Approaching hydro-equivalent ignition in laser direct-drive via target design optimization using novel statistical modeling

A. Lees,[1,2] R. Betti,[1,2,3] V. Gopalaswamy,[1,2] J. P. Knauer,[1] L. A. Ceurvorst,[1] D. Patel,[1,2] R. Ejaz,[1,2] K. S. Anderson,[1] K. A. Bauer,[1] M. J. Bonino,[1] D. Cao,[1] K. Churnetski,[1] T. J. Collins,[1] P. S. Farmakis,[1,2] C. J. Forrest,[1] V. N. Goncharov,[1] D. R. Harding,[1] P. V. Heuer,[1] I. V. Igumenshchev,[1] R. T. Janezic,[1] S. P. Regan,[1] M. J. Rosenberg,[1] S. Sampat,[1] R. C. Shah,[1] C. Stoeckl,[1] C. A. Thomas,[1] K. M. Woo,[1] M. Gatu-Johnson,[4] C. Li,[4] J. A. Frenje,[4] and C. W. Wink[4]

[1)] *Laboratory for Laser Energetics, University of Rochester, Rochester, NY 14623, USA*
[2)] *Department of Mechanical Engineering, University of Rochester, Rochester, NY 14611, USA*
[3)] *Department of Physics and Astronomy, University of Rochester, Rochester, NY 14627, USA*
[4)] *Plasma Science and Fusion Center, Massachusetts Institute of Technology, Cambridge, MA, USA*

(*Electronic mail: alee@lle.rochester.edu)

(Dated: 24 February 2025)

Laser direct-drive offers significant advantages in terms of target simplicity, improved energy coupling and large fuel masses over indirect drive. However, performance degradations from hydrodynamic and laser-plasma instabilities seeded and driven by the direct illumination pose limitations on the parameter space available for achieving ignition. In this paper, new design improvements are identified to forge a path forward for a hydro-equivalent ignition demonstration. The first is related to a new formulation of the statistical model (SM) used to accurately predict target performance directly from input parameters such as laser pulse shape and target specifications. This new SM formulation provides direct guidance on target dimensions and laser beam to target radius to achieve the highest fusion yield on the OMEGA laser. The second improvement comes from cooling the deuterium-tritium (DT) ice layer below the triple point right before shot time leading to lower DT vapor densities and higher convergence. Guided by these design improvements, a Bayesian optimization algorithm was used to design an implosion that is predicted to reach a Lawson triple product that hydrodynamically scales to ignition if equivalent laser-target coupling is achieved at laser energies typical of the National Ignition Facility.

I. INTRODUCTION

In laser-driven inertial confinement fusion (ICF)[1] a spherical capsule containing a deuterium-tritium (DT) fuel mixture is imploded in order to achieve the conditions required for fusion ignition. The capsule consists of a shell, typically composed of an outer mid-Z layer (ablator) and a solid DT ice layer surrounding a central low-density DT vapor region. The goal is to accelerate the shell inwards and form a high-temperature compressed core (hotspot) via pdV work done by the imploding shell and initiate D+T fusion reactions. The performance of an implosion is measured by a generalized Lawson parameter[2], which can be defined in terms of experimentally measurable quantities as[3]

$$\chi_{\mathrm{no}\alpha} = \left(\frac{0.12Y \times 10^{-16}}{M_{\mathrm{stag}}}\right)^{0.34} (\rho R)^{0.61}, \qquad (1)$$

where the fusion yield Y is measured as the number of neutrons produced in the D+T reactions, M_{stag} is the mass stagnating at peak neutron rate in milligrams and $\rho R = \int \rho \mathrm{d}R$ is the areal density in g/cm^2. The subscript noα indicates that the quantities on the right-hand-side are to be measured in the absence of alpha-heating. A burning plasma regime, where the energy deposited into the hotspot by the alpha-particles produced in the D+T fusion exceeds the pdV work done by the shell is achieved at $\chi_{\mathrm{no}\alpha} > 0.80$, whereas ignition occurs at $\chi_{\mathrm{no}\alpha} \approx 1.0$.

The two main approaches to laser-driven ICF are: (1) direct drive[4], where the laser light is focused directly on the target surface and (2) indirect drive[5], where the implosion is driven by x-rays that are produced by heating a high-Z *Hohlraum* with the laser. The latter approach is used at the National Ignition Facility (NIF)[6], where ignition and fusion yields up to 5.2 MJ have been demonstrated with 2.2 MJ of laser energy[7]. Direct drive offers advantages in terms of improved energy coupling, large fuel masses and target simplicity over indirect drive. However, hydrodynamic and laser-plasma instabilities limit the parameter space available for achieving ignition. Direct drive experiments are currently pursued on the OMEGA[8] laser at 30 kJ of energy on target, which is not capable of producing a hotspot of sufficient scale to demonstrate meaningful alpha-heating. In Ref.[9], it was shown that the highest-performing OMEGA implosions would reach the burning plasma regime with a generalized Lawson parameter $\chi^*_{\mathrm{no}\alpha} = 0.86 \pm 0.02$, when hydro-equivalently scaled in size by a scale-factor of 4.2. When equivalent laser-target coupling is assumed, this scale-factor corresponds to a driver energy of 2.15 MJ, typical of the NIF laser. In this paper we present design improvements that led to a small improvement in the hydro-equivalently scaled Lawson parameter to $\chi^*_{no\alpha} = 0.88 \pm 0.05$: (1) a new statistical prediction model that enables direct optimization of design parameters, (2) subcooling, i.e., lowering the temperature of the target below the DT triple point prior to shot to reduce the ini-

tial vapor density and achieve higher convergence. The new statistical model was then used to identify a potential path to hydro-equivalent ignition $\chi^*_{noa} \approx 1.0$ when scaled to a laser energy of 2.15 MJ.

The method of statistical modeling (SM)[10] has been used to guide the design of cryogenic implosions on OMEGA and was a major component that led to development of the high-performance implosion designs that reached hydro-equivalent burning plasma[9] and hotspot gain[11]. The statistical model attempts to bridge the gap between imperfect simulations and experiments by finding a functional mapping from simulation outputs to experimental observables. In this paper, we present a new form of the statistical model, which defines the mapping in terms of "design parameters", i.e., independent input parameters to the experiment that can be freely chosen by the designer. Examples of such include the target dimensions, laser energy and shell adiabat. The new formulation of the SM enables direct optimization of the design parameters without having to run simulations over a vast parameter space. The new model and its findings for optimizing the fusion yield are discussed in Sec. II. The model suggests that the fusion yield on OMEGA can be maximized: (1) at a target outer radius of $R_t = 480$ μm, (2) using an Si-doped outer ablator layer and (3) reducing the target thickness compared to current best designs.

Applying the same SM techniques to modeling the areal density has proved more difficult due to its anisotropic nature and higher measurement uncertainty. In Sec. III, a mapping of the areal density to measurements of target convergence is used to infer the dependencies of the areal density on a subset of design parameters, namely the target layer thicknesses. The model suggests that for fixed convergence, the measured areal density is proportional to both ice and ablator thicknesses in good agreement with 1D simulations. As the yield dependence on target thickness is inverse to the areal density dependence it follows that in order to achieve a significant increase in the Lawson parameter $\chi_{\mathrm{no}\alpha}$ it is necessary to increase the convergence.

ICF implosions designed for high convergence traditionally rely upon low entropy of the shell to ensure high compressibility. While these types of designs perform well in 1D simulations, they suffer from high in-flight aspect ratios (IFAR) and low ablative stabilization[12,13] of the Rayleigh-Taylor instability seeded by laser imprint, capsule nonuniformities and engineering features. An alternative method to increase convergence that does not modify the in-flight stability properties of the shell, is to lower the density of the vapor by subcooling the target below the DT triple point. As the temperature is lowered a larger fraction of the DT vapor is frozen onto the inner ice surface and the initial vapor density is decreased by roughly a factor of two for every 1 K of subcooling. While subcooling the target is expected to increase the convergence it introduces a number of potential issues that can adversely affect the performance. These are discussed in Sec. IV where experimental results are presented from several recent campaigns studying the effects of subcooling on cryogenic implosions. The best-performing subcooled implosions that exhibited good symmetry achieved a $\sim 10\%$ increase in the generalized Lawson parameter.

The next major milestone of the cryogenic optimization campaign on OMEGA is achieving hydro-equivalent ignition, i.e. a $\chi^*_{\mathrm{no}\alpha} \approx 1.0$. Putting together all the findings from this paper, the design changes needed to optimize the fusion yield in isolation would inevitably lead to a reduction in areal density. Thus the thinner layer thicknesses for both the ice and ablator as suggested by the yield predictions can be combined with subcooling to "recover" the areal density through higher convergence. Section V presents the results from a Bayesian optimization algorithm[14] guided by the statistical model to find the best pulse shape for the optimized target defined in Sec. II. The new implosion design is predicted to approach ignition when hydro-equivalently scaled to 2.15 MJ of laser energy.

II. OPTIMUM DESIGN PARAMETERS TO MAXIMIZE THE FUSION YIELD

Predictive models based on the statistical mapping method have been used successfully to guide performance improvements in cryogenic implosions on OMEGA[10,15,16]. In Ref.[15], a statistical model was constructed to predict the yield-over-clean (YOC), i.e., the fusion yield in an experiment Y_{exp} as a fraction of the calculated yield Y_{1D} from a one-dimensional radiation-hydrodynamics simulation (LILAC[17]). The prediction Y_{pred} for the fusion yield in the experiment can then be calculated as:

$$Y_{\mathrm{pred}} = \mathrm{YOC}_{\mathrm{pred}} \times Y_{\mathrm{1D}}, \tag{2}$$

where the model for $\mathrm{YOC}_{\mathrm{pred}}$ is defined in terms of simple piece-wise power-laws, while Y_{1D} as an integrated simulation output is a very complicated non-linear function of the inputs to the simulation. Such a formulation can be interpreted as a model of degradation from an ideal 1D implosion, and each of the terms in the statistical model for $\mathrm{YOC}_{\mathrm{pred}}$ were constructed as parametrizations of specific physical degradation mechanisms. This type of model is very useful for quantifying the effects of the various non-idealities in experiments and comparing the behavior of isolated degradations to multi-dimensional simulations.

From an implosion design standpoint, the statistical model formulation in Eq. (2) is difficult to use for optimization purposes, since the simulated yield Y_{1D} responds to the same parameters as $\mathrm{YOC}_{\mathrm{pred}}$ in a complicated way and therefore it is not trivial to determine the inputs that optimize their product Y_{pred}. Finding the implosion design that produces the highest yield would thus require an expansive search over a vast parameter

space using a large number of simulations. Instead, a new formulation of the statistical model was developed to predict the experimental yield directly from the initial conditions, here referred to as the "design statistical model" or DSM. This new model is written in terms of design parameters, i.e., a reduced set of input parameters that can be used to model the outcome of the implosion with sufficient accuracy and are ideally independent and can be freely chosen by the designer. The target is characterized by the outer radius R_t, the thickness of the DT ice layer $\Delta_{\rm ice}$, the total thickness of the plastic ablator layers $\Delta_{\rm abl}$ and the material composition of the ablator. The laser inputs are parametrized by the size of the beam R_b, the total laser energy delivered E_L as well as a set of parameters used to characterize the shape of the laser pulse.

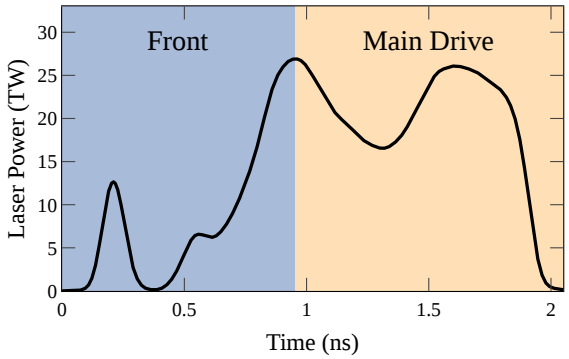

FIG. 1: Laser pulse shape from a high-performance OMEGA implosion. The front part of the pulse is used to launch two controlled shocks to set the shell entropy before adiabatically ramping up to peak power. The "main drive" accelerates the shell while balancing the trade-offs between coupling efficiency, in-flight stability and coasting.

A typical ICF laser pulse shape can be divided into two separate parts, each serving a specific purpose as shown in Fig. 1. The front part of the pulse is used to launch a series of controlled shocks to set the entropy of the shell before carefully ramping up to peak power. The majority of the energy is coupled to the target in the "main drive" part of the pulse to adiabatically accelerate the shell to a high implosion velocity. The challenge lies in defining an encoding consisting of a small number of design parameters that are able to capture the effect that the pulse shape features have on the implosion dynamics. For instance, a fixed amount of energy delivered to a target of a given mass can result in dramatically different fusion yields due to the shape of the main drive – a short, high intensity pulse couples most of the energy to the target early in the implosion when the target has not converged much resulting in much better coupling efficiency than a long, low intensity pulse. The higher implosion velocity V_i in the shorter pulse leads to much higher yields as[18] $Y_{\rm 1D} \sim V_i^6$. On the other hand, a pulse shape that is too short where the laser turns off long before the target reaches peak compression leads to coasting and decompression of the shell that causes poor convergence[19].

In the DSM, the front part of the pulse shape is parametrized by an in-flight measure of the shell entropy referred to as the shell adiabat $\alpha_F = P_{\rm sh}/P_F$, where $P_{\rm sh}$ is the shell pressure and P_F is the Fermi degenerate pressure. The intensity in the beginning of the main drive is characterized by the in-flight aspect ratio (IFAR). It is useful to combine the shell adiabat and IFAR into a single stability parameter[20]. $S = (\alpha_F/5)^{1.3}/({\rm IFAR}/47)$, which characterizes the growth of short wavelength perturbations[13] and can therefore be related to the yield degradation caused by such perturbations[15]. The same parameter, for reasons unrelated to stability, characterizes aspects of 1D implosion physics with lower values of S corresponding to either higher IFAR or lower adiabat that both lead to higher $Y_{\rm 1D}$. The laser intensity at the end of the pulse and the effect of the pulse shape on coasting are characterized by the 1D-simulated convergence ratio CR – ratio of initial target radius to hotspot radius at peak compression. While the convergence ratio and IFAR encode a substantial amount of the effect that the pulse shape has on the implosion, these parameters appear cumbersome to use as design parameters since they are not known ahead of time without running a simulation. In addition, both CR and IFAR are not independent of many of the other design parameters we identified earlier such as the shell dimensions and adiabat. As will be shown later, however, the highest fusion yields are achieved in a region of the parameter space where the yield is largely insensitive to simulated convergence and hydrodynamic stability and therefore these parameters do not need to be known with high precision during the design process.

The design parameters identified above account for a large fraction of the yield dependencies, however, a power-law model does not have enough capacity to fully represent the complicated dependency of $Y_{\rm 1D}$ on the pulse shape. Two multiplicative correction terms that are independent of any power-law of the design parameters are defined below to accommodate this "non-power-law" aspect of pulse shaping. The core idea of the DSM is that as long as these correction terms remain close to unity then the power-law dependencies will capture most of the effect of pulse shaping and therefore they can be relied upon by the designer when optimizing the implosions. On average, over all the shots in the OMEGA database, the correction terms defined below amount to an approximately 14% adjustment.

As mentioned earlier, because of the strong dependency of $Y_{\rm 1D}$ on the implosion velocity it is essential that the model is able to capture how pulse shaping influences the velocity. In order to achieve that, a power-law model was trained to predict the simulated implosion velocity V_i using the same set of parameters described above that will be used to predict the experimental yield. The velocity correction term $\hat{V}$ is then defined as the ratio of the simulated implosion velocity V_i and the predicted ve-

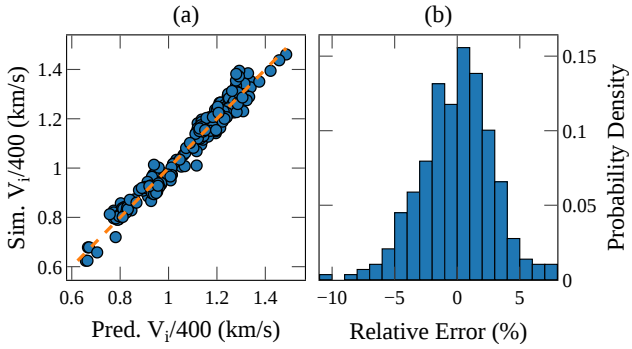

FIG. 2: (a) Piece-wise power-law fit of the simulated peak implosion velocity V_i using the same design parameters as the experimental yield model. The residual standard error of the fit is 2.9 %. (b) A histogram of the relative errors of the fit shows that only a small correction to the power-law is needed to reproduce V_i.

locity $V_i^{\rm pred}$ from the power-law model $\hat{V} = V_i/V_i^{\rm pred}$.

Note that this correction term is independent of any power-law of the parameters used in the regression. The power-law fit is shown in Fig. 2(a) and a histogram of the relative error in the velocity prediction for each shot is shown in Fig. 2(b). As seen in the figure, the power-law model is able to reproduce the simulated velocity with only a 2.9% error from the ground truth. Given such a small error, the dependencies inferred can be considered an excellent approximation of the true dependencies in the data.

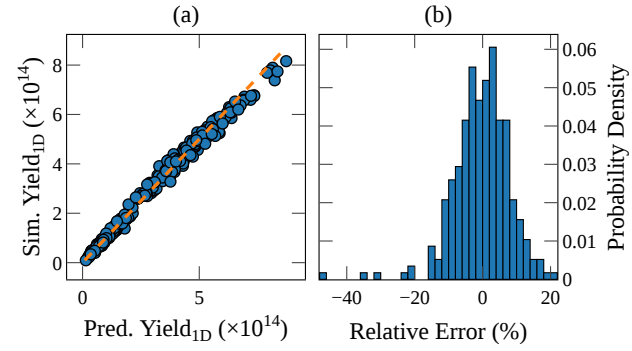

FIG. 3: (a) Piece-wise power-law fit of the simulated fusion yield $Y_{\rm 1D}$ using the same design parameters as the experimental yield model. The residual standard error of the fit is 8.2 %. (b) A histogram of the relative errors of the fit shows that for a majority of shots $Y_{\rm 1D}$ is well-reproduced by the power-law. A significant deviation in a particular shot can indicate a systematic shift in the design, e.g., a shock-merger inside the DT ice layer.

The same procedure was used to define an analogous correction term $\hat{Y} = Y_{\rm 1D}/Y_{\rm 1D}^{\rm pred}$ for the simulated yield. This term is again independent of any power-law of the regression parameters. The power-law fit of the 1D yield is shown in Fig. 3(a) where the velocity correction term $\hat{V}$ has now been included in the model. Fig. 3(b) shows a histogram of the relative error of the predicted 1D yield for each shot with an error of 9.7 %. Once again, since the magnitude of the yield correction term $\hat{Y}$ is relatively small, the power-laws inferred from the regression can be considered a good approximation of the true dependencies in the data. With the addition of the $\hat{V}$ and $\hat{Y}$ terms to the parametrization of the pulse shape it is ensured that the statistical model has the capacity to reproduce the simulated yield with good accuracy while the major dependencies are still extracted in the form of easy-to-interpret power-laws.

Modeling the fusion yield measured in experiments $Y_{\rm exp}$ requires two additional parameters that account for the physics not included in the 1D simulations as discussed in Ref.[16]. The effect of mode $\ell = 1$ asymmetry that varies from shot-to-shot is included post-shot through the apparent ion temperature asymmetry denoted as $\hat{R}_T$. The effect of ^{3}He accumulation in the central vapor region caused by the decay of tritium in the fuel over the fill age, i.e. the time between when the target is filled and shot-time, is characterized by the simulated yield degradation caused by the additional ^{3}He in the vapor. This fill age parameter is calculated by taking the ratio of yields in the ^{3}He-contaminated simulation and the "clean" simulation: $\mathrm{YOC_{He}} = Y_{\rm 1D,He}/Y_{\rm 1D}$ After including all of the aforementioned design parameters, the statistical prediction model for the fusion yield becomes:

$$Y_{\rm exp} \approx \mu_0 C_{\rm abl} E_L^{\mu_1} R_t^{\mu_2} A^{\mu_3} \hat{M}^{\mu_4} \left(\frac{R_b}{R_t}\right)^{\mu_5} \mathrm{CR}_{\rm 1D}^{\mu_6} S^{\mu_7} \alpha_F^{\mu_8} \hat{V}^{\mu_9} \hat{Y}^{\mu_{10}} \hat{R}_T^{\mu_{11}} \mathrm{YOC}_{\rm He}^{\mu_{12}}, \quad (3)$$

where $C_{\rm abl}$ is a discrete categorical parameter that represents the material composition of the ablator, E_L is the laser energy, R_t is the outer radius of the target, $A = R_t/(\Delta_{\rm ice} + \Delta_{\rm abl})$ is the initial aspect ratio, $\hat{M} = M_{\rm ice}/M_{\rm total}$ is the ratio of initial ice mass and total target mass, R_b is the radius of the laser beams, $CR_{\rm 1D}$ is the simulated convergence ratio, S is the stability parameter, α_F is the shell adiabat, $\hat{V}$ and $\hat{Y}$ are the aforementioned pulse-shape correction terms for simulated implosion velocity and fusion yield. The regression coefficients μ_0 through μ_{12} are determined from a fit to 289 experiments.

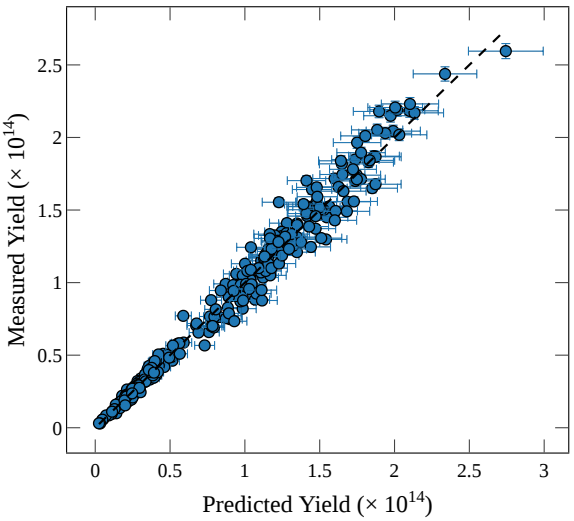

FIG. 4: Piece-wise power-law fit of the measured D+T neutron yield using the design parameters in Eq. (3). The power-law coefficients and threshold parameters are given in Table I. The residual standard error of the fit is 9.7 %.

Results from the DSM trained on the experimental yield are shown in Fig. 4, where all the parameters discussed above are included as shown in Eq. (3) with the coefficients listed in Table I. Also listed in the table are 1-σ confidence intervals for all the coefficients. Similar to Ref.[16], the regression is performed over piece-wise power laws, where some of the parameters such as the beam ratio Rb/Rt and hydro-stability term include physically motivated thresholds where the values of the coefficients change. The threshold locations are determined by a best fit to the experimental data and are listed in Table I. The quality of the fit in Fig. 4 with a residual standard error of 9.7% is approximately equal to the YOC-model in Ref.[16] whereas the main advantage of the approach in this paper is the ability to directly determine the optimum design parameters.

Writing the power-law model for the measured yield in a general form as $Y_{\rm exp} \approx \mu_0 \times \prod a_i^{\mu_i}$, where a are the regression parameters, the yield dependence (YD) on a single parameter a_j can be extracted from the model as

$$\mathrm{YD}(a_j) = \frac{Y_{\rm exp}}{\prod_{i\neq j} a_i^{\mu_i}} \qquad (4)$$

TABLE I: Piece-wise power-law coefficients for the statistical model of the measured fusion yield in Eq. (3). Where indicated, the exponents take on piece-wise constant values that change at the given thresholds determined from a best fit to data. The parameter $C_{\rm abl}$ is a constant whose value depends on the material composition of the ablator layer. The uncertainties listed indicate a 1-σ confidence interval.

Parameter	Coefficient
$C_{\rm abl}$	$C_{\rm abl} = 1.00$, if GDP CD
	$C_{\rm abl} = 1.05 \pm 0.03$, if Polystyrene
	$C_{\rm abl} = 1.45 \pm 0.03$, if SiCH
E_L	$\mu_1 = 2.64 \pm 0.08$
R_t	$\mu_2 = -3.60 \pm 0.28$
A	$\mu_3 = 4.80 \pm 0.15$
$\hat{M}$	$\mu_4 = 4.12 \pm 0.27$
R_b/R_t	$\mu_5 = 1.27 \pm 0.39$, if $R_b/R_t < 0.86$
	$\mu_5 = -2.69 \pm 0.29$, if $0.86 \leq R_b/R_t$
$\mathrm{CR_{1D}}$	$\mu_6 = -0.22 \pm 0.09$
S	$\mu_7 = 0.78 \pm 0.07$, if $S < 0.85$
	$\mu_7 = 0.01 \pm 0.05$, if $S \geq 0.85$
α_F	$\mu_8 = -0.17 \pm 0.11$
$\hat{V}$	$\mu_9 = 3.79 \pm 0.23$
$\hat{Y}$	$\mu_{10} = 1.02 \pm 0.10$
$\hat{R}_T$	$\mu_{11} = 0.00$, if $\hat{R}_T < 1.14$
	$\mu_{11} = -1.30 \pm 0.09$, if $\hat{R}_T \geq 1.14$
$\mathrm{YOC_{He}}$	$\mu_{12} = 1.35 \pm 0.05$

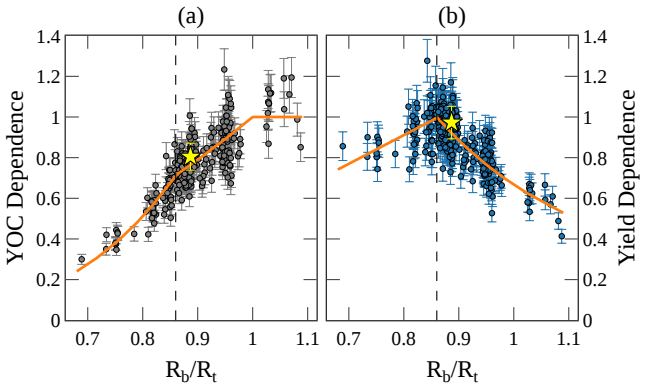

FIG. 5: Experimentally inferred dependence on R_b/R_t of (a) yield-over-clean from Ref.[16] and (b) fusion yield from Eq. (3). The optimum value of R_b/R_t that maximizes the fusion yield can be directly read from the plot of the yield model (b) whereas the underlying dependence of $Y_{\rm 1D}$ is not explicit in the YOC model (a). The current highest-performing shot on OMEGA (shot nr. 106493) is highlighted with a yellow star and is close to the optimum with $R_b/R_t = 0.89$.

Figure 5 shows the extracted dependence on the beam size term R_b/R_t for both (a) the YOC-model in Ref.[16]

and (b) the yield model defined here. The figure illustrates the advantage of the DSM that shows an optimum value of $R_b/R_t = 0.86$, where the predicted yield is maximized. In contrast, to find the optimum R_b/R_t using the YOC-model requires multiplication by Y_{1D}, the dependence of which on R_b/R_t is unknown in the YOC-formulation. Highlighted with the yellow star in Fig. 5 is OMEGA shot nr. 106493, which achieved the highest measured Lawson parameter to-date using the SG5-850 distributed phase plates[21] ($R_b = 415$ μm) on a target with a radius of $R_t = 465$ μm. The model in Fig. 5(b) suggests that the shot was close to the optimum R_b/R_t with a slight improvement in yield possible when increasing the target size to $R_t = 480$ μm.

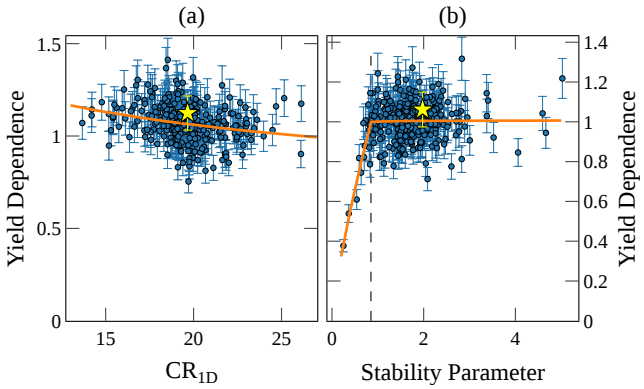

FIG. 6: Experimentally inferred dependencies of the fusion yield on (a) simulated convergence CR_{1D} and (b) hydrodynamic stability S. The dependence of Y_{1D} on the convergence is approximately canceled in the experimental data. The fusion yield shows negligible dependence on the stability parameter above the critical threshold of $S = 0.85$. The current highest-performing shot on OMEGA (shot nr. 106493) is highlighted with a yellow star and is located in the hydrodynamically stable region.

Figure 6 shows the inferred fusion yield dependence on (a) the simulated convergence CR_{1D} and (b) the hydrodynamic stability parameter S. While the yield in 1D simulations shows a clear dependence on both these parameters, the yield in experiments has a very weak dependence on CR_{1D} over the entire dataset and negligible dependence on S above the critical stability threshold $S_{crit} = 0.85$. This lack of dependence in the experimental data could be caused by 3D degradation effects that approximately cancel the improvements in the 1D dynamics but could also be caused by inaccuracies in 1D simulations, e.g., if the shell adiabat in the experiments is significantly higher than the simulations predict then that would result in a weaker apparent dependence on the simulated convergence.

Figure 7(a) shows the inferred fusion yield dependence on the initial aspect ratio A. Given that, at fixed R_t, a smaller initial aspect ratio corresponds to lower target mass and therefore a higher velocity it is not surprising that the model finds a very strong dependence with thin-

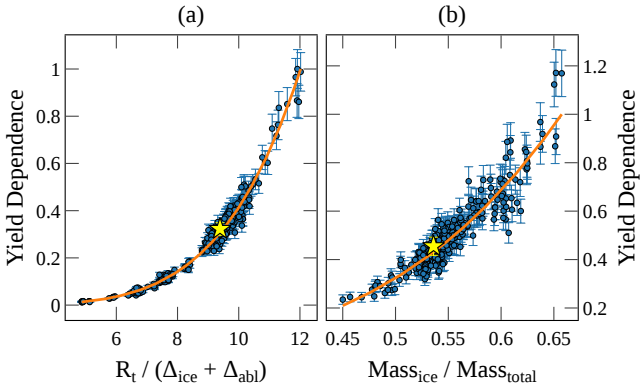

FIG. 7: Experimentally inferred dependencies of the fusion yield on (a) initial aspect ratio A and (b) ratio of ice mass to total target mass $\hat{M}$. The current highest-performing shot on OMEGA (shot nr. 106493) is highlighted with a yellow star, indicating that higher yield can be achieved by reducing the thickness of both the DT ice and ablator layers.

ner targets leading to higher yield. The yield dependence on the ratio of initial ice mass to the total target mass is shown in Fig. 7(b), indicating that higher yields are expected as the ablator mass is reduced with respect to the mass of the ice layer. It is important to note that there are limits to reducing the layer thicknesses that are not reflected in Fig. 7. As the ablator mass is reduced beyond a certain point, DT begins to reach the critical surface during the course of the implosion. This leads to a smaller fraction of the laser energy being absorbed in the plastic ablator with comparatively higher atomic number Z resulting in less driver energy coupled to the target. Eventually, as the ablator layer becomes very thin, DT will extend out to the quarter-critical surface leading to higher levels of two-plasmon decay and preheating of the target by hot electrons. In addition, as the target layers are made thinner it is important to keep the in-flight aspect ratio low enough to stay above the critical stability threshold S_{crit}.

Finally, Fig. 8 shows the experimental yield dependence on the ablator material, isolated from the effects of all the other design parameters via Eq. (4). The OMEGA database includes cryogenic targets with three main types of ablators: (1) glow-discharge polymer deuterated plastic (labeled GDP CD in Fig. 8), (2) polystyrene and (3) a two-layer ablator with an inner layer of GDP CD and an outer layer of Si-doped CH plastic. The effect of the ablator material is modeled in Eq. (3) as a constant multiplicative factor for each of the ablator types. Both the GDP CD and polystyrene ablators behave similarly in terms of the fusion yield while targets with an Si-doped outer CH layer show an approximately 45% increase in yield. This improvement in yield is attributed to the introduction of the higher-Z Si-doped material in the coronal plasma, leading to increased collisional absorption and reduction of cross-beam energy

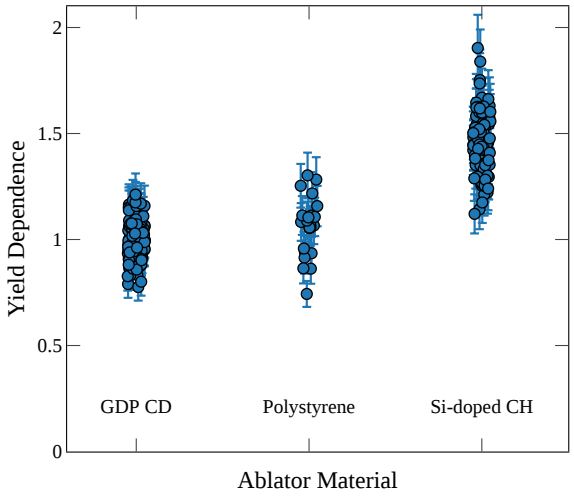

FIG. 8: Experimentally inferred dependence of the fusion yield on ablator material, indicating a 45% increase in yield, on average, when using Si-doped CD ablators over GDP CD and Polystyrene shells when all other parameters are fixed. A small jitter has been added to individual points in the horizontal axis to aid data visualization.

transfer[9]. In focused experiments, an outer Si-doped layer has been shown to reduce the growth of short wavelength perturbations[22] and mitigate hot-electron preheat from the two-plasmon decay instability[23,24]. Other dopants with higher Z numbers are being considered for future campaigns.

To conclude the discussion on optimizing yield, it follows from the DSM that the fusion yield on OMEGA can be maximized as follows: (1) using ablators with an Si-doped CH outer layer, (2) $R_t = 960$ μm with SG5-850 DPPs, (3) reducing the ablator and ice layer thicknesses to approximately $\Delta_{\rm ice} \approx 35$ μm and $\Delta_{\rm abl} \approx 7.0$ μm. However, it is important to keep in mind that simply reducing the target mass is expected to decrease the areal density. How these high-yield, thin-layer target designs can be used without sacrificing areal density and therefore improve the Lawson parameter is the topic of the next sections.

III. INFERRING AREAL DENSITY DEPENDENCIES

Building a statistical model of the type used in Sec. II to predict the areal density of an ICF implosion from initial conditions is a much more difficult task than predicting the yield for a number of reasons – higher measurement uncertainty, limited number of lines-of-sight to measure an anisotropic quantity, high sensitivity of areal density to minute details in shock timing. Instead, statistical modeling was used to relate the measured areal density to a combination of a subset of design parameters

TABLE II: Piece-wise power-law coefficients for the statistical model relating the areal density to target convergence. In the "Experiment" column the convergence is measured from time-integrated hotspot x-ray images. Synthetic x-ray images generated from 1D simulation data were used to calculate the simulated convergence in a consistent manner. The uncertainties listed indicate a 1-σ confidence interval.

Parameter	Experiment	1D Simulation
CR	$\mu_1 = 1.07 \pm 0.14$	$\mu_1 = 1.48 \pm 0.04$
$\Delta_{\rm ice}$	$\mu_2 = 1.03 \pm 0.14$	$\mu_2 = 1.02 \pm 0.04$
$\widehat{\rho R_0}$	$\mu_3 = 1.19 \pm 0.18$	$\mu_3 = 1.19 \pm 0.06$

and measured convergence. While this type of model is not capable of predicting the areal density from initial conditions since it includes the measured convergence as a parameter, it is nevertheless useful for two reasons: (1) the model provides a consistency check that indicates whether the observed variance in areal density is correlated to variance in the hotspot size or is caused by asymmetry and measurement errors; and (2) the model can be used to infer areal density dependencies on design parameters at fixed convergence.

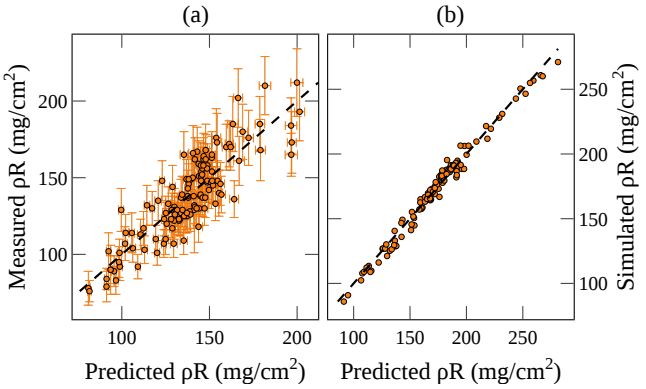

FIG. 9: Power-law fits of the (a) measured and (b) simulated areal density using the target parameters and convergence in Eq. (5). The power-law coefficients are given in Table II. The residual standard error is 9.7 % in the fit to experiments and 2.9 % in the fit to simulations.

Figure 9(a) shows the results from a statistical model fit of the areal density $\rho R_{\rm exp}$ measured with the magnetic recoil spectrometer[25] using the measured convergence ${\rm CR}_{\rm exp}$ from time-integrated hotspot x-ray images and two design parameters: the initial ice thickness $\Delta_{\rm ice}$ and the ratio of the initial areal density of the ablator to the total initial areal density $\widehat{\rho R}_0 = \rho R_{0,{\rm abl}}/\rho R_{0,{\rm total}}$

$$\rho R_{\rm exp} \approx \mu_0 {\rm CR}_{\rm exp}^{\mu_1} \Delta_{\rm ice}^{\mu_2} \left(\widehat{\rho R}_0\right)^{\mu_3} \qquad (5)$$

A corresponding statistical model relating the areal density from 1D simulations to simulated convergence is shown in Fig. 9(b), where the simulated convergence

is calculated based on the size of synthetic hotspot x-ray images in order to maintain consistency with the experimental model. The power-law coefficients for the three terms in Eq. (5) for both the experimental model and the simulation model are given in Table II. The experimentally inferred coefficients of both Δ_{ice} and $\widehat{\rho R}_0$ match the 1D simulations while the convergence term has a stronger effect in simulations. The difference in coefficients could be caused by uncertainties in the measurements or differences in density and temperature profiles between simulations and experiments leading to a different relationship between apparent hotspot size and areal density. Nevertheless, the clear dependence of the areal density on convergence indicates good consistency between the two measurements.

A comparison of the areal density dependencies to the results in Sec. II reveals the source of the difficulty in developing new designs to improve the Lawson parameter. According to the coefficients in Table II, the design changes required to reach higher yields – thinner DT ice layer and thinner ablator – would also lead to a reduction in areal density unless the convergence is increased. Achieving high convergence is one of the most difficult challenges in ICF. Recent results demonstrating improved convergence via subcooling the target below the triple point are discussed in the next section.

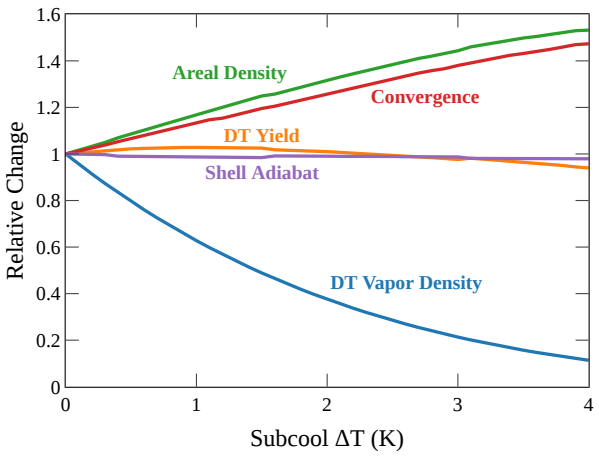

FIG. 10: Results from 1D simulations showing trends as the vapor density (blue line) is reduced by lowering the initial target temperature below the DT triple point. In-flight characteristics such as the shell adiabat are not affected by subcooling, while convergence and areal density increase with lower vapor density. Despite higher convergence, the yield in simulations stays approximately constant when subcooling by up to 4 K.

IV. SUBCOOLING

The convergence of an ICF implosion is dependent on the initial density of the DT vapor in the central region of the target, which can be modified via subcooling, i.e., reducing the temperature of the target below the DT triple point right before shot. This method of reaching higher convergence is advantageous since it does not affect the in-flight stability properties of the shell. However, subcooling can introduce other potentially detrimental side-effects. Namely, the material composition of the DT vapor changes in a way that has a negative effect on the yield and the stresses introduced in the DT ice can result in degraded uniformity of the layer.

A series of 1D LILAC simulations were run with varying degrees of subcooling to study the behavior of ideal implosions. The results of these simulations are shown in Fig. 10, where the subcooling ΔT on the horizontal axis corresponds to the difference between the DT triple point temperature and the temperature of the DT layer in the simulations. The initial vapor density at a given DT layer temperature, relative to the vapor density at the triple point is shown with the blue line in the figure. As expected, because subcooling only affects the deceleration phase of the implosion, the in-flight quantities such as the shell adiabat remain constant as the layer temperature is reduced while both the convergence and areal density monotonically increase as the initial vapor density decreases. Also shown in Fig. 10 is the simulated fusion yield, which remains approximately constant with subcooling. The lack of improvement in the simulated yield despite higher convergence is explained by the lower DT number density in the hotspot. In 1D simulations, a large fraction of the fusion yield is produced in the material that was initially in the vapor region (39% at 0 K subcooling in the simulations shown in Fig. 10) and the rest is produced in the DT that is ablated into the hotspot during deceleration. At 2.5 K of subcooling the fraction of the total yield that comes from the initial vapor region reduces to 18% because of the reduced vapor mass. In addition, because the heavier molecules that include tritium (TT and DT) will freeze before the DD molecules, the vapor becomes more deuterium-rich, thereby slightly reducing the reaction rate (from 52% deuterium at the triple point to 58% deuterium at 2.5 K of subcooling).

The effect of subcooling was studied in a number of implosion experiments on OMEGA over several shot days. The results of the experiments are shown in Fig. 11, with (a) the neutron yield, (b) convergence and (c) areal density plotted as a function of the subcooling temperature. The measurements are normalized to a non-subcooled reference shot with nominally the same pulse shape and target parameters that was taken on the same day to remove variations caused by detector calibrations across days. Contrary to simulations, the fusion yield in experiments does not stay constant but instead decreases with subcooling. Possible explanations for this discrepancy include the degradation caused by increased ice layer nonuniformity as visible in images of the targets before they were shot (see Fig. 11(d)) or a difference in the vapor composition compared to the simulations. Both con-

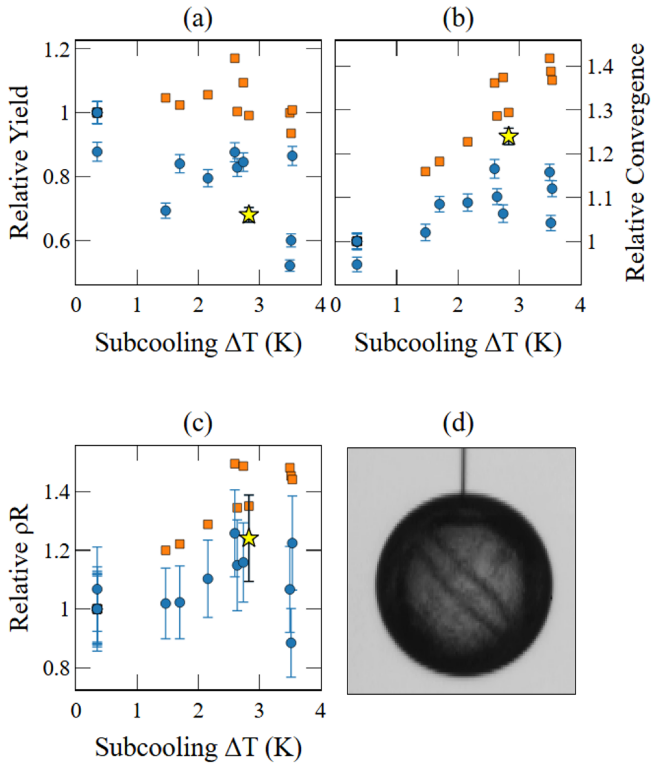

FIG. 11: Experimental results from subcooled implosions (blue points) compared to post-shot simulations (orange squares). Trends with respect the subcooling temperature are shown for (a) fusion yield, (b) target convergence and (c) areal density. The current highest-performing shot on OMEGA (shot nr. 106493) was subcooled by 2.8 K and is highlighted with a yellow star. All implosion data are normalized to a non-subcooled shot taken on the same day. (d) High speed video image taken shortly prior to the shot of the subcooled target from shot 106493 showing features in the DT ice layer caused by subcooling.

vergence and areal density show improvement with subcooling in the experiments but not as much as predicted by 1D simulations. However, not all subcooled shots exhibited the same increase in convergence. In particular, none of the shots at 3.5 K of subcooling showed improvement over < 3.0 K of subcooling. These 3.5 K subcooled shots all showed signatures of large mode-1 amplitude (hotspot flow velocity above 100 km/s, large apparent ion temperature asymmetry and asymmetric shape of x-ray images), while the best-performing shots at 2.5 K appeared round. It is clear that fielding targets at very low temperatures requires more accurate mitigation of low modes due to higher sensitivity at high convergence. A mode-1 mitigation strategy via an imposed target offset from the target chamber center was demonstrated in[26]. A novel analysis that takes into account measurements of laser beam mispointing and energy imbalance, as well as a systematic mode-1 caused by polarized cross-beam energy transfer[27] has been developed and used to refine the target offset procedure. This improved mode-1 mitigation strategy has been successfully employed in recent OMEGA cryo experiments and will be detailed in a forthcoming publication.

In terms of overall performance, subcooling leads to an increase in the Lawson parameter in shots with good low-mode symmetry. The current highest performing cryogenic implosion on OMEGA (shot 106493) is highlighted in Fig. 11 with a yellow star and was fielded at 2.8 K below the triple point. The performance improvement can be understood as follows: in a thin shell approximation the Lawson parameter can be re-written in terms of the hotspot size $R_{\rm hs}$ as $\chi \sim (Y/R_{\rm hs}^2\rho\Delta)^{1/3}(\rho\Delta)^{2/3} \sim [Y\rho\Delta({\rm CR})^2]^{1/3}$, where Δ is the shell thickness. If the reduction in yield caused by subcooling and the increase in areal density approximately cancel then the Lawson parameter for a subcooled shot increases as $\chi \sim {\rm CR}^{2/3}$. Shot 106493 achieved an approximately 10% increase in the Lawson parameter compared to the non-subcooled reference shot 106483.

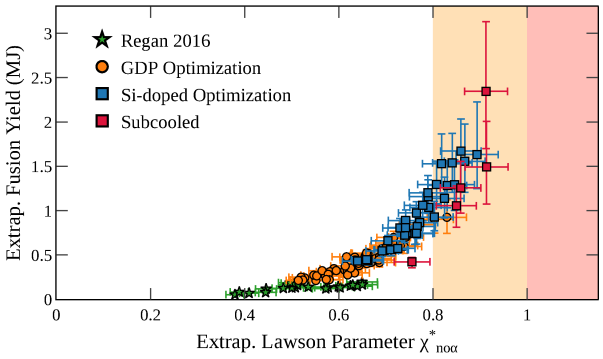

FIG. 12: Experimental results from cryogenic optimization campaigns on OMEGA since 2014. The generalized Lawson parameter was hydro-equivalently scaled to 2.15 MJ of laser energy. The extrapolated fusion yield is calculated by applying a yield amplification factor due to alpha-heating based on the scaled $\chi^*_{{\rm no}\alpha}$.

Figure 12 shows results from the best-performing cryogenic implosion experiments on OMEGA. The quantity plotted on the horizontal axis is the Lawson parameter $\chi^*_{{\rm no}\alpha}$ extrapolated to the NIF driver energy. The $\chi^*_{{\rm no}\alpha}$ is calculated by hydrodynamically scaling the hotspot from OMEGA experiments by a scale factor of 4.2. If laser-target coupling efficiency is assumed to be equivalent to OMEGA then the scale factor corresponds to a driver energy of 2.15 MJ. The figure shows progress made in the high-performance cryogenic implosion optimization campaigns since 2014, leading up to hydro-equivalent burning plasma ($\chi^*_{{\rm no}\alpha} > 0.8$) as reported in Ref.[9]. Figure 12 also shows that two of the subcooled shots with good low-mode symmetry reached higher Lawson parameters than the best non-subcooled shots. The highest performance was recorded on shot 106493 with $\chi^*_{{\rm no}\alpha} = 0.88 \pm 0.05$ It should be noted that shot 106493 is somewhat more

TABLE III: Implosion metrics for the current highest-performing OMEGA shot 106493 compared to the proposed new design. The values without uncertainties for shot 106493 are calculated from a post-shot 1D LILAC simulation.

	Shot 106493 (measured)	New Design (predicted)
V_i (km/s)	508	593
α_F	5.3	7.6
Stability	2.3	3.3
Neutron Yield	$(1.29 \pm 0.06) \times 10^{14}$	2.52×10^{14}
ρR (mg/cm^2)	165 ± 13	142
$M_{\rm stag}$ (μg)	7.1 ± 1.6	6.9
$\chi^*_{\rm noa}$ ($E_L = 2.15$ MJ)	0.88 ± 0.05	0.98

degraded by ^{3}He accumulation than most of the high-performance shots because of and 8-day DT fill age as opposed to the typical 3-day fill age. Correcting for the fill age effect using the statistical model in Eq. (3) suggests that if the shot were repeated with a 3-day old DT fill it would achieve a Lawson parameter of $\chi^*_{\rm noa} \approx 0.92$.

V. OPTIMIZED IMPLOSION DESIGN

In Ref.[14], a parallel Bayesian optimization algorithm was developed to search the vast parameter space of ICF implosion designs. The algorithm has an efficient inner loop that optimizes the laser pulse shape for a given target with a prescribed in-flight shell adiabat. In principle, the search can then be performed over an outer loop of all possible targets and shell adiabats. The size of the search-space can be significantly reduced by leveraging the findings from the DSM discussed in this paper. The optimum target dimensions that maximize the fusion yield on OMEGA were given in Sec. II. In addition, imposing the constraint that the design must lie above the critical stability threshold severely restrict the range of viable values for the shell adiabat.

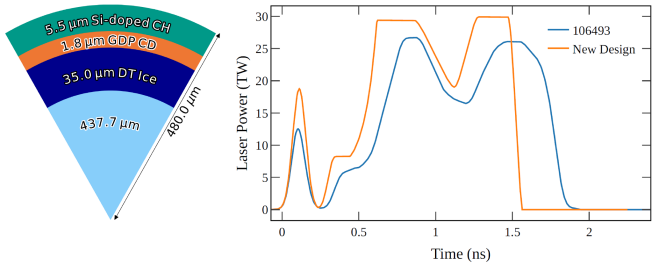

FIG. 13: (a) A schematic showing the optimized target dimensions from Sec. II and (b) the new laser pulse shape (blue) proposed to be used on the optimized target compared to the current highest-performing design (orange) from OMEGA shot 106493.

The Bayesian optimization algorithm was applied in this constrained region of parameter space to find the best pulse shape for the optimum target defined above[28]. In order to recover the areal density lost because of lower mass of the target, the optimization was performed with a DT ice layer subcooled below the triple point by 2.6 K. The resulting design is shown in Fig. 13 and compared to the current best performer (shot 106493). The performance predictions and important implosion metrics for both designs are given in Table III. The new design is predicted to reach a very high yield owing to significantly higher implosion velocity (Vi = 593 km/s) compared to shot 106493 (Vi = 508 km/s). The DSM in Sec. II was used to predict neutron yield while the areal density was predicted using the statistical model defined in Ref.[14]. The effect of subcooling was then modeled by applying a degradation of 27% to the yield prediction and a 20% increase to the areal density, according to the experimental data in Fig. 11. If the design in Fig. 13 achieves an areal density of The implosion is predicted to reach an extrapolated Lawson parameter of $\chi^*_{\rm noa} = 0.98$, when scaled to 2.15 MJ of laser energy, approaching hydro-equivalent ignition.

VI. CONCLUSIONS

Achieving hydro-equivalent ignition is a major milestone for the OMEGA optimization campaign. The highest-performing OMEGA implosions have reached a generalized Lawson parameter that is 88% of the value required for ignition when hydro-equivalently scaled to a laser energy of 2.15 MJ. In this paper, a new version of the statistical model for the fusion yield was developed by mapping the measured yield in terms of design parameters. This parametrization of the model allows direct optimization of the target parameters to maximize the fusion yield on OMEGA. It was found that the highest yield can be reached by using targets with Si-doped outer layers with an outer radius of $R_t = 480$ μm when using the standard SG5-850 DPPs and reducing the DT ice layer and ablator thicknesses to approximately 35 μm and 7.0 μm, respectively. A statistical mapping of the measured convergence to the areal density was used to show that in order to achieve higher performance with the proposed high-yield, thin-target designs it is necessary to increase the convergence. Experiments performed with subcooling the target below the DT triple point demonstrated higher convergence and a $\sim 10\%$ improvement in the generalized Lawson parameter. Finally, a new implosion design was developed by first strongly restricting the parameter space using the optimized target suggested by the new formulation of the statistical model and utilizing subcooling. A Bayesian optimization algorithm was then used to find the best laser pulse shape that maximizes the Lawson parameter. The new design is predicted to approach hydro-equivalent ignition when scaled to a laser energy of 2.15 MJ.

[1] John Nuckolls, Lowell Wood, Albert Thiessen, and George Zimmerman. Laser compression of matter to super-high densities:

Thermonuclear (ctr) applications. *Nature*, 239(5368):139–142, 1972.

[2]John D Lawson. Some criteria for a power producing thermonuclear reactor. *Proceedings of the physical society. Section B*, 70(1):6, 1957.

[3]R Betti, AR Christopherson, BK Spears, R Nora, A Bose, J Howard, KM Woo, MJ Edwards, and J Sanz. Alpha heating and burning plasmas in inertial confinement fusion. *Physical review letters*, 114(25):255003, 2015.

[4]RS Craxton, KS Anderson, TR Boehly, VN Goncharov, DR Harding, JP Knauer, RL McCrory, PW McKenty, DD Meyerhofer, JF Myatt, et al. Direct-drive inertial confinement fusion: A review. *Physics of Plasmas*, 22(11), 2015.

[5]John D Lindl. Inertial confinement fusion: the quest for ignition and energy gain using indirect drive. *(Springer-Verlag)*, 1998.

[6]EI Moses. Ignition on the national ignition facility. In *Journal of Physics: Conference Series*, volume 112, page 012003. IOP Publishing, 2008.

[7]AL Kritcher, DJ Schlossberg, CR Weber, CV Young, E Dewald, AB Zylstra, OA Hurricane, A Allen, B Bachmann, KL Baker, et al. Design of first experiment to achieve fusion target gain¿ 1. In *High-Power Laser Ablation VIII*, volume 12939, pages 25–31. SPIE, 2024.

[8]TR Boehly, DL Brown, RS Craxton, RL Keck, JP Knauer, JH Kelly, TJ Kessler, SA Kumpan, SJ Loucks, SA Letzring, et al. Initial performance results of the omega laser system. *Optics communications*, 133(1-6):495–506, 1997.

[9]V Gopalaswamy, CA Williams, R Betti, D Patel, JP Knauer, A Lees, D Cao, EM Campbell, P Farmakis, R Ejaz, et al. Demonstration of a hydrodynamically equivalent burning plasma in direct-drive inertial confinement fusion. *Nature Physics*, pages 1–7, 2024.

[10]V Gopalaswamy, R Betti, JP Knauer, N Luciani, D Patel, KM Woo, A Bose, IV Igumenshchev, EM Campbell, KS Anderson, et al. Tripled yield in direct-drive laser fusion through statistical modelling. *Nature*, 565(7741):581–586, 2019.

[11]CA Williams, R Betti, V Gopalaswamy, JP Knauer, CJ Forrest, A Lees, R Ejaz, PS Farmakis, D Cao, PB Radha, et al. Demonstration of hot-spot fuel gain exceeding unity in direct-drive inertial confinement fusion implosions. *Nature Physics*, pages 1–7, 2024.

[12]H Takabe, K Mima, L Montierth, and RL Morse. Self-consistent growth rate of the rayleigh–taylor instability in an ablatively accelerating plasma. *The Physics of fluids*, 28(12):3676–3682, 1985.

[13]Huasen Zhang, R Betti, Rui Yan, and Hussein Aluie. Nonlinear bubble competition of the multimode ablative rayleigh–taylor instability and applications to inertial confinement fusion. *Physics of Plasmas*, 27(12):122701, 2020.

[14]V Gopalaswamy, A Lees, R Ejaz, CA Thomas, TJB Collins, KS Anderson, W Ebmeyer, and R Betti. Automated and highly parallelized bayesian optimization scheme for direct drive fusion experiments on omega. *Physical Review Research*, 7(1):013009, 2025.

[15]A Lees, R Betti, JP Knauer, V Gopalaswamy, D Patel, KM Woo, KS Anderson, EM Campbell, D Cao, J Carroll-Nellenback, et al. Experimentally inferred fusion yield dependencies of omega inertial confinement fusion implosions. *Physical review letters*, 127(10):105001, 2021.

[16]Arnee Lees, Riccardo Betti, James P Knauer, Varchas Gopalaswamy, Dhrumir Patel, Ka Ming Woo, Ken S Anderson, E Michael Campbell, Duc Cao, Jonathan Carroll-Nellenback, et al. Understanding the fusion yield dependencies in omega dt-layered implosion experiments using a physics-based statistical mapping model. *Physics of Plasmas*, 30(1), 2023.

[17]J Delettrez, R Epstein, MC Richardson, PA Jaanimagi, and BL Henke. Effect of laser illumination nonuniformity on the analysis of time-resolved x-ray measurements in uv spherical transport experiments. *Physical Review A*, 36(8):3926, 1987.

[18]Ricardo Betti, PY Chang, BK Spears, KS Anderson, J Edwards, M Fatenejad, JD Lindl, RL McCrory, R Nora, and D Shvarts. Thermonuclear ignition in inertial confinement fusion and comparison with magnetic confinement. *Physics of Plasmas*, 17(5), 2010.

[19]OA Hurricane, A Kritcher, DA Callahan, O Landen, PK Patel, PT Springer, DT Casey, EL Dewald, TR Dittrich, T Döppner, et al. On the importance of minimizing "coast-time" in x-ray driven inertially confined fusion implosions. *Physics of Plasmas*, 24(9), 2017.

[20]VN Goncharov, TC Sangster, R Betti, TR Boehly, MJ Bonino, TJB Collins, RS Craxton, JA Delettrez, DH Edgell, R Epstein, et al. Improving the hot-spot pressure and demonstrating ignition hydrodynamic equivalence in cryogenic deuterium–tritium implosions on omega. *Physics of Plasmas*, 21(5):056315, 2014.

[21]Y Lin, TJ Kessler, and GN Lawrence. Design of continuous surface-relief phase plates by surface-based simulated annealing to achieve control of focal-plane irradiance. *Optics Letters*, 21(20):1703–1705, 1996.

[22]G Fiksel, SX Hu, VA Goncharov, DD Meyerhofer, TC Sangster, VA Smalyuk, B Yaakobi, MJ Bonino, and R Jungquist. Experimental reduction of laser imprinting and rayleigh–taylor growth in spherically compressed, medium-z-doped plastic targets. *Physics of Plasmas*, 19(6), 2012.

[23]RK Follett, JA Delettrez, DH Edgell, VN Goncharov, RJ Henchen, J Katz, DT Michel, JF Myatt, J Shaw, AA Solodov, et al. Two-plasmon decay mitigation in direct-drive inertial-confinement-fusion experiments using multilayer targets. *Physical Review Letters*, 116(15):155002, 2016.

[24]K Churnetski, D Patel, W Theobald, R Betti, D Cao, L Ceurvorst, MJ Rosenberg, AA Solodov, C Stoeckl, and SP Regan. Mitigation of hot-electron preheat from the two-plasmon-decay instability using silicon-doped plastic shells in direct-drive implosions on omega. *Physics of Plasmas*, 31(11), 2024.

[25]M Gatu Johnson, J Katz, C Forrest, JA Frenje, V Yu Glebov, CK Li, R Paguio, CE Parker, C Robillard, TC Sangster, et al. Measurement of apparent ion temperature using the magnetic recoil spectrometer at the omega laser facility. *Review of Scientific Instruments*, 89(10), 2018.

[26]OM Mannion, IV Igumenshchev, KS Anderson, R Betti, EM Campbell, D Cao, CJ Forrest, M Gatu Johnson, V Yu Glebov, VN Goncharov, et al. Mitigation of mode-one asymmetry in laser-direct-drive inertial confinement fusion implosions. *Physics of Plasmas*, 28(4):042701, 2021.

[27]DH Edgell, PB Radha, J Katz, A Shvydky, D Turnbull, and DH Froula. Nonuniform absorption and scattered light in direct-drive implosions driven by polarization smoothing. *Physical Review Letters*, 127(7):075001, 2021.

[28]Varchas Gopalaswamy, Riccardo Betti, Aarne Lees, Cliff Thomas, Timothy Collins, and Kenneth Anderson. Physics informed, automated and highly parallel bayesian optimization of direct-drive implosions. *Bulletin of the American Physical Society*, 2024.