

Simulating high-harmonic generation with ultracold atoms

Daniel M. Harrington*

*Department of Physics, University of California, Santa Barbara, California 93106, USA and
Department of Physics & Astronomy, Tufts University, Medford, Massachusetts 02155, USA*

Anna Dardia, Yifei Bai, Siddharth Mukherjee, Toshihiko Shimasaki, and David Weld
Department of Physics, University of California, Santa Barbara, California 93106, USA

Analog quantum simulation using trapped ultracold atoms can be used to investigate natural processes at more accessible time and energy scales with highly tunable parameters. Ultrafast electronic dynamics of atoms in strong laser fields is one such area that can be challenging to investigate experimentally. In particular, recent theoretical proposals have shown how high-harmonic generation spectra can be accessed with cold atom simulators to address outstanding questions and theoretical inconsistencies. Here we describe an experimental apparatus for simulating strong field interactions using trapped ultracold atoms and structured light fields that is robust against noise and aberration. We detail the use and optimization of a digital micromirror device to shape a linear intensity gradient and demonstrate how an iterative algorithm can be used to improve light pattern fidelity and achieve target accuracy in the chemical potential applied to the atoms.

I. INTRODUCTION

A deep understanding of ultrafast electronic dynamics is increasingly important in a wide range of disciplines, where these dynamics provide descriptions of phenomena seen at angstrom and attosecond levels. The motion of electrons in these realms is responsible for emergent behavior in nearly every composite system, the importance of which is emphasized by the 2023 Nobel prize in physics for advances in techniques to probe behavior at these scales. However, the inherent difficulty of accessing configurable, interesting experimental regimes has motivated alternate approaches to the study of ultrafast dynamics.

In recent years, quantum simulation has been demonstrated as a useful technique by which to investigate various physical systems of scientific interest [1]. In particular, neutral atom analog quantum simulators have been used to study ultrafast processes [2–4], where relative time scales are generally increased by many orders of magnitude. These simulators are comprised of a system whose behavior is described by a Hamiltonian directly mappable to that of the process being simulated and are able to reconstruct a wide variety of dynamics and measured observables.

One such process of interest is high-harmonic generation (HHG) [5–7] wherein an atom illuminated with a strong laser light pulse emits harmonics of the incident light frequency upon recombination to the ground state. Accessing HHG experimentally in a tunable way with high time and energy resolution is non-trivial, and analog quantum simulation promises an easier platform with which to explore this phenomenon [8]. Such a platform

allows for new capabilities in studying HHG in settings including solid-state media and non-classical light sources where traditional theoretical models have shown inconsistent results [8, 9]. Natural HHG experiments typically utilize a macroscopic cloud of a noble gas, making precise study of HHG in a single atom difficult. Being able to simulate HHG in a single atom directly is a valuable tool to access challenging experimental regimes and to benchmark developing models with tunability not often possible in direct experiments.

By simulating the nuclear potential and electric field-dipole interaction over a few-cycle light pulse using structured light incident on a trapped Bose–Einstein condensate (BEC), we can reconstruct the electronic dynamics that lead to HHG. The harmonic emission spectrum can then be accessed through position or momentum measurements of the BEC that map to dipole moment oscillations responsible for emission. Here we describe the physics of HHG, how it can be simulated, and a portion of the experimental setup needed to generate simulated forces using a digital micromirror device (DMD) on a ^{84}Sr quantum simulator.

II. THEORY

A. Strong field dynamics

We first introduce a qualitative treatment of the interaction of a bound electron with a strong electric field [5, 8, 10]. When high-intensity coherent radiation is incident on an atom, the resultant interactions with the electronic wavefunction give rise to a rich variety of highly non-classical phenomena. In the strong-field regime, the laser electric field interacts with the electron charge in a way that cannot be treated perturbatively. If the elec-

* dharr24@tufts.edu

tron position is denoted by \mathbf{r} , then its potential energy is given (in atomic units) under the dipole approximation as

$$V(\mathbf{r}, t) = V_{\text{atom}}(\mathbf{r}) + \mathbf{r} \cdot \mathbf{E}(t) \quad (1)$$

where $\mathbf{E}(t)$ is the laser electric field and $V_{\text{atom}}(\mathbf{r})$ refers to the effective nuclear Coulombic potential. Assuming a pulse waveform linearly polarized in the x direction, the field can be written $E(t) = E_0 f(t) \sin(\omega t + \phi)$, where $f(t)$ is a some pulse envelope function.

When the field is strong (i.e., E_0 is large), the initial interaction with the electron effectively lowers the nuclear Coulombic potential, allowing the electron to tunnel or escape from its parent atom. Under the strong field approximation, the escaped electron behaves as a free particle with zero initial velocity. As the system evolves in time and the pulse cycle changes sign, the electric field accelerates the electron back toward the nucleus and they interact. During this process, a portion of the electron wavefunction interferes with itself, resulting in an oscillation of the dipole moment. This leads to harmonic emissions at higher frequencies that depend on the pulse form, intensity, and nuclear potential. This process of electron excitation, excursion, and emission is called high-harmonic generation.

B. Optical dipole force

Trapped ultracold atoms used in the simulator must be configured in such a way that their collective wavefunction mimics another physical system. For simulating the combined potential of the nuclear interaction and the dipole interaction in eq. (1), the technique that we employ is based on the optical dipole force arising from the a.c. Stark effect. Here we briefly review how this force arises and its use in quantum simulation experiments.

As noted previously, the interaction of an electric field with the induced dipole moment of the atom results in a potential energy change [11]. The dipole moment arising from the separation of charges in the atom is itself proportional to the electric field and can be expressed as $\epsilon_0 \chi_a \mathbf{E}$, where $\epsilon_0 \chi_a$ is the scalar polarizability. The optical dipole potential can then be written as

$$V = -\frac{1}{2} \epsilon_0 \chi_a E^2 \quad (2)$$

A perturbation theory treatment expresses eq. (2) in terms of atomic properties [11] as

$$V(\mathbf{r}, t) = \frac{\hbar \Gamma^2}{8\delta} \frac{I(\mathbf{r}, t)}{I_{\text{sat}}} \quad (3)$$

The optical dipole force is the conservative force arising from the gradient of this potential and is proportional to the gradient of intensity:

$$F_{\text{ODF}}(\mathbf{r}, t) = \nabla V(\mathbf{r}, t) = \frac{\hbar \Gamma^2}{8\delta I_{\text{sat}}} \nabla I(\mathbf{r}, t) \quad (4)$$

By deterministically configuring the spatio-temporal intensity distribution of light incident on trapped atoms, an arbitrary potential energy landscape can be configured to match that of another physical system. Values of physical observables in the system of interest can then be extracted through measurements on the atoms used in the simulator.

C. Simulating HHG

To simulate the HHG process, the BEC wavefunction must experience similar forces to a bound electron experiencing a strong laser pulse. The Coulombic potential is approximated by a Gaussian and can be simulated using a simple Gaussian beam incident on the BEC. For linear polarization, the potential from the dipole interaction with the laser pulse takes the form $E(t)x$. At some time t , this force can be reconstructed in the simulator using a linear intensity gradient. The intensity is modulated in time to match the pulse waveform modulation. Arbitrary modulation capabilities allow for tuning the duration, frequency, amplitude, and cycle number of this pulse, parameters which cannot be tuned arbitrarily with real lasers in direct HHG experiments. In this way, the BEC potential landscape can be dynamically sculpted to mimic the electron potential in a strong field environment.

The HHG spectrum can then be accessed through measurements of the dipole moment. The emission yield is proportional to the dipole acceleration, which can be reconstructed from the dipole velocity. The equivalent observable to the dipole moment velocity is the momentum distribution of the BEC. This can be retrieved through a time of flight measurement, wherein the BEC trap is turned off and the atoms are allowed to expand for a finite time before being imaged and the velocity calculated.

III. OPTICS AND HARDWARE

A. Digital micromirror devices

In order to create a beam with a linear intensity transverse cross section, we employ a DMD to selectively transmit light according to a prescribed pattern [12, 13]. The DMD is composed of an array of addressable microscopic mirrors that articulate about their diagonal, directing light either to further collection optics or to be discarded. Mirrors are by default in a “parked state” where they are oriented at a flat 0° . The “on” and “off” states correspond to a respective positive and negative 12° tilt from the normal.

The DMD is non-trivial to treat theoretically, with the traditional blazed grating assumption failing to give accurate predictions [13–15]. Contrary to the prediction yielded by a simple blazed grating model, ref. [13] found maximum first order diffraction efficiency was achieved when the incident beam angle was 12° , the same as the DMD mirror tilt angle. In our setup, the input beam is close to this angle, and we achieve efficiencies within 5% of the theoretical maximum described in [13].

The full optical setup for testing and development is described in appendix A. More information on DMD control is available on the Weld lab internal wiki [16].

B. Power considerations

The DMD arrangement must be such that the linear gradient opposes the direction of gravity and imparts a force greater than that of gravity on the atoms. An acousto-optic modulator (AOM) is then used to modulate the intensity in time to follow the amplitude of a pulse waveform. The pulse zero crossing is at a force that corresponds to an acceleration of $1g$. The force imparted on the atoms by the incident light must then be strictly greater than the force of gravity.

We can calculate the force generated by an incident beam following eq. (2). We express the intensity as

$$I(\mathbf{r}, t) = T(t) \frac{P_0}{A} \frac{y}{y_0} \quad (5)$$

where P_0 is the incident power, A is the cross-sectional image area, $T(t) = E(t)/E_0$ is the time variation of the pulse waveform, and y_0 is the vertical distance over which the gradient extends. From eqs. (2) and (4), the force is

$$F(y, t) = \frac{\hbar\Gamma^2}{8\delta I_{\text{sat}}} \frac{P_0 T(t)}{A y_0} \quad (6)$$

Since the force is proportional to the beam power P_0 , the power considerations of all optical components is important to ensure that the force is high enough to get above the $1g$ waveform zero crossing.

Source	Measured eff.
AOM	0.8
Spatial filter	0.95
DMD main order	0.349
Gradient	0.30
Total Eff.	0.080

TABLE I: Measured efficiencies for test setup.

With an initial 10W output from the laser, this restricts the atom incident power to $P_0 = 0.8\text{W}$. Assuming an area of $A = 50 \times 100\mu\text{m}^2$, this results in a force of 3.48×10^{-24} N or an acceleration of $2.54g$ experienced by ^{84}Sr atoms from eq. (6).

C. DMD imaging

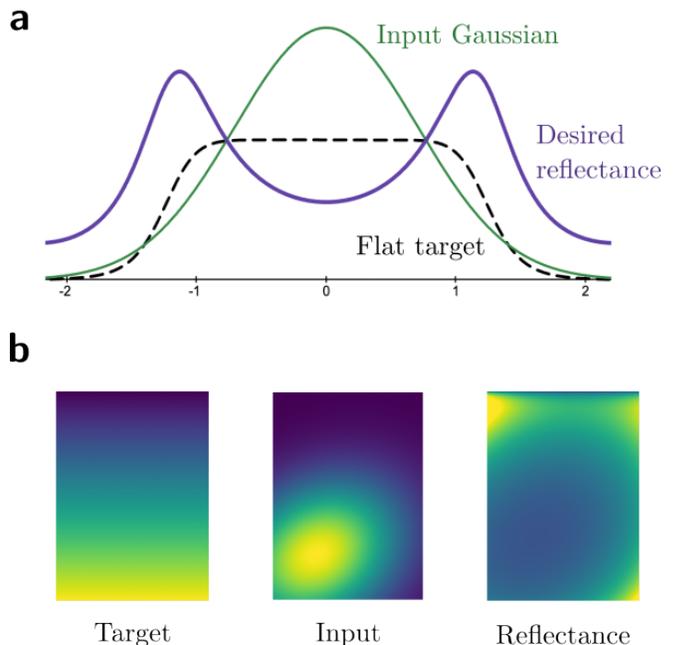


FIG. 1: (a) Beam cross section diagram showing correction for beam curvature required to create a flat top profile. (b) Target intensity, fitted input beam intensity, and computed desired reflectance profiles for linear gradient.

To project a linear intensity gradient using the DMD, the continuous target image (single-channel 8-bit integer valued, in our setting) is dithered into a binary array. This is accomplished through the Floyd-Steinberg error diffusion algorithm, which is natively implemented in Python’s PIL library. This algorithm has been shown to perform more favorably than random dithering [17] but can lead to large scale periodic structures in the dithered image that are not present in the original continuous one. Other groups have seen improvements using a modified Floyd-Steinberg algorithm which adds random noise in the error diffusion step [18]. We do not observe any macroscopic periodicity in the linear gradient, so currently this solution is not implemented.

The DMD illumination beam is a Gaussian mode with characteristic curvature over the scale of the DMD window. This results in non-uniformity that appears in the reflected DMD image. To account for this, the displayed DMD image $D_0(x, y)$ is modified to divide out the beam profile [19]. The beam profile is computed by imaging the DMD with a solid field displayed and computing a 2D Gaussian fit. The fit is used instead of the raw image to reduce the impact of noise, fringing, and normalization issues from bad pixels. Since the intensity is the measured quantity, the DMD image accounts for the amplitude reflection through a square root. The final expression for the initial continuous DMD image (before dithering) is [19]

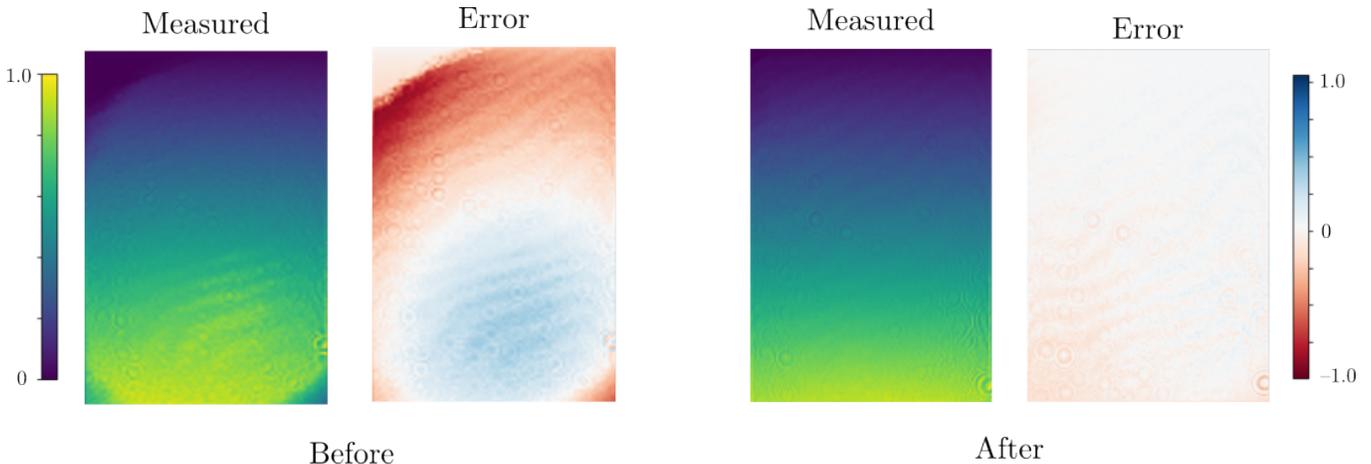


FIG. 2: Feed forward optimization. (Left) Initial measured intensity and error map for linear gradient. (Right) Measured intensity and error map after seven feed forward iterations with step size $\eta = 0.2$.

$$D_0(x, y) = \left(\frac{\text{target}(x, y)}{\text{beam}(x, y)} \right)^{1/2} \quad (7)$$

where $\text{beam}(x, y)$ is the 2D Gaussian fit of the beam and $\text{target}(x, y)$ is the desired intensity distribution (here a linear gradient). This expression leads to a much more accurate measured gradient than simply displaying a gradient image and provides adequate initial conditions for further optimization.

The DMD controller chip updates the mirror states every $105\mu\text{s}$. The mirrors briefly unlatch from their current state before being flipped to their next state. However, this occurs even when the mirror state does not change on the next frame, leading to undesired motion of mirrors even when a static pattern is displayed. This results in high frequency flickering of the DMD output image visible on a photodiode in the 10kHz range which can lead to undesired heating effects when projected onto the atoms [20]. This issue is addressed by interrupting the mirror clock pulse (MCP) sent from the DMD controller chip to the DMD itself. Further detail of our implementation is given in Appendix B.

With this fix and the DMD pattern computed from eq. (7), the linear gradient is improved but still not at a high enough fidelity. To enhance the quality of the projected gradient, we use software solutions to optimize for best quality.

IV. SOFTWARE OPTIMIZATION

A. Initial processing

Before characterization and optimization of the DMD images can be performed, they must first be processed. This consists of transformation, fringe removal, and normalization steps.

In the image captured on the beam profiler, the DMD output is shown as a scaled, rotated, and noisy version of the image displayed on the DMD. For comparison to the target image, the captured image must be transformed back to the same resolution rectangle as the DMD display. The rotation, translation, and scaling information is encoded in an affine transformation represented by a 2×3 matrix M that contains both the linear map and the translation.

$$M = [A \ B] = \begin{bmatrix} a_{00} & a_{01} & b_0 \\ a_{10} & a_{11} & b_1 \end{bmatrix} \quad (8)$$

Encoding a two dimensional input point (x, y) as an augmented vector $\mathbf{v} = [x \ y \ 1]^\top$, the corresponding transformed point $\mathbf{w} = [x' \ y' \ 1]^\top$ is given by the product $\mathbf{w} = M\mathbf{v}$.

The matrix M can be defined by three pairs of points in the domain and the codomain. For our application, the three input points are manually selected as three of the corners of the DMD image as seen on the beam profiler, and the second set of points are the corresponding corners of the DMD image itself. Other works have shown an automated method of computing the coefficients of M through gradient descent by maximizing the overlap between the DMD displayed image and the image captured on the beam profiler [21]. We have found good stability with the manual method and are not concerned with optimizing high resolution structures in the projected pattern, so this procedural mapping method was deemed unnecessary.

Once the mapping between pixels in the measured and target images is found and the transformation is performed, the two images can be compared. However, aberrations and fringing in the imaging path prevent unbiased quantitative comparisons. Several methods for fringe removal were attempted, and one was found to be successful for our application.

The most straightforward approach was a spatial low

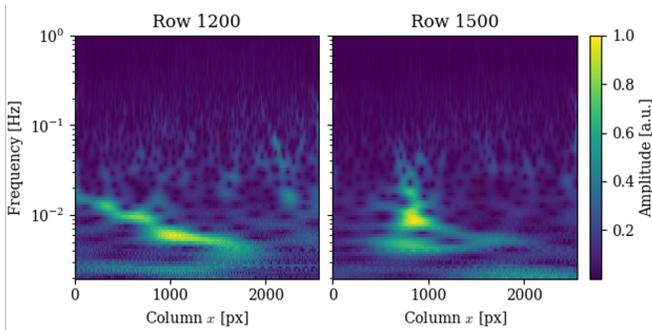


FIG. 3: Continuous Morlet wavelet transforms for two rows showing row variation in fringe signal.

pass filter, which only became successful at removing fringes at cutoff frequencies high enough that they were contributing to the formation of the gradient. Filtering out fringes would also result in unwanted distortions at the edges of the image. A second approach was attempted that was based on the continuous wavelet transform and introduced in [22]. A wavelet is a function $\psi : \mathbb{C} \rightarrow \mathbb{C}$ which defines a basis for $L^2(\mathbb{C})$ composed of transformations of ψ [23]. Coefficients of some arbitrary function in this basis can provide information about its frequency components that vary over the function domain.

This is conceptually similar to the Fourier basis, and indeed wavelets admit a decomposition similar to the Fourier transform known as the continuous wavelet transform, defined for a sequence $\{x_n\}$ (which corresponds to the image) as [23]

$$W_n(s) = \sum_{n'=1}^N x_{n'} \psi^* \left(\frac{(n' - n)\delta t}{s} \right) \quad (9)$$

where s is a wavelet scale and here n indexes along the image columns. Practically, a convenient choice of wavelet such as the Morlet wavelet (a Gaussian-windowed complex plane wave) can yield a transformation that provides spatially-resolved frequency information. Ref. [22] introduced the application of wavelet transforms to remove fringes. The wavelet transform displays high intensity regions which correspond to fringe spatial frequencies that can be subtracted out. An inverse wavelet transform then recovers the image without these fringes present. Wavelet transforms for two rows of a gradient image are shown in fig. 3. The high amplitude frequencies where fringes are present visibly differ from row to row. The observed fringe spectra have a large degree of variation across row and column, making automated fitting and identification of the fringe components for removal as in [22] not straightforward.

The last and most successful method for fringe and noise removal is to use a computed flat field. A solid field image is captured that can then be used to normalize other images and remove persistent fringes. Time-independent intensity variations of the DMD input beam

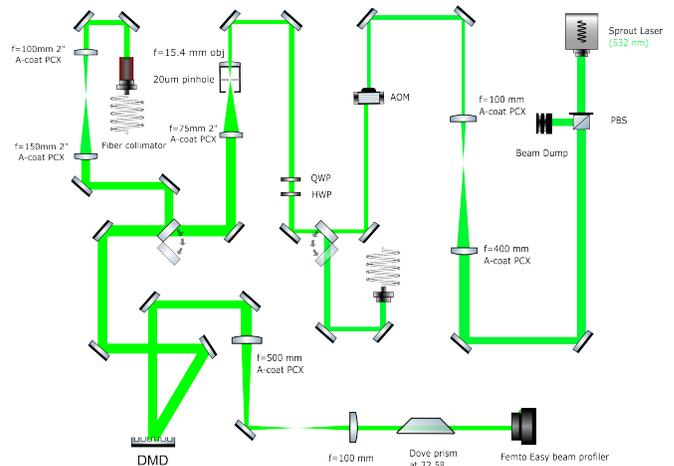


FIG. 4: DMD test setup optical diagram.

and collection optics, however, want to be preserved so that they can be corrected for in later optimization steps. Speckle patterns resulting from random interference of coherent light also pose an issue but are harder to detect in our imaging setup than scattering from larger dust particles. Reduction in undesirable speckle and interference can be attained by use of an incoherent illumination source [24], but with our coherent source setup further noise reduction methods were found to be necessary. A naïve flat field collection approach of simply displaying a solid field on the DMD or replacing it with a mirror will reflect the intensity profile of the input beam.

One solution to combat this issue is to use a computed flat field image. Computing this flat field is first done by displaying a solid field on the DMD and capturing the image. This image is then smoothed using a 10×10 convolution kernel before fitting to a 2D Gaussian. The original image is then divided by the Gaussian fit to produce an image that is approximately flat in intensity while still carrying information about fringes and static noise in the imaging optics. Before any optimization is performed, the measured images are divided by this computed flat field to remove fringes associated with the imaging path. It is worth noting that in the final installation on the full experiment, only fringes from the imaging path, not any arising from the DMD or subsequent optics on the path to the atoms, want to be removed before analysis.

B. Feed forward

After the initial setting of the DMD pattern following eq. (7) and fringe removal, there still exist deviations from the desired gradient in the measured field. Mitigation of these adverse effects is accomplished through an iterative feed forward optimization procedure [21]. The output from the DMD is imaged onto a detector, and the difference between the captured and target images is used to compute an update map that is then applied to

the DMD. This process repeats until convergence.

The normalized error for each iteration k is computed as

$$E_k(i, j) = \frac{n_M(i, j) - n_T(i, j)}{\max\{n_M\}} \quad (10)$$

where n_M and n_T are the measured and target (continuous) intensity distributions, respectively. This error is used to compute an updated pixel map through a simple first order method:

$$D_{k+1}(i, j) = D_k(i, j) - \eta E_k(i, j) \quad (11)$$

for some step size η . This gives a continuous desired reflectance array which is then dithered and displayed on the DMD. The process is repeated until the current iteration performs worse than the previous iteration. Performance is evaluated by using a normalized root mean square (RMS) error [17]:

$$\varepsilon_k = \sqrt{\sum_{i,j} \left(\frac{n_M(i, j)}{n_T(i, j)} - 1 \right)^2} \quad (12)$$

where the sum is over pixels with values greater than a threshold value to prevent unwanted fluctuations.

V. DISCUSSION AND OUTLOOK

The methods presented in this work are able to accurately produce a linear intensity gradient with high fidelity. A combination of fine tuning alignment, characterization of the input beam, and iterative feed forward optimization are found to be necessary in order to produce a high quality image. This involves a combination of techniques reported in earlier works [19–21, 25] for best results. The process and control code used easily generalize to distributions other than linear intensity gradients and successfully enable a highly reconfigurable platform for designer chemical potential landscapes.

Python code, schematics, and documentation relating to this project is available on the Weld lab Github [26].

ACKNOWLEDGMENTS

I would like to thank my faculty advisor Prof. David Weld for giving me the opportunity to work with his lab and contribute to a meaningful and interesting project, as well as for valuable direction and mentorship. Thank you also to graduate students Anna Dardia and Yifei Bai for their invaluable guidance and support, to fellow undergraduate Siddharth Mukherjee for sharing his knowledge of DMDs and prior contributions to this project, and to Dr. Toshihiko Shimasaki for additional guidance and help. I would also like to thank REU site director Prof. Sathya Guruswamy for making this experience

possible. This work was supported by NSF REU grant PHY-2349677.

Appendix A: Optics Testing Layout

The optics diagram for the current testing layout is shown in fig. 4. The output mode from the laser (Sprout 10W 532nm) is resized to fit the AOM aperture, where the first order is collected. A flip mirror path exists directing the beam into a fiber coupler that is sent to a fiber collimator output port. This path was used for testing and comparison but cannot be used at high power. If not sent into the fiber, two waveplates are used for polarization control before the beam is sent through a spatial filter and resized.

For the second lens in this resizing telescope, we tried focal lengths of both 75mm and 100mm. The 100mm focal length lens gives a larger image that covers the entire DMD area and extends partially outside the edges. The beam from the 75mm lens is completely contained within the DMD array but does not illuminate the whole region as well as the 100mm lens path. The first order DMD efficiency was measured at 44.9% using the 75mm lens and 34.9% using the 100mm lens, where the loss in power comes from excess light hitting the region around the DMD. Due to the better illumination uniformity, we use the 100mm lens for the corresponding better image size. After reflection from the DMD, the beam is resized for imaging.

We do not observe any change in DMD power transmittance at different polarizations. The DMD is rotated at 45° so that the mirror’s axis of rotation is perpendicular to the table surface. The projected image is then also rotated, so a dovetail prism is used to rotate the image 45° in the opposite direction.

Because the incident beam is coming in from the left side of the DMD, the default on and off states are switched.

Appendix B: MCP Interrupter

Our DMD is a DLP9000 contained in the LC9000 EVM package driven by the DLPC900 controller chip. The MCP output from the DLPC900 is labeled DLP_STRB/AF5 in the documentation and schematics and is accessible at test point TP15 on the LC9000 EVM board. To control the transmission of the MCP, we removed resistor R66 before TP15, leaving the circuit open. Two leads were then soldered to either side of the old resistor contacts and connected to a MOSFET (2N7000) switching circuit. When the MCP should be interrupted, the MOSFET is pulled high and the pulse is sent to ground.

We found that using the same 10Ω resistance value as in the documentation led to undesired small signals still being transmitted. This is likely due to a combination of

the on resistance of the MOSFET and contact resistance through the solder connections and breadboard testing setup. The small transmission still led to switching of the mirrors, so a 150Ω resistor between the MCP input and output was found to be necessary to suppress this issue. The previous 10Ω resistor is part of an RC filter before transmission through the ribbon cable connecting the EVM to the DMD chip, and we have not observed any negative effects with the added resistance included.

In addition to updating mirror states, the MCP is designed to prevent mirrors from getting stuck in their previous states. When the MCP is interrupted with our circuit for a long enough time, there is a risk of the mirrors getting stuck. This occurred once during testing when the MCP was only interrupted for a few seconds. Various other testing has involved interrupting the MCP for tens of seconds with no issues. When mirror stick-

ing occurs, leaving the MCP uninterrupted for several minutes has resolved the issue in the previous recorded case. To prevent further risk of mirror sticking, we use a microcontroller-based interlock which routes the MCP if it has been off for longer than a specified amount of time.

The interlock and control circuit consists of a TTL input signal to be generated from Cicero when a stable DMD pattern is needed. The microcontroller (Arduino Nano) monitors this input and pulls the MOSFET high sending the MCP to ground whenever the TTL signal is high. If the interrupting TTL signal has been on for longer than a specified amount of safe time, the MOSFET is pulled low and the MCP is allowed to pass for a specified time before returning to the input state again. A full schematic and PCB layout is available on Github [26].

-
- [1] I. M. Georgescu, S. Ashhab, and F. Nori, Quantum simulation, *Reviews of Modern Physics* **86**, 153 (2014).
- [2] R. Senaratne, S. V. Rajagopal, T. Shimasaki, P. E. Dotti, K. M. Fujiwara, K. Singh, Z. A. Geiger, and D. M. Weld, Quantum simulation of ultrafast dynamics using trapped ultracold atoms, *Nature Communications* **9**, 2065 (2018).
- [3] S. Sala, J. Förster, and A. Saenz, Ultracold-atom quantum simulator for attosecond science, *Physical Review A* **95**, 011403 (2017).
- [4] S. Arlinghaus and M. Holthaus, Driven optical lattices as strong-field simulators, *Physical Review A* **81**, 063612 (2010).
- [5] M. Lewenstein, P. Balcou, M. Y. Ivanov, A. L’huillier, and P. B. Corkum, Theory of high-harmonic generation by low-frequency laser fields, *Physical Review A* **49**, 2117 (1994).
- [6] J. L. Krause, K. J. Schafer, and K. C. Kulander, High-order harmonic generation from atoms and ions in the high intensity regime, *Physical Review Letters* **68**, 3535 (1992).
- [7] S. Ghimire and D. A. Reis, High-harmonic generation from solids, *Nature Physics* **15**, 10 (2019).
- [8] J. Argüello-Luengo, J. Rivera-Dean, P. Stammer, A. S. Maxwell, D. M. Weld, M. F. Ciappina, and M. Lewenstein, Analog Simulation of High-Harmonic Generation in Atoms, *PRX Quantum* **5**, 010328 (2024).
- [9] M. Even Tzur, M. Birk, A. Gorlach, M. Krüger, I. Kammer, and O. Cohen, Photon-statistics force in ultrafast electron dynamics, *Nature Photonics* **17**, 501 (2023).
- [10] F. Krausz and M. Ivanov, Attosecond physics, *Reviews of Modern Physics* **81**, 163 (2009).
- [11] C. Foot, *Atomic Physics* (Oxford University Press, 2005).
- [12] G. Gauthier, I. Lenton, N. M. Parry, M. Baker, M. J. Davis, H. Rubinsztein-Dunlop, and T. W. Neely, Direct imaging of a digital-micromirror device for configurable microscopic optical potentials, *Optica* **3**, 1136 (2016).
- [13] S. Scholes, R. Kara, J. Pinnell, V. Rodríguez-Fajardo, and A. Forbes, Structured light with digital micromirror devices: A guide to best practice, *Optical Engineering* **59**, 041202 (2019).
- [14] Y.-X. Ren, R.-D. Lu, and L. Gong, Tailoring light with a digital micromirror device, *Annalen der Physik* **527**, 447 (2015).
- [15] S. M. Popoff, Y. Bromberg, M. W. Matthès, and R. Gutiérrez-Cuevas, *A practical guide to Digital Micromirror Devices (DMDs) for wavefront shaping* (2024), [arXiv:2311.17496 \[physics\]](https://arxiv.org/abs/2311.17496).
- [16] Weld lab wiki dmd page (permission req.), <https://github.com/weldlabucsb/weldwiki/wiki/DMD-Control-and-Operation>.
- [17] C. Dorrer and J. D. Zuegel, Design and analysis of binary beam shapers using error diffusion, *JOSA B* **24**, 1268 (2007).
- [18] A. Mazurenko, *Probing Long Range Antiferromagnetism and Dynamics in the Fermi-Hubbard Model*, Ph.D. thesis (2017).
- [19] J. Liang, J. Rudolph N. Kohn, M. F. Becker, and D. J. Heinzen, 1.5% root-mean-square flat-intensity laser beam formed using a binary-amplitude spatial light modulator, *Applied Optics* **48**, 1955 (2009).
- [20] K. Hueck, A. Mazurenko, N. Luick, T. Lompe, and H. Moritz, Note: Suppression of kHz-frequency switching noise in digital micro-mirror devices, *Review of Scientific Instruments* **88**, 016103 (2017).
- [21] G. Gauthier, *Transport and Turbulence in Quasi-Uniform and Versatile Bose-Einstein Condensates*, Ph.D. thesis, The University of Queensland, Brisbane (2020).
- [22] P. M. Rojo and J. Harrington, A method to remove fringes from images using wavelets, *The Astrophysical Journal* **649**, 553 (2006).
- [23] C. Torrence and G. P. Compo, A Practical Guide to Wavelet Analysis, *Bulletin of the American Meteorological Society* (1998).
- [24] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice, *Nature* **462**, 74 (2009).
- [25] J. Liang, J. Rudolph N. Kohn, M. F. Becker, and D. J. Heinzen, High-precision laser beam shaping using a binary-amplitude spatial light modulator, *Applied Op-*

tics **49**, 1323 (2010).

[26] Dmd control code, <https://github.com/weldlabucsb/>

[dmd-feedforward](#).