, has an upper laser level lifetime of 240 microseconds, less than ¼ as long. This lifetime represents the length of time an atom will remain in the excited state, storing energy that can be usefully extracted in the form of laser light energy. Diode lasers efficiently convert electricity into optical power (energy per unit time), whereas pulsed lasers emit stored energy. To extract 100 J from the Mercury laser gain crystals, a *minimum* of 115 J must be absorbed in the crystals, due to the 15% energy loss between the pump-diode photon energy and the laser emission photon energy. Enough diodes must be purchased so that at least 115 J can be delivered in the available

1,100 microseconds (> 100,000 W). Diodes are continuous wave devices—they do not store energy. The diode arrays on the Mercury laser system would exceed 100,000 W if they were run continuously. Alternatively, if the upper laser level were four times shorter than 1,100 microseconds (as is the case for Nd:YAG), the system would require over four times as many diodes to store the same amount of energy.

Mercury uses diode bars purchased in 2002 that are rated at 120 W per bar at 45% electrical efficiency (200-W bars at 65% efficiency are now standard products from many manufacturers). The diode bars are assembled into tiles (23 diode bars per tile). Thirty-six tiles are assembled into a single array, and eight arrays, rated at over 100 kW of continuous diode light, provide the gain for the Mercury laser system. (Laser diode divergence, losses from diode light transport, amplified spontaneous emission losses, and residual energy left in the medium drive the need for more diode light than the theoretical 100 kW minimum mentioned above). A laser gain crystal with a shorter laser lifetime requires even more diode bars. Yb:S-FAP was chosen for the exceptionally long lifetime, minimizing the cost of the laser diodes needed for the system. As diode prices continue to drop in price due to their increased popularity with the industrial laser market, upper state lasing lifetime becomes less of an issue, and we anticipate that Nd:glass or a transparent ceramic host may be a suitable gain material for future laser systems.

A final laser design issue concerns the tradeoff between the stimulated emission cross section, the maximum convenient laser aperture, the maximum extractable energy per aperture, and the fluence needed to extract this energy. For efficient extraction of laser energy, a laser is typically operated at three times the saturation fluence. The saturation fluence, a parameter that is specific to the individual laser gain crystal, is the laser energy per unit area needed to extract 63% of the stored energy in a single pass. LLNL, in combination with Konishima and Baikowski, recently demonstrated the largest aperture Nd:YAG ceramic laser crystals in the world. Nd:YAG has a saturation fluence of 0.62 J/cm^2. On the other hand, the Nd:phosphate glass used on NIF has a saturation fluence of 4.3 $Joules/cm^2$. The nearly eight times lower saturation fluence of Nd:YAG has implications regarding the storage of energy. It is difficult to store energy in a laser system based on the Nd:YAG laser gain crystal. As the aperture of the Nd:YAG gain crystal is increased, it becomes more probable that amplification of spontaneous emission will occur across the aperture, emitting energy through the edges of the gain crystal, rather than storing the energy until the desired laser pulse extracts the stored light energy. NIF has a 40-cm aperture beamline storing over 36 kilojoules of energy. If ceramic Nd:YAG were to replace the Nd:glass, multiple small apertures (each aperture smaller than the current 40-cm aperture of NIF) would be needed. In contrast, Yb:S-FAP has a saturation fluence of 3 J/cm^2, more similar to Nd:glass than Nd:YAG.

The saturation fluence is one example where the lasing efficiency drives an aspect of the laser design. A high saturation fluence (well above that of Nd:glass) can store tremendous amounts of energy per beamline, but requires a high laser fluence (possibly high enough to cause optical damage) in order to extract the energy. In contrast, a low saturation fluence leads to many beamlines. It is likely that a contemporary facility will use a material with a saturation fluence similar to that of NIF, using the tremendous advances

in beamline design, fabrication methods, and optical reliability data to fullest advantage. The Mercury laser system uses this developed expertise. Yb:S-FAP was chosen for its similarity to Nd:glass in many aspects (except availability). The combination of long lifetime, moderate drive fluence and aperture scaling, and, especially, the lower energy defect drove the choices from the viewpoint of reducing the cost of diodes required for a 100-J-class laser. A different pricing structure for diodes could make a different material more attractive.

Cooling, average power operation, and frequency conversion

Roughly 15% to 26% of the energy stored in the gain medium is left in the laser crystal as heat, depending upon whether the laser ion is Ytterbium or Neodymium, respectively. For continuous operation, the residual heat must be removed. Mercury uses high-pressure helium gas, which flows over the laser crystal in a manner similar to the way air flows around the wing of an airplane. Helium is chosen for two unique properties—light passing through helium gas experiences little refraction, and helium has high thermal conductivity. Refraction is the property observed on a hot summer day as the light shimmers near the road or roof of a car. The shimmering is caused by small changes in the angle (refraction) of the light direction due to density differences between hot and cold air. As a very light gas, helium does not affect the beam quality of the laser light. Secondly, helium also has high thermal conductivity, the ability to remove heat. A properly designed gas cooling system allows the heat to be efficiently removed from the laser gain crystal by transfer to the helium gas. The closed-cycle helium gas passes through a water-cooled heat exchanger, and is then cycled back to the laser gain medium. One advantage of gas cooling is that the technology is scaleable in aperture. Once a system design has been formulated to remove a certain amount of heat per square centimeter, it is possible to scale in aperture, allowing a small-energy (100 J) laser system to operate with the same characteristics as a system with 100 times more energy (i.e. 10,000 J). This heat removal method is applicable not only to the laser gain medium, but also to the frequency conversion system, where the infrared laser light emission is converted to green or ultraviolet light (frequency doubling or tripling).

Operations

A hidden but critical aspect of a laser facility is the control system. NIF has a control system that allows for remote inspection and alignment of all the optical subsystems, as well as the data collection of the critical beam parameters. The Mercury laser system also depends on a system of remote alignment and inspection diagnostics. In particular, critical laser beam optics are under continuous inspection. It is not possible to shut down a laser system manually when operating at a repetition rate beyond a few shots per second. Instead, control regarding optical inspection is handed over to an image processing system. Data from the previous laser shot is compared to the current laser shot. Any significant differences are flagged, allowing the system to be safely stopped (if necessary) before the next laser shot is initiated, preventing potential damage to more than a single optic. Such a data collection and examination system will be the real brains behind a potential multiple-beamline power-generating system running at a high repetition rate, allowing the safe shutdown and refurbishment of any optical components, as needed.

The future looks bright

We have examined the design of the Mercury laser system in comparison to its much larger parent, NIF. Mercury is a diode-pumped, gas-cooled, scaleable laser that has operated to date for 0.3 million shots at energy levels above 50 J at a repetition rate of 10 shots per second. Engineering advances are under way to produce diode pump arrays with a higher cost-effectiveness and advanced laser gain materials more suitable to the construction of a larger-aperture, higher-energy facility. When NIF ignites a fusion target, it is hoped that these technological achievements will pave the way to a reduced-carbon-impact future, based on a high-repetition-rate laser system. As always, the future looks bright for advanced laser applications!

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

	Mercury	NIF
Beamlines	1	192
Energy per beamline at 1ω (design)	100 J	18.75 kJ
Fundamental wavelength	1047 nm	1053 nm
Gain medium	Yb:SFAP	Nd:Glass (phosphate)
Gain medium aperture	3 cm x 5 cm	40 cm x 40 cm
Number of slabs	14 (4 pass)	11 (4 pass) + 5 (2 pass)
Slab thickness	0.75 cm	4.2 cm (effective 5 cm)
Total gain medium thickness	42 cm	270 cm
Repetition rate	10 shots per second	4-6 shots per day
Cooling	High velocity helium gas	Radiative and Convective
Beam propagation	Image relayed	Image relayed
Pockels cell	KD*P	KDP
Spot size	5 times diffraction limited	5 times diffraction limited
Gain excitation	Diode (900 nm)	Flashlamp
Optical-Optical efficiency	7%	0.7%
Building area	Handball court (800 sq.ft.)	3 football fields (230,000 sq. ft.)

Table 1. Comparison of the Mercury laser system to the National Ignition Facility
The goal of the National Ignition Facility is to achieve fusion breakeven. In contrast, the goal of the Mercury laser system is demonstration of a reliable, diode-pumped, solid-state laser system capable of scaling in aperture to the equivalent energy of a single beamline of the 192 beamlines of the National Ignition Facility. Diode pumping reduces the heat deposited into the Mercury laser, while gas cooling allows the residual heat to be efficiently removed.
Error!

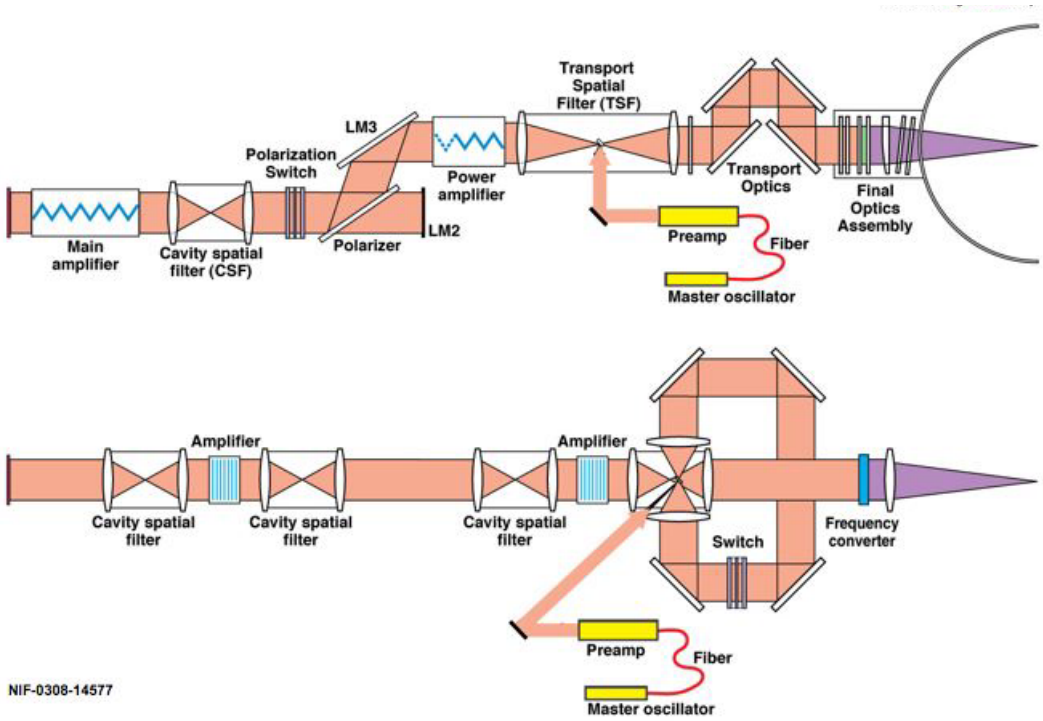

Figure 1. A schematic of a single beamline of the flashlamp-pumped National Ignition Facility (top) compared to the single beamline of the diode-pumped Mercury laser (bottom).

In NIF, the flashlamps provide the gain for the main and power amplifiers. A fiber-optic laser provides a pulse with the precise temporal shape. This is amplified in the preamplifier to an energy level of 5 Joules and directed towards the power amplifier. After passing through the power amplifier, the Plasma Electrode Pockels Cell (PEPC switch) traps the light in the main amplifier cavity for four passes. The PEPC is turned off and the pulse passes through the power amplifier. The light that has been amplified from picojoules to 20 kilojoules is transported to the target chamber, frequency doubled and tripled, before finally interacting with the target. NIF has the capability to perform this feat six times a day.

Many elements in the single beam line of Mercury (shown below) are similar in concept to the NIF beamline. To reduce thermal loading, diode lasers provide the gain for the two amplifiers. Again, a fiber-optic laser provides a pulse with the precise temporal shape. This is amplified in the preamplifier to an energy level of 0.5 Joules and injected in towards the first amplifier. The beam passes through the first and second amplifier, and is reflected back through both amplifiers. The beam path is reversed for two more passes through both amplifiers. The laser beam, after being amplified from picojoules to the 60-Joule level is then directed towards the frequency converter. Using high-velocity gas cooling to remove the residual heat, the Mercury laser system has operated for more than

0.3 million shots at an energy level over 50 Joules at a repetition rate of 10 shots per second.

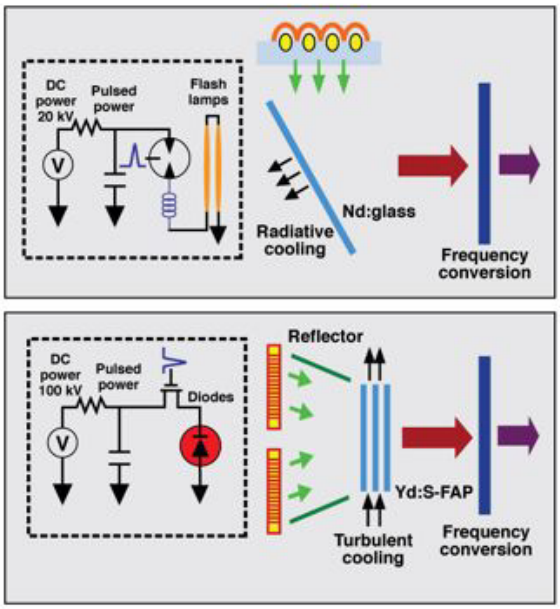

Figure 2. As shown in the top figure in the schematic above, the National Ignition Facility is a flashlamp-pumped laser system, intended to fire up to six shots per day (0.0001 Hz).

To increase the repetition rate by 100,000 (10 Hz), the Mercury laser system, shown in the bottom schematic, must use active cooling. Use of diode lasers (as opposed to flashlamps) reduces the thermal loading on the laser gain medium. High-speed helium gas is circulated across the amplifiers, allowing efficient removal of the residual heat.

Figure 3. The Mercury laser system, a gas-cooled, diode-pumped, solid-state laser, operates in a facility about the size of a handball court.