

Model-based real-time synthesis of acousto-optically generated laser-beam patterns and tweezer arrays

Marcel Mittenbühler¹, Lukas Sturm¹, Malte Schlosser¹, and Gerhard Birkel^{1,2,*}

¹*Technische Universität Darmstadt, Institut für Angewandte Physik, Schlossgartenstraße 7, 64289 Darmstadt, Germany*

²*Helmholtz Forschungsakademie Hessen für FAIR (HFHF), GSI Helmholtzzentrum für Schwerionenforschung, 64291 Darmstadt, Germany*

 (Received 21 April 2025; revised 22 August 2025; accepted 10 November 2025; published 17 December 2025)

Acousto-optic deflectors (AODs) enable spatiotemporal control of laser beams through diffraction at an ultrasonic grating that is controllable by radio-frequency (rf) waveforms. These devices are a widely used tool for high-bandwidth random-access scanning applications, such as optical tweezers in quantum technology. A single AOD can generate multiple optical tweezers by multitone rf input in one dimension. Two-dimensional (2D) patterns can be realized with two perpendicularly oriented AODs. As the acousto-optical response depends nonlinearly on the applied frequency components, phases, and amplitudes, and in addition experiences dimensional coupling in 2D setups, intensity regulation becomes a unique challenge. Guided by coupled-wave theory and experimental observations, we derive a compute-efficient model which we implement on a graphics processing unit. Only one-time sampling of single-tone laser-power calibration is needed for model parameter determination, allowing for straightforward integration into optical instruments. We implement and experimentally validate an open-loop diffraction efficiency control system that enables programmable 2D multibeam trajectories with intensity control applied at every time step during digital signal generation, overcoming the limited flexibility, pattern-size constraints, and bandwidth limitations of methods using precalculation and precalibration of a predefined pattern set or closed-loop feedback. The system is capable of stable real-time waveform streaming of arrays with up to 50×50 tweezers with minimal time resolution of 1.4 ns (700 MS/s) and a peak latency below 257 μ s for execution of newly requested patterns. Reactive, real-time 2D multibeam laser patterning and scanning with strict intensity matching will substantially benefit parallelization and increasing data rates in materials processing, microscopy, and applications of optical tweezers, where the work presented is of immediate relevance for time-critical conditional operations with minimal latency.

DOI: [10.1103/d3tx-3tg8](https://doi.org/10.1103/d3tx-3tg8)

I. INTRODUCTION

Advanced methods for spatiotemporal control of laser beams are driving innovation in a wide range of applications from quantum technology and advanced manufacturing to biology [1–19]. Among the physical principles utilized for dynamic light modulation, the electro-optic and acousto-optic effects facilitate solid-state devices with no moving parts. Compared to inertia-limited spatial light modulators and mirror-based scanners,

these optical modulators excel in high-bandwidth applications [20]. On the basis of the acousto-optic effect (i.e., the diffraction of a laser beam by an ultrasound acoustic wave), optical parameters of the diffracted light can be accessed through the applied waveform of the acoustic wave in the radio-frequency (rf) domain. This allows for precise control of phase, frequency, amplitude, and direction of propagation of the diffracted light. Optimized designs accommodate the basic constraints of phase matching (i.e., momentum conservation) for various applications [6,21,22], combining intensity modulation with beneficial characteristics for rf modulation, tunable wavelength filtering, and beam deflection.

In the work presented in this paper, we focus on two-dimensional (2D) intensity and position control, employing two acousto-optical deflectors (AODs) in a mutually perpendicular configuration. Our methods can be adapted

*Contact author: apqpub@physik.tu-darmstadt.de, <https://www.iap.tu-darmstadt.de/apq>

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International* license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

to other acousto-optical designs and applications. Unlike electro-optic beam scanners, AODs enable parallelization by applying multifrequency waveforms (i.e., multi-tone patterns of acoustic waves) that generate multiple diffracted beams with a single device. For each beam, the angle of deflection is linearly linked to the specific rf frequency, giving access to straightforward position control. The ratio of light intensities of diffracted and incident beams, in general, is nonlinearly dependent on the rf power spectrum, relative phases of the individual tones, and geometrical effects. In consequence, intensity regulation presents a considerable challenge, particularly for multitone operation. Practical realizations have resorted to calibration of previously known patterns, which carries at minimum a quadratic overhead with the number of tones, or slow iterative closed-loop feedback, significantly reducing the usable bandwidth.

We have developed and implemented a method for the spatiotemporal control of single- and multibeam patterns that harnesses the full acousto-optical bandwidth. The control system presented mitigates nonlinearities of the acousto-optical response in real time for strict intensity matching in user-defined beam patterns and vector trajectories of laser spots in a 2D plane (Fig. 1). For this purpose, we derived a model of the frequency-dependent transfer function that maps the nonuniform diffraction efficiency, the interference of frequency components in multi-tone patterns, and the implications of using two AODs in series. Model parameters are independent of the number of rf tones and can be determined via single-tone calibration measurements. We implement a data-efficient tuning procedure based on model assumptions, that is applied at every time step during digital signal generation using a central processing unit (CPU) as pattern controller, a graphics processing unit (GPU) for waveform computation, and an arbitrary waveform generator (AWG) for signal output [Fig. 1(a)]. Our open-loop approach is valid for arbitrary configurations and highly scalable.

Figures 1(b) and 1(c) depict exemplary realizations of 2D multibeam control. In Fig. 1(b) a uniform 50×50 array of regularly spaced focused Gaussian beams is shown, while the image in Fig. 1(c) was created by rastering 184 lines of individual one-dimensional (1D) horizontal multi-beam patterns for which we implemented an open-loop version of the composite beam method for beam shaping [23]. The system presented is particularly suited for improving time-critical random-access scanning applications [20,24–26]. This includes increasing data rates in materials processing by direct laser writing for electronic microfabrication [1] and nanoscale additive manufacturing by multiphoton lithography [2,3], as well as in advanced imaging techniques of laser-scanning microscopy [4–6] and superresolution imaging [7]. Another prominent example is resolved laser addressing for the precise manipulation of individual atomic, molecular, and biological




FIG. 1. Real-time synthesis of AOD-generated 2D beam patterns and laser spot arrays. (a) Schematic illustration of digital signal synthesis and optical setup (PM, laser power meter; PD, photodiode). The second dimension of the beam pattern is omitted for clarity. (b) Camera image ($3.1 \text{ mm} \times 3.1 \text{ mm}$) of an intensity-equalized array of 50×50 optical tweezers in the focal plane. For a targeted integral diffraction efficiency of 50% an efficiency of $(49.5 \pm 2.3)\%$ is achieved. The frequency span is 36.75 MHz per dimension with a 0.75 MHz spacing between tones. (c) Two-dimensional spatiotemporal control was used to imprint the logo of TU Darmstadt depicting the Greek goddess Athene.

systems with so-called optical tweezers [8,9], enabling significant methods in cytology [10,11], sensing [12–14], chemistry [15,16], and quantum technology [17–19].

In our research work, we target quantum science with atomic and molecular systems in optical tweezer arrays [27–29]. We use this application to describe and assess our method. In this field [17], AOD-generated movable optical tweezers have become the *de facto* standard for realizing transport operations [27,30–38] and laser addressing of quantum systems in a site-selective fashion [32,35,39–43]. These techniques are essential for the formation of individual molecules [15] and the assembly of long-range interacting quantum systems [44] of Rydberg atoms [45] or polar molecules [46], continuous operation of the created quantum arrays [33,34,47,48], and the ability to perform algorithms and implement error correction in neutral-atom quantum processors [32,49–51]. Future progress will significantly benefit from parallelization by flexible multitone operation of AODs with reliable diffraction efficiency. This is exemplified by the rearrangement of quantum systems for the formation of defect-free arrays, where the number of transport operations shows polynomial scaling with increasing array size. The polynomial degree is above

linear for sequential single-tweezer operations [52]. Optimized multitweezer algorithms potentially achieve below square-root scaling, but practical advantage has been limited by inhomogeneity effects when using multitone waveforms [53].

A second parameter of severe practical importance, in particular for conditional mid-circuit operations [32, 54], is the latency between request and execution of a transport or addressing operation. Although a hard-wired implementation on a field programmable gate array (FPGA) potentially can achieve superior performance for a fixed algorithm [55,56], the accessibility and flexibility of software-based control, scalable GPU compute power, and straightforward AWG-based output have proven to be essential for the presented work. In contrast to previous approaches on a similar platform [57,58], which inevitably add output delay at the millisecond scale by processing sequences of entire waveforms, our method uses a streaming mode of operation that requires the waveform to be processed continuously in small blocks only. Depending on the size of the block, short latencies of less than 257 μ s are realized. To fulfill these strict real-time requirements on signal generation, we put a highly optimized fixed-core software system into operation. Moreover, the modular programming approach enables the extension to additional applications in areas other than qubit transport by loading additional modules to serve the requirements of specific tasks.

The paper is organized as follows. Section II summarizes the theory of acousto-optics for the generation of dynamic 2D arrays. Section III describes the model for intensity adjustment via 2D decomposition and single-axis diffraction efficiency control, as well as details on waveform streaming performance and the real-time signal generation framework. In Sec. IV the test apparatus is explained and measurements of the single- and multitone efficiency control are presented.

II. THEORY OF ACOUSTO-OPTICS FOR DYNAMIC TWEEZER ARRAYS

The interaction between an incident light field and the acoustic wave within the crystal of an AOD can be understood as a scattering process. The acoustic wave changes the refractive index, creating a periodic grating that diffracts the incoming light wave [21,22]. Assuming plane waves and a fixed optical wavelength, modeling the diffraction efficiency in the first diffraction order needs mainly to consider two dependencies. First, maximum diffraction occurs as a result of wave-vector matching within the crystal and therefore is a function of the acoustic waveform [22,59,60]. Second, coupled-wave theory provides a foundation for the dependence of the diffraction efficiency $\eta(\nu)$ on the power $P_a(\nu)$ of the acoustic wave with frequency ν and the properties of the AOD crystal

parameterized in the acousto-optical figure of merit η_c . Under the above assumptions, η_c is constant for a particular device. The single-tone diffraction efficiency $\eta(\nu)$ is given by Ref. [22] using the normalized sinc function

$$\eta(\nu) = \eta_c P_a(\nu) \text{sinc}^2 \left(\frac{1}{2\pi} \sqrt{4\eta_c P_a(\nu) + (\Delta k(\nu)L)^2} \right) \quad (1)$$

with crystal length L and wave vector mismatch $\Delta k(\nu) = |\vec{k} + \vec{q}(\nu) - \vec{k}_d(\nu)|$ of the incident light wave $\vec{k}$, the acoustic wave $\vec{q}(\nu)$, and the diffracted light wave $\vec{k}_d(\nu)$.

Typically, AODs are operated in a power regime up to the first maximum of $\eta(\nu)$. In negligence of the oscillatory behavior, this maximum is commonly termed saturation. We approximate $\eta(\nu)$ by introducing the experimentally accessible quantities of peak diffraction efficiency $\alpha(\nu)$, rf power $P(\nu)$, and rf saturation power $\beta(\nu)$ as

$$\eta(\nu) \simeq \alpha(\nu) \sin^2 \left(\frac{\pi}{2} \sqrt{\frac{P(\nu)}{\beta(\nu)}} \right). \quad (2)$$

The error incurred by this approximation up to saturation, where $P(\nu) = \beta(\nu)$ yields $\eta(\nu) \simeq \alpha(\nu)$, is below 0.03%. Since the fraction $P(\nu)/\beta(\nu)$ describing the power scaling is dimensionless, the expression can be rescaled to any experimentally accessible power measure, for example, the output power of the signal generator or the digital power level set to drive it.

For waveform synthesis, it is central to determine the value of $P(\nu)$ required to achieve a targeted diffraction efficiency. The approximate model can be inverted efficiently. Moreover, using coupled-wave theory, a simple analytical solution for multitone acoustic waveforms (i.e., operation with multifrequency rf signals) can be derived. Let $\eta(\nu) \in [0, 1]^N$ be a list of N diffraction efficiencies at respective frequencies $\nu \in \mathbb{R}^N$ and rf powers $P(\nu) \in \mathbb{R}^N$. Assuming no coupling between diffracted components and operation close to the Bragg regime ($\Delta k \approx 0$), we find

$$\eta_i := \eta(\nu_i) = \frac{\alpha(\nu_i) P(\nu_i)}{\beta(\nu_i) \sum_{l=1}^N \frac{P(\nu_l)}{\beta(\nu_l)}} \sin^2 \left(\frac{\pi}{2} \sqrt{\sum_{l=1}^N \frac{P(\nu_l)}{\beta(\nu_l)}} \right). \quad (3)$$

Taking into account normalization of all $P(\nu_i)$ to $\beta(\nu_i)$, the global diffraction efficiency depends only on the total rf power. In addition, the global diffraction efficiency is divided among all tones ν_i by their normalized fraction of the total rf power.

Next, we consider two AODs in a sequential, mutually perpendicular configuration for 2D deflection, as depicted in Fig. 1(a). The angle of the incident beam on the second AOD is altered by the deflection of the first, depending

on the frequency components of the applied rf signal. Thus, the 2D diffraction efficiency must account for the deflection in both AODs, which gives rise to a reduction in the peak diffraction efficiency of the 2D system due to the induced wave vector mismatch. We model the 2D diffraction efficiency $\eta_{ij}^{(xy)}$ by single-axis efficiencies $\eta_i^{(x)}$, $\eta_j^{(y)}$ and an interaction term Ξ_{ij} that expresses the frequency-dependent efficiency loss:

$$\eta_{ij}^{(xy)} = \eta_i^{(x)} \eta_j^{(y)} \Xi_{ij}. \quad (4)$$

For deflection with N tones $\nu_i^{(x)}$ in x and M tones $\nu_j^{(y)}$ in y , $\Xi \in [0, 1]^{N \times M}$. It is considered an error term which leads to a deviation from two ideal single-axis systems. To minimize the error between the measurement and our model, we analyzed the dominant terms in a Taylor series, finding empirically that the residuals in $\eta_{ij}^{(xy)}$ can be minimized assuming

$$\Xi_{ij} = \left(1 - \xi_0 \left(\nu_i^{(x)} - \xi_x \right)^2 \left(\nu_j^{(y)} - \xi_y \right)^2 \right)^2, \quad (5)$$

with interaction parameters ξ_x , ξ_y , and ξ_0 . The parameters ξ_x and ξ_y are frequencies related to the geometry of the optical setup, while ξ_0 quantifies the interaction strength. For different applications and AOD devices, a revised form of Ξ may be chosen for minimizing the residuals.

III. REAL-TIME MODELING AND SIGNAL GENERATION

In a typical application, a target diffraction efficiency pattern $\eta_T^{(xy)} \in [0, 1]^{N \times M}$ should be generated with frequencies $\nu^{(x)} \in \mathbb{R}^N$, $\nu^{(y)} \in \mathbb{R}^M$ defining the spatial structure of the pattern. The goal of the method presented is the computation of the required rf power values for all frequency components that minimize the error between the target efficiency pattern and the measured diffraction efficiencies. As a simplification, we assume that $\eta_T^{(xy)}$ can be decomposed into an outer product defined by the single-axis target diffraction efficiencies $\eta_T^{(x)} \in [0, 1]^N$, $\eta_T^{(y)} \in [0, 1]^M$. Due to the interaction term, it is impossible to achieve the exact pattern in a 2D array if $N > 1 \wedge M > 1$, the interaction can be exactly compensated for a single beam and a line of beams only. However, it is possible to compute sets of adjusted diffraction efficiencies $\eta_i^{(x)}$ and $\eta_j^{(y)}$ to minimize the error caused by Ξ . Consequently, our approach splits the model-based diffraction control into two parts. First, the adjusted $\eta_i^{(x)}$ and $\eta_j^{(y)}$ are calculated to minimize the error in $\eta_{ij}^{(xy)}$ compared to the targeted 2D diffraction efficiency $\eta_T^{(xy)}$. Second, the 1D diffraction efficiency model is used to compute the rf power values needed for achieving the adjusted diffraction efficiencies for each axis.

A. Two-dimensional decomposition

The actual diffraction efficiency given in Eq. (4) contains the interaction term that restricts an exact decomposition when controlling multiple tones in both dimensions. The goal is to minimize the error in the achieved diffraction efficiency caused by Ξ_{ij} . Therefore, adjusted values for the 1D diffraction efficiency $\eta_i^{(x)}$, $\eta_j^{(y)}$ need to be found, as a minimization problem with

$$\arg \min_{\eta_i^{(x)}, \eta_j^{(y)}} \sum_{i=1}^N \sum_{j=1}^M \left(\eta_{T,i}^{(x)} \eta_{T,j}^{(y)} - \eta_i^{(x)} \eta_j^{(y)} \Xi_{ij} \right)^2. \quad (6)$$

An approximate solution given by a 2D decomposition can be computed by

$$\eta_i^{(x)} = \frac{\eta_{T,i}^{(x)}}{\sqrt{c_0 c_x} \sum_{j=1}^M \eta_{T,j}^{(y)} \Xi_{ij}}, \quad (7)$$

$$\eta_j^{(y)} = \frac{\eta_{T,j}^{(y)}}{\sqrt{c_0 c_y} \sum_{i=1}^N \eta_{T,i}^{(x)} \Xi_{ij}}, \quad (8)$$

$$c_0 = \left[\sum_{i=1}^N \eta_{T,i}^{(x)} \sum_{j=1}^M \eta_{T,j}^{(y)} \Xi_{ij} \right]^{-1}. \quad (9)$$

The solution is not uniquely specified as it can be altered without consequences by a constant ratio of the efficiencies c_x/c_y with $c_x c_y = 1$. This degree of freedom mirrors the experimental implementation since diffraction efficiencies are measured after both AODs only. Finding the exact ratio can be challenging as this would require a secondary optimization goal to introduce a constraint, for example, measuring the diffraction efficiency for both AODs independently. It is a strength of our approach that these additional measurements are not required as the ratio can be chosen to distribute the mean deviations from the target diffraction efficiency approximately equally between both axes, by setting

$$c_x = \frac{1}{c_y} = \sqrt{\frac{\sum_{i=1}^N \sum_{j=1}^M \eta_{T,j}^{(y)} \Xi_{ij}}{\sum_{i=1}^N \sum_{j=1}^M \eta_{T,i}^{(x)} \Xi_{ij}}}. \quad (10)$$

In general, computation of Eqs. (7) and (8) would require a summation over the target diffraction efficiency and the interaction term Ξ for each tone, resulting in quadratic scaling in the number of frequency components. Moreover, Eqs. (9) and (10) would require a double summation, resulting in quadratic scaling of the computation time concerning the number of frequency components. The form of Ξ in Eq. (5) on the other hand, allows a decomposition

such that a summation can be rewritten as

$$\sum_{j=1}^M \eta_{T,j}^{(y)} \Xi_{i,j} = a_y^{(0)} + \left(v_i^{(x)} - \xi_x\right)^2 a_y^{(1)} + \left(v_i^{(x)} - \xi_x\right)^4 a_y^{(2)}, \quad (11)$$

$$a_y^{(0)} = \sum_{j=1}^M \eta_{T,j}^{(y)}, \quad (12)$$

$$a_y^{(1)} = \sum_{j=1}^M -2\eta_{T,j}^{(y)} \xi_0 \left(v_j^{(y)} - \xi_y\right)^2, \quad (13)$$

$$a_y^{(2)} = \sum_{j=1}^M \eta_{T,j}^{(y)} \xi_0^2 \left(v_j^{(y)} - \xi_y\right)^4. \quad (14)$$

By performing this decomposition, the y -axis constant factors $a_y^{(0,1,2)}$ are computed just once for each time step, requiring linear time. Computation of Eq. (7) requires constant time, so it can be computed for all frequency components in linear time. For the x axis, the constant factors $a_x^{(0,1,2)}$ can be computed analogously. Thus, the normalization factors c_0, c_x, c_y and the adjusted 1D diffraction efficiencies $\eta_i^{(x)}, \eta_j^{(y)}$ can be computed efficiently. In general, such a decomposition can be performed for other forms of Ξ as long as they consist of polynomials.

B. Single-axis diffraction efficiency control

For each axis, the required adjusted diffraction efficiencies are achieved by tuning the rf power for each tone. Equation (3) admits an exact solution for the required rf power of each component:

$$P(v_i) = \underbrace{\frac{\eta_i \beta(v_i)}{\alpha(v_i)}}_{\text{local scaling factor}} \underbrace{\frac{4}{\pi^2 \sum_{l=1}^N \frac{\eta_l}{\alpha(v_l)}} \arcsin^2 \sqrt{\sum_{l=1}^N \frac{\eta_l}{\alpha(v_l)}}}_{\text{global scaling factor}}. \quad (15)$$

For the signal computation it is beneficial to initially consider the local power terms only and then rescale the summed signal. This requires a single arcsin computation and a single iteration over all tones.

We found that both the peak diffraction efficiency α and the saturation power β can be modeled well by a power series (see Sec. IV B). To efficiently compute α and β , the values for the functions are discretized and stored in constant memory (256 points per dimension).

For readout, linear interpolation is performed for a smooth transition between discretization points. As the frequency of each tone changes only slowly, it is almost constant over short time scales. Therefore, spatially related

memory regions are accessed for computing α and β , resulting in fast read performance.

It is important to note that, if not properly optimized, multitone signals can easily surpass the regime of linear amplification due to constructive interference of frequency components. Depending on the relative phases, different peak-to-mean amplitude ratios are incurred, also called crest factors. In general, low crest factors are desirable as they result in better signal-to-noise ratios reducing intermodulation [61] and may not exceed power specifications of the AOD. Crest factor optimization is beyond the scope of this work as it is a fundamental problem in rf signal engineering [62]. In applications of optical tweezers using static beam patterns known ahead of time, it has proven useful to iteratively optimize the phases [63–65] as it significantly improves signal quality. As it is not possible to compute all possible frequency combinations in advance, a direct method needs to be used for phase optimization [66–70]. For the reactive open-loop control system presented, we use the precalculation described in Ref. [69] commonly known as the Narahashi phases, as the computational effort is manageable, but provides a better crest factor reduction compared to other direct methods.

C. Real-time system

The system used in this work [Fig. 1(a, left)] consists of a controller on the CPU that manages which 2D beam patterns should be generated, the signal computation on the GPU, and finally the signal output using an AWG. The data transfer between GPU and AWG is performed via remote direct memory access (RDMA) to minimize latency and CPU load. The AWG runner is highly isolated due to the strict real-time requirements of the AWG in streaming mode. Timing jitter above 10 μ s can result in a buffer underflow, crashing the AWG. The RDMA is performed in blocks that get transferred at a time before they can be output. On the GPU, at least two data blocks should be allocated such that while one is transferred to the AWG, the other can be written with the subsequent waveform. The buffer sizes directly impact the achievable latency of the system, as in the worst case all buffers need to be output before a newly requested waveform gets generated. Therefore, the buffer size should be chosen as small as possible while keeping the system stable. The time bin defined by buffer size and sample rate must suffice for all communication between CPU, GPU, and AWG, as well as the computation of a new block and the data transfer. To avoid buffer underflow, a real-time system is crucial. As only one short part of the waveform is computed at a time, the reference points for GPU computation are regularly updated by the CPU with each requested waveform block. This approach prevents loss of precision caused by the use of single-precision floating-point numbers on the GPU.

D. Waveform streaming performance

The hardware for this real-time system consists of an AMD Ryzen Threadripper PRO 5955WX CPU, an Nvidia Ada 4500 graphics card containing the GPU for waveform computation, and an additional graphics card for rendering tasks. Moreover, the system houses the Spectrum M4i.6631-x8 AWG (two channels, 16 bit, 1.25 GS/s). All real-time devices are attached to NUMA node 1 which is isolated from other processes, with the GPU and AWG being within the same root complex for ideal data transfer conditions. We do not reach the maximum output sample rate of the AWG, as the PCI Express implementation (Gen2 x8) limits continuous data streaming to 2.8 GB/s, resulting in 700 MS/s per channel for our 2D system. Hence, at each channel the signal is output with a time resolution of 1.4 ns.

The user defines the requested waveform through instruction sets that comprise initial conditions for frequency, phase, and diffraction efficiency of all tones, as well as their derivatives defining waveform evolution for a specified time step. By stitching multiple sets, dynamic pattern sequences defined by flexible multitone parameter trajectories can be programmed. The queue is managed by the CPU and continuously transmitted to the GPU. Subsequent queue entries are output at the full time resolution of 1.4 ns without additional latency between arbitrary tone patterns.

Assuming an empty queue, the peak latency between requesting a new pattern and the output of this pattern as analog waveform is defined by the signal buffer sizes on the GPU and AWG. Minimization of the transfer size to 256 kB for the AWG with double buffering on the GPU and a single buffer on the AWG results in a peak latency of 257 μ s and a minimal latency of 164 μ s.

In this configuration, our software implementation for multitone rf signal processing on the GPU achieves sustained computation of waveforms including model-based efficiency control for arbitrary patterns of up to 50×50 tones. The system handles static and dynamic patterns identically. Thus, all patterns presented in this paper have been generated in real-time mode. In addition, the computation handles all patterns as if they consist of 50 tones, some of which may have zero amplitude. This approach ensures that the system is stable independently of the pattern. It also simplifies code optimization because an empty pattern has the same computational overhead as a dynamic 50×50 pattern. The system is stable over multiple days without any buffer underflows. All performance parameters can be further improved with more powerful state-of-the-art hardware.

IV. EXPERIMENTAL VALIDATION

We now describe the implementation and experimental evaluation of this system. Measured diffraction efficiencies

always reflect the combined response of both AODs and are labeled $\eta_M^{(\text{cy})}$. Absolute rf power levels are given as root-mean-square (rms) values. The AWG output for each channel is amplified by a Minicircuits ZHL-03-5WF+ power amplifier. We calibrated the AWG's output range such that for an rms rf power of 2 W of a single-tone signal at a frequency of 101 MHz the AWG power fraction is defined as 1. For parameter sampling and measurements, we parametrize the applied rf power by the output power fraction of the AWG card that is digitally set in the interval $[0, 1]$. Unless stated otherwise, uncertainties are given as the 95th percentile (i.e., 2σ deviation) throughout this paper.

A. Optical setup: 2D AOD tweezers

The optical setup is schematically depicted in Fig. 1(a). The AA Opto-Electronics DTSXY-400-800 combines two single-axis AODs that are oriented perpendicularly to deflect along the y and x dimensions (in that order). The AODs are based on a TeO_2 crystal in shear-mode configuration with an active aperture of $7.5 \text{ mm} \times 7.5 \text{ mm}$. The optical transmission per AOD is 99% according to the data sheet.

Light at a wavelength of 800 nm is supplied via a single-mode polarization-maintaining fiber with linear polarization matched to the input polarization axis of the first AOD. The laser beam is collimated to a beam waist of 1.7 mm ($1/e^2$ radius) at the position of the AOD. Behind the AOD, a lens with focal length of 63.5 mm focuses the diffracted beam such that the angular deflection is turned into a displacement in the back focal plane. An aperture blocks any order other than the desired first order. The front focal plane of the lens is positioned approximately in the middle of the two AOD crystals. The selected optical parameters, specifically the beam size at the AOD crystals, are predetermined by our targeted application. We expect no performance limitations to arise when using the full available aperture of the AODs.

An initial evaluation of the AOD showed that the incident beam and the diffracted first-order beam are collinear at a frequency of 104 MHz. The minimum of the central dip in diffraction efficiency, typical for wide-angle AODs using birefringent crystals [see Fig. 2(a)], is at $\nu_c = 101 \text{ MHz}$. We use ν_c as center frequency for the scanning interval and corresponding measurements in the work presented. The approximate bandwidth of 42 MHz results in a scan angle of 50.2 mrad and translates into a 2D scan range of $3.16 \text{ mm} \times 3.16 \text{ mm}$ in the back focal plane for the selected focal length. The focused spot waist is 11 μm ($1/e^2$ radius). For applications that require a modified addressable area, spatial resolution, or working distance, the back focal plane can be relayed using reimaging optics (not shown). A power meter with a thermopile sensor measures the laser power. The total measurement




FIG. 2. Parameter determination for model-based diffraction efficiency control using randomly sampled single-tone data. (a) Peak diffraction efficiency α_N and (b) saturation power β in units of the AWG power fraction were modeled for the x -axis AOD (solid blue) and the y -axis AOD (dashed yellow). (c) The interaction term Ξ resolves correlations between the two crossed AODs. (d) Histogram of the residual prediction error. The measured data are modeled with an error below $\pm 0.7\%$ (2σ deviation).

uncertainty, including calibration uncertainty, repeatability, linearity, and uniformity errors is specified to $\pm 3\%$. For relative measurements, the calibration uncertainty can be neglected. In this case, the resulting uncertainty is $\pm 2.3\%$. All measurements use random sampling to reduce measurement errors due to thermal drifts that depend on previously measured power levels. When applicable, we performed a baseline measurement after heating the sensor to compensate for this effect. In addition, fits include an offset term.

The laser power was permanently stabilized using a photodetector and digital intensity controller [71] (not shown). The reference power level was 750 mW, measured directly in front of the AOD by the reference power meter [see Fig. 1(a)]. The measured diffraction efficiencies $\eta_M^{(xy)}$ are defined as the ratio of the measured power of the first diffracted order in the back focal plane and the reference power. Images of the beam pattern in the back focal plane [Figs. 1(b), 1(c), and 7(a)] were recorded with a CMOS camera (sensor size 5.60 mm $\times$ 3.16 mm, pixel size 1.45 $\mu\text{m} \times$ 1.45 μm) without additional imaging optics. For camera measurements, the intensity stabilization was disabled and the input laser power was reduced to avoid saturation of the CMOS sensor.

B. Model parameter determination

The model-based control requires predetermined values of α , β , and the interaction parameters ξ_0 , ξ_x , and ξ_y .

For this purpose, the AOD response in terms of $\eta_M^{(xy)}$ was sampled in the relevant frequency range in a randomized fashion and a least-squares estimator was used to determine the model parameters. As per measurement protocol, losses at optical surfaces, apertures, and within optical components are captured in $\eta_M^{(xy)}$ and thereby incorporated in α and the interaction parameters of the model.

We performed sample acquisition using single-tone signals on both AODs with frequencies uniformly drawn from continuous values between 80 MHz and 122 MHz and AWG power fractions also uniformly between 0 and 1. A total of 20×10^3 measurements were collected for sufficient statistics. For α and β polynomials of different degrees were fitted and rated by the model error on a validation dataset. We settled on polynomials of degree $N = 40$ as the error was minimized and no overfitting was observed. Our approximation of the diffraction efficiency is only valid up to the first diffraction maximum of the AOD (see Sec. II). Collected samples exceeding saturation adversely impact the model accuracy. For this reason, we first fitted the model with all data points and used the approximate values of β to remove samples beyond saturation. Then the fit was repeated for the remaining 14×10^3 samples.

It is important to note that during parameter sampling, but also for typical applications of the model-based control, only the product $\alpha^{(x)}\alpha^{(y)}$ needs to be determined. From this combined measurement, the individual peak diffraction efficiency of each axis can be extracted only up to a multiplicative constant. We introduce the normalized peak diffraction efficiencies

$$\alpha_N^{(x)} = c_\alpha \alpha^{(x)}, \quad (16)$$

$$\alpha_N^{(y)} = \frac{1}{c_\alpha} \alpha^{(y)}, \quad (17)$$

with a normalization constant c_α in analogy to the discussion of Eqs. (7)–(10). In practice, our approach uses the identical fraction of the peak diffraction efficiency for both axes when determining the required rf power for a specific requested single-axis diffraction efficiency [Eq. (15)]. Therefore, the normalization constant does not influence the model performance. For the complete 2D expression, the normalization cancels in Eq. (4) due to $c_\alpha \times \frac{1}{c_\alpha} = 1$. As such, we fitted the normalized peak diffraction efficiencies $\alpha_N^{(x)}$, $\alpha_N^{(y)}$ instead of $\alpha^{(x)}$, $\alpha^{(y)}$ and used their values for the model. The additional degree of freedom introduced by normalization is eliminated by setting

$$\alpha_N^{(x)}(\nu_c) = \alpha_N^{(y)}(\nu_c). \quad (18)$$

This choice maximizes the usable scanning range for our AOD model, assuming that we request equal diffraction efficiencies per axis. Different normalizations may be beneficial for other devices.

Figures 2(a)–2(c) visualize the final fit parameters. The peak diffraction efficiency α_N [Fig. 2(a)] expresses the quality of wave vector matching and therefore reflects the expected band shape of the acousto-optic interaction in the case of birefringent diffraction [22]. The saturation power β [Fig. 2(b)] shows a more complex behavior. While β is expected to decrease with increasing wave vector mismatch, which is the case for the lowest and highest frequencies, additional frequency-dependent transfer functions, such as the frequency response of the power amplifier, the piezoelectric transducer, and the impedance matching circuit within the AOD, strongly influence the obtained shape. Our approach directly adjusts the digital output power of the AWG card to compensate for frequency-dependent rf power changes. From the presentation of the interaction term Ξ [Fig. 2(c)] it becomes evident that the geometric factors in the 2D system account for a reduction in diffraction efficiency of the order of 20% at the edges of the accessible frequency range. The extracted value of the interaction strength $\xi_0 = 0.94 \pm 0.01$. The dashed lines in Fig. 2(c) mark the frequencies for least geometric impairment, $\xi_x = 108.02(8)$ MHz and $\xi_y = 102.20(3)$ MHz. The deviation from ν_c indicates that the alignment, (i.e., the wave vector matching) was better along the y axis than the x axis. This is also visible in the stronger asymmetry observed in α_N along the x axis in Fig. 2(a). The following subsections prove that our model-based control can easily compensate for potential wave vector mismatches. Figure 2(d) shows a histogram of the residual prediction error of the fit model, which is defined as the difference between the measured value $\eta_M^{(xy)}$ and the model-predicted diffraction efficiency $\eta_{i,j}^{(xy)}$ for each sample. On basis of the determined model parameters, 95% of the residuals fall within $\pm 0.7\%$ diffraction efficiency error for the full frequency range (2σ deviation).

C. Real-time performance

All beam patterns presented in this paper are generated in real-time operation. To validate the latency specification derived from digital data flow analysis in Sec. III D in the optical signal, the actual latency between submitting a new pattern and recording the resulting light signal was measured. For this, the pattern queue was emptied, so that a zero-volt signal level is output to the AOD, blanking the deflected beam pattern. Then the control program emits an output trigger and directly thereafter submits a pattern to the queue that generates a beam pattern illuminating a fast photodiode. An oscilloscope measures the delay between the control trigger and a light signal on the photodiode for 10^4 samples. The cumulative distribution of latencies in Fig. 3 reveals that the median latency is below $207 \mu\text{s}$ and peak latency below $247 \mu\text{s}$. The measured upper bound is even lower than the specified peak latency, confirming real-time operation with peak latency of $257 \mu\text{s}$.




FIG. 3. Measured latency between submission of a pattern to an empty data queue ($t = 0$) and detection of the requested light pattern on a photodiode (10^4 samples). A peak latency below $247 \mu\text{s}$ is measured.

D. Single-tone efficiency control

In a second set of measurements, we validated the capabilities of the diffraction efficiency control system for single-tone applications. For comparison, we performed characteristic measurements of the AOD response as a function of the applied rf power without diffraction efficiency control.

Figure 4 depicts the measured diffraction efficiency when driving the y axis with a single tone at ν_c and the x axis with a single tone at 116 MHz. The behavior of $\eta_M^{(xy)}$ as a function of the x -axis AWG power fraction is shown in Fig. 4(a) for disabled efficiency control. The y -axis AOD was driven at a constant AWG power fraction of 0.9 and the x -axis power fraction was sampled between 0 and 1. The nonlinear AOD response is well described by Eq. (2) (dashed line), where all constant losses, including the effect of the static y -axis AOD, are incorporated into α_N .

In comparison, we used our model-based control system to set the target diffraction efficiency $\eta_T^{(y)} = 0.9$, while




FIG. 4. Measured diffraction efficiency with and without efficiency control. A single tone is applied to both axes. (a) Without efficiency control, a constant AWG power fraction of 0.9 is set for the y axis while sweeping the x -axis AWG power fraction between 0 and 1. (b) Efficiency control set to 0.9 diffraction efficiency for the y axis while sampling the x -axis diffraction efficiency from 0 to 1 (i.e., $\eta_T^{(xy)} \in [0, 0.9]$) results in a linear behavior up to saturation.




FIG. 5. Measured 2D diffraction efficiency for single-tone operation at both axes. (a) Unregulated operation with fixed AWG power fraction 0.9 for both AODs as reference. (b) Operation with the control system set to generate a flat response of 80% target diffraction efficiency. The color bar on the right holds for (a),(b). (c) The uniformity of the controlled 2D diffraction efficiency is analyzed as a function of the 2D frequency span around ν_c . Mean value and variation (shaded areas) are plotted for 40–80% target diffraction efficiency. (d) Mean diffraction error (i.e., the deviation from the target efficiency) and its variation visualizing the precision and accuracy for three selected target values.

sampling $n_T^{(x)} \in [0, 1]$. This defines the measurement range $\eta_T^{(xy)} \in [0, 0.9]$ depicted in Fig. 4(b). The linear scaling of $\eta_M^{(xy)}$ as a function of $\eta_T^{(xy)}$ under efficiency control is evident and well described by the identity relation (dashed line) up to the point where the AOD reaches its maximum efficiency and the response saturates.

For 2D pattern generation with arbitrary frequencies, one has to acknowledge that α , β , and Ξ are frequency dependent (Fig. 2). A simplified approach without efficiency control is to choose global rf power levels that run the AODs close to saturation. Figure 5(a) shows the measured 2D diffraction efficiency landscape using this approach with the AWG power fraction set to 0.9 for both AODs. Contour lines indicate regions in which the AOD system operates above the labeled thresholds. The maximum diffraction efficiency is 92%. While the continuous response allows the user to select local frequency patches with minor variation, the measured landscape reveals a considerable variation even for the central region within the 79% contour line. This clearly shows the limits of the unregulated approach.

For many applications, ensuring genuine uniformity of diffraction efficiencies for a given working range of the

AOD is crucial. To benchmark the presented control system, we sampled the 2D diffraction efficiency at target levels $\eta_T^{(xy)}$ between 40% and 80% in steps of 10% over the frequency span of 42 MHz per axis. For all realizations, both single-axis diffraction efficiencies were set equally: $\eta_T^{(x)} = \eta_T^{(y)} = (\eta_T^{(xy)})^{1/2}$. Figure 5(b) gives the result for $\eta_T^{(xy)} = 80\%$. In comparison to the measurement with fixed rf power, the flatness is strongly improved. Within the 79% contour line, 95% of the measured efficiencies are in the range 79.2% to 80.9%. Figure 5(c) shows the measured performance for all targeted values $\eta_T^{(xy)}$. The mean value and the variation of the data for $\eta_M^{(xy)}$ are plotted as a function of the 2D frequency span around ν_c . Results confirm that our approach remains linear with flat response independent of the frequencies, as long as the requested efficiency does not exceed the achievable peak efficiency. The latter effect causes a reduction in the mean value of $\eta_M^{(xy)}$ and an increased variation for large frequency spans. As displayed in Fig. 5(d), the edges of the error bands intersect a diffraction efficiency error of 1% for a 2D span above 33 MHz for the 80% data, above 37 MHz for the 70% data, and above 40 MHz for the 60% data. Noting that the apparent precision in this measurement series is better than the specified uncertainty of the laser power meter, we observe two trends for data within the frequency range where the requested efficiency $\eta_T^{(xy)}$ is still accessible. First, the variation given by the error bands increases with larger frequency spans. Second and more prominently, the curves show a nonzero slope. We attribute both effects to shortcomings of the model used for the interaction term Ξ [Eq. (5)] including only up to fourth-order corrections. These effects gain increasing significance for larger deviations from ξ_x and ξ_y due to an increased wave vector mismatch [see also Fig. 2(c)].

E. Two-dimensional multibeam efficiency control

The nonlinear acousto-optic response is additionally evolved for multitone waveforms. All frequency components interfere with the diffraction of each individual tone due to saturation effects. This makes multitone operation highly impractical without efficiency control. To demonstrate the effect, we conducted a measurement of the dual-tone crosstalk on the x -axis AOD while driving the y -axis AOD with a single tone of constant power at ν_c . The setup is illustrated in Fig. 6(a). The x -axis AOD received a dual-tone signal with one tone remaining static at $\nu_{\text{static}} = 116$ MHz, while the other tone was randomly sampled in the range $\nu_{\text{random}} \in [82 \text{ MHz}, 106 \text{ MHz}]$. The rf power of both tones was randomly sampled. The diffracted beams up to 106 MHz were blocked in front of the power meter. Only the static beam was measured.

The static y -axis AOD was driven at AWG power fraction 0.9. Without efficiency control, the power sampling




FIG. 6. Characterization of the crosstalk for dual-tone operation. (a) Schematic setup. The y -axis (out of plane) AOD is driven with a constant single-tone signal. For the x -axis AOD, the diffraction efficiency of one tone at the static frequency of 116 MHz is recorded while the power and frequency of a second tone are sampled randomly. (b) Measured diffraction efficiency with disabled efficiency control as a function of the static-tone AWG power fraction. Crosstalk-induced spreading of the efficiency is observed. The data are color-coded according to their normalized incidence (2×10^4 samples, 100 bins per axis). (c) With efficiency control enabled, a linear response as a function of the requested target efficiency is obtained. (d) With efficiency control the residual 2σ deviation from the target diffraction efficiency is $\pm 0.5\%$.

of both x -axis tones was restricted by limiting the combined dual-tone x -axis AWG power fraction to 0.5. This constraint accommodates the crest factor, as a dual-tone signal has a peak power twice as high as a single-tone sine wave of equal mean power (rms). Figure 6(b) shows $\eta_M^{(xy)}$ as a function of the x -axis AWG power fraction of the static tone at 116 MHz. Without a second tone, the response follows Eq. (2), as observed for the measurement in Fig. 4(a). For the dual-tone measurement, this corresponds to the samples with no power in the second tone and gives the upper boundary of the data in Fig. 6(b). The interference of the second tone causes a reduction of $\eta_M^{(xy)}$ for the measured component. The main reduction stems from saturation with increased power values of the second tone. This effect is further enhanced by the frequency dependence of β [see Fig. 2(b)]. At 0.25 AWG power fraction, the spread in $\eta_M^{(xy)}$ is 35% to 47%. The lower boundary arises due to the power constriction of the second tone defined by the maximally allowed dual-tone AWG power fraction, limiting the second tone's power fraction range to 0–0.25 at this point.

Figure 6(c) depicts the respective measurement with the efficiency control system in operation. The y -axis AOD was set to $\eta_T^{(y)} = 0.9$. The dual-tone waveforms for the x -axis AOD were sampled up to a limit of $\eta_T^{(x)} = 0.75$ per tone while constricting samples to combinations that comply with $\eta_{T,\text{static}}^{(x)} + \eta_{T,\text{random}}^{(x)} \leq 0.75$ for the combined target efficiency of both tones. This defines the depicted measurement range of $\eta_{T,\text{static}}^{(xy)} \in [0, 0.675]$. The data closely follow the sought relation $\eta_M^{(xy)} = \eta_{T,\text{static}}^{(xy)}$ (dashed). Figure 6(d) shows the deviation from the expected linear response function. The residual error is $\pm 0.5\%$ (2σ) and below $\pm 1\%$ for all samples, demonstrating how effectively our model can compensate for crosstalk between two tones.

Multibeam efficiency control is applied to generate the 2D beam patterns of Figs. 1(b) and 1(c) in real-time mode. A quantitative analysis of the efficiency control of large 2D laser spot arrays is performed for the pattern of Fig. 1(b). For this intensity-equalized 50×50 array, the measured integral diffraction efficiency of 0.495 ± 0.023 agrees well with the targeted efficiency of 0.5. When analyzing the homogeneity of all 2500 laser spots, we observe a relative rms spread of 5.3%. This performance is comparable to significantly more involved methods that employ extensive calibration of a static array implementation [72].

F. Intermodulation distortion

A major source of disturbances in multitone pattern generation is intermodulation. Although control of intermodulation is outside the scope of the applied model, we investigated its effect since we attribute part of the intensity variation in Fig. 1(b) to intermodulation. Due to nonlinearities in signal processing, including amplification and signal transduction to the AOD's crystal, spurious signals are created. A poorly optimized crest factor strongly favors intermodulation. The leading order is found at the sum and difference frequencies of pairs of programmed tones for the respective axis which are usually outside the range of operation. In next order, third-order intermodulation gives rise to frequency components that lie within the frequency range of the original waveform.

For a multitone intermodulation analysis, we recorded camera images of 2D tweezer patterns that were deliberately chosen to be prone to intermodulation by using grids of equally spaced frequencies and setting the amplitudes for the central frequency components (i.e., the target power for the central row and column) to zero. Intermodulation of neighboring tones could create a spurious signal at these exact frequencies and corresponding beam positions. Figure 7(a) shows a realization with 11 tones in both dimensions spaced by 0.75 MHz, creating four clusters of 5×5 equidistant tweezers that are separated by a central 1.5 MHz gap. An integral target diffraction efficiency of 0.3 was set for the grid and 0.245 ± 0.023 was measured.




FIG. 7. Intermodulation analysis of a 2D multispot optical tweezer array. Tweezer positions are defined by a multitone waveform of 11 equally spaced frequencies per axis. Diffraction efficiencies of all tones are equalized while the central rf component is blanked. (a) Camera image of the intensity pattern in the back focal plane. The tweezer array covers an area of $1\text{ mm} \times 1\text{ mm}$. (b) Camera signal summed along the y axis to reveal any intermodulation-induced signal. The gap shows a third-order intermodulation peak at -23 dB .

Although we attribute this reduction of efficiency to intermodulation, no spurious spots are immediately visible on the 2D camera image of Fig. 7(a). To quantify the magnitude the intermodulation in the light field, we integrated the camera signal along the y axis as shown in Fig. 7(b). For the blanked frequency at pixel 300, an optical signal with an intensity reduction of 23 dB is visible with similar peaks left and right of the target beam pattern. The uniformity of the spectrum is also affected, with a relative rms spread of 9.8% . These measurements highlight the importance of taking into account the characteristics of the specific device for high-fidelity multitone operation. Crest factor optimization with direct methods is particularly suitable for multitone patterns with fixed frequency spacing, but time-varying frequency spacings may result in reduced suppression of intermodulation, which has to be closely monitored and presumably mitigated by predistortion methods [73,74]. Since predistortion and advanced phase tuning for crest factor reduction strongly depend on the acousto-optic device as well as the application, their optimization is beyond the scope of this work.

V. CONCLUSION

Building on the theoretical framework of the acousto-optic response that is described by coupled-wave theory, we deduced a resource-efficient model for the computation of the diffraction efficiencies of multitone signals in AODs. The description enfolds 2D pattern generation by incorporating geometric effects that arise from operating two mutually perpendicular AODs in series. Model-parameter determination only requires measurements of the laser power after 2D deflection with a single-tone signal at each axis. This allows the straightforward integration in existing setups and optical instruments since model calibration

can be realized by addition of a simple beam splitter and a power meter.

On this basis, we developed an open-loop control system that ensures strict intensity matching for user-programmable random-access scanning applications with acousto-optically deflected laser beams. Our tailored software implementation inverts the model in real time on commodity hardware to determine the required rf power per frequency component to achieve a specific diffraction efficiency target pattern, enabling on-the-fly control of intensity profiles for arbitrary spatiotemporal multibeam trajectories. In the configuration presented, it is capable of streaming tweezer arrays with up to 2500 spots in continuous mode. Waveforms are generated with adequate time resolution for controlling common AODs at their full acousto-optical bandwidth.

The performance of the approach presented was validated in a 2D AOD setup. Linear control over diffraction efficiencies in single-tone and multitone operation was achieved. Residual errors were within the measurement uncertainty. Intermodulation was found to be sufficiently suppressed by direct methods for crest factor optimization that are feasible for real-time operation. The presented real-time system for reactive 2D multibeam laser scanning will benefit abundant applications that employ optical forces by optical tweezers or precise deposition of light energy with exquisite spatiotemporal resolution. In our work, we focus on quantum science with arrays of atomic qubits, where laser addressing of individual qubits is essential at every stage of operation. Parallelized qubit transport will strongly improve the preparation of defect-free many-body systems [53,55,56,58,75,76], their continuous operation [33,34,47,48,77,78], as well as increasing data rates in sensing applications [28]. Moreover, time-critical parallelized optical control at the individual qubit level becomes exceedingly momentous in quantum information processing. The application of quantum algorithms and error correction protocols [32,49–51] necessitates local qubit addressing for coherent quantum-state control, mid-circuit measurements, conditional quantum operations, and memory transfer in zoned architectures [79]. High-fidelity rendering of programmed intensities, as demonstrated in this work, is vital. Deviations directly impact the quantum-state phases of the qubits, scrambling quantum information. For large-scale implementations [27,36,75,80], high-bandwidth multibeam steering is of immediate relevance for future progress.

ACKNOWLEDGMENTS

We thank S. de Léséleuc for insightful discussions. We acknowledge financial support by the Federal Ministry of Education and Research (BMBF) [Grant No. 13N15981], by the Federal Ministry of Research, Technology and Space (BMFT) [Grant No. 13N17366], and

by the Deutsche Forschungsgemeinschaft (DFG – German Research Foundation) [Grant No. BI 647/6-1 and BI 647/6-2, Priority Program SPP 1929 (GiRyd)].

DATA AVAILABILITY

The data that support the findings of this article are not publicly available because they contain commercially sensitive information. The data are available from the authors upon reasonable request.

- [1] T. Pinheiro, M. Morais, S. Silvestre, E. Carlos, J. Coelho, H. V. Almeida, P. Barquinha, E. Fortunato, and R. Martins, Direct laser writing: From materials synthesis and conversion to electronic device processing, *Adv. Mater.* **36**, 2402014 (2024).
- [2] E. Skliutas, M. Lebedevaite, E. Kabouraki, T. Baldacchini, J. Ostrauskaite, M. Vamvakaki, M. Farsari, S. Juodkazis, and M. Malinauskas, Polymerization mechanisms initiated by spatio-temporally confined light, *Nanophotonics* **10**, 1211 (2021).
- [3] C. Arnoux, L. A. Pérez-Covarrubias, A. Khaldi, Q. Carlier, P. L. Baldeck, K. Heggarty, A. Banyasz, and C. Monnereau, Understanding and overcoming proximity effects in multi-spot two-photon direct laser writing, *Addit. Manuf.* **49**, 102491 (2022).
- [4] P. O. Bayguinov, D. M. Oakley, C.-C. Shih, D. J. Geanon, M. S. Joens, and J. A. J. Fitzpatrick, Modern laser scanning confocal microscopy, *Curr. Protoc. Cytometry* **85**, e39 (2018).
- [5] A. D. Elliott, Confocal microscopy: Principles and modern practices, *Curr. Protoc. Cytometry* **92**, e68 (2020).
- [6] M. Duocastella, S. Surdo, A. Zunino, A. Diaspro, and P. Saggau, Acousto-optic systems for advanced microscopy, *J. Phys.: Photonics* **3**, 012004 (2021).
- [7] D. Mahecic, D. Gambarotto, K. M. Douglass, D. Fortun, N. Banterle, K. A. Ibrahim, M. Le Guennec, P. Gönczy, V. Hamel, P. Guichard, and S. Manley, Homogeneous multifocal excitation for high-throughput super-resolution imaging, *Nat. Methods* **17**, 726 (2020).
- [8] A. Ashkin, Optical trapping and manipulation of neutral particles using lasers, *Proc. Natl. Acad. Sci.* **94**, 4853 (1997).
- [9] J. Gieseler, J. R. Gomez-Solano, A. Magazzù, I. P. Castillo, L. P. García, M. Gironella-Torrent, X. Viader-Godoy, F. Ritort, G. Pesce, A. V. Arzola, K. Volke-Sepúlveda, and G. Volpe, Optical tweezers—from calibration to applications: A tutorial, *Adv. Opt. Photon.* **13**, 74 (2021).
- [10] J. R. Moffitt, Y. R. Chemla, S. B. Smith, and C. Bustamante, Recent advances in optical tweezers, *Annu. Rev. Biochem.* **77**, 205 (2008).
- [11] P. Ricci, G. Sancataldo, V. Gavryusev, F. S. Pavone, P. Saggau, and M. Duocastella, Acousto-optic deflectors in experimental neuroscience: overview of theory and applications, *J. Phys.: Photonics* **6**, 022001 (2024).
- [12] C. Gonzalez-Ballester, M. Aspelmeier, L. Novotny, R. Quidant, and O. Romero-Isart, Levitodynamics: Levitation and control of microscopic objects in vacuum, *Science* **374**, eabg3027 (2021).
- [13] J. Yan, X. Yu, Z. V. Han, T. Li, and J. Zhang, On-demand assembly of optically levitated nanoparticle arrays in vacuum, *Photon. Res.* **11**, 600 (2023).
- [14] B. Siegel, G. Afek, C. Lowe, J. Wang, Y.-H. Tseng, T. W. Penny, and D. C. Moore, Optical levitation of arrays of microspheres, *Phys. Rev. A* **111**, 033514 (2025).
- [15] Y. Liu and K.-K. Ni, Bimolecular chemistry in the ultracold regime, *Annu. Rev. Phys. Chem.* **73**, 73 (2022).
- [16] A. A. Manshina, I. I. Tumkin, E. M. Khairullina, M. Mizoshiri, A. Ostendorf, S. A. Kulinich, S. Makarov, A. A. Kuchmizhak, and E. L. Gurevich, The second laser revolution in chemistry: Emerging laser technologies for precise fabrication of multifunctional nanomaterials and nanostructures, *Adv. Funct. Mater.* **34**, 2405457 (2024).
- [17] A. M. Kaufman and K.-K. Ni, Quantum science with optical tweezer arrays of ultracold atoms and molecules, *Nat. Phys.* **17**, 1324 (2021).
- [18] B. Cheng, X.-H. Deng, X. Gu, Y. He, G. Hu, P. Huang, J. Li, B.-C. Lin, D. Lu, Y. Lu, *et al.*, Noisy intermediate-scale quantum computers, *Front. Phys.* **18**, 21308 (2023).
- [19] T. M. Graham, E. Oh, and M. Saffman, Multiscale architecture for fast optical addressing and control of large-scale qubit arrays, *Appl. Opt.* **62**, 3242 (2023).
- [20] G. Römer and P. Bechtold, Electro-optic and acousto-optic laser beam scanners, *Phys. Procedia* **56**, 29 (2014).
- [21] B. E. A. Saleh and M. C. Teich, in *Fundamentals of Photonics* (John Wiley & Sons, Ltd, Hoboken, 1991), Chap. 20.
- [22] I. Chang, in *Handbook of Optics: Volume V* (McGraw-Hill Professional, 2010), Chap. 6. (see also: Chap. 12 in *Handbook of Optics: Volume II*, New York, 1995).
- [23] D. Trypogeorgos, T. Harte, A. Bonnin, and C. Foot, Precise shaping of laser light by an acousto-optic deflector, *Opt. Express* **21**, 24837 (2013).
- [24] L. R. B. Picard and M. Endres, A three-dimensional acousto-optic deflector, [arXiv:2510.07633](https://arxiv.org/abs/2510.07633).
- [25] Z. Guo, R. A. H. van Herk, E. J. D. Vredenburg, and S. J. J. M. F. Kokkelmans, Acousto-optic lens for 3D shuttling of atoms in a neutral atom quantum computer, [arXiv:2510.09398](https://arxiv.org/abs/2510.09398).
- [26] Y.-H. Lu, N. Song, T. Xiang, J. Ho, T.-C. Lee, Z. Yan, and D. M. Stamper-Kurn, Astigmatism-free 3D optical tweezer control for rapid atom rearrangement, [arXiv:2510.11451](https://arxiv.org/abs/2510.11451).
- [27] L. Pause, L. Sturm, M. Mittenbühler, S. Amann, T. Preuschoff, D. Schäffner, M. Schlosser, and G. Birkel, Supercharged two-dimensional tweezer array with more than 1000 atomic qubits, *Optica* **11**, 222 (2024).
- [28] D. Schäffner, T. Schreiber, F. Lenz, M. Schlosser, and G. Birkel, Quantum sensing in tweezer arrays: Optical magnetometry on an individual-atom sensor grid, *PRX Quantum* **5**, 010311 (2024).
- [29] M. Schlosser, D. O. de Mello, D. Schäffner, T. Preuschoff, L. Kohfahl, and G. Birkel, Assembled arrays of Rydberg-interacting atoms, *J. Phys. B: At. Mol. Opt. Phys.* **53**, 144001 (2020).
- [30] H. Hwang, A. Byun, J. Park, S. de Léséleuc, and J. Ahn, Optical tweezers throw and catch single atoms, *Optica* **10**, 401 (2023).

- [31] C. M. Holland, Y. Lu, and L. W. Cheuk, On-demand entanglement of molecules in a reconfigurable optical tweezer array, *Science* **382**, 1143 (2023).
- [32] D. Bluvstein *et al.*, Logical quantum processor based on reconfigurable atom arrays, *Nature* **626**, 58 (2024).
- [33] M. A. Norcia *et al.*, Iterative assembly of ^{171}Yb atom arrays with cavity-enhanced optical lattices, *PRX Quantum* **5**, 030316 (2024).
- [34] F. Gyger, M. Ammenwerth, R. Tao, H. Timme, S. Snigirev, I. Bloch, and J. Zeiher, Continuous operation of large-scale atom arrays in optical lattices, *Phys. Rev. Res.* **6**, 033104 (2024).
- [35] L. R. B. Picard, G. E. Patenotte, A. J. Park, S. F. Gebretsadkan, and K.-K. Ni, Site-selective preparation and multistate readout of molecules in optical tweezers, *PRX Quantum* **5**, 020344 (2024).
- [36] H. J. Manetsch, G. Nomura, E. Bataille, K. H. Leung, X. Lv, and M. Endres, A tweezer array with 6100 highly coherent atomic qubits, *Nature* **647**, 60 (2025).
- [37] S. Hwang, H. Hwang, K. Kim, A. Byun, K. Kim, S. Jeong, M. P. Soegianto, and J. Ahn, Fast and reliable atom transport by optical tweezers, *Optica Quantum* **3**, 64 (2025).
- [38] C. Cicali, M. Calzavara, E. Cuestas, T. Calarco, R. Zeier, and F. Motzoi, Fast neutral-atom transport and transfer between optical tweezers, *Phys. Rev. Appl.* **24**, 024070 (2025).
- [39] S. Kim, R. R. Mcleod, M. Saffman, and K. H. Wagner, Doppler-free, multiwavelength acousto-optic deflector for two-photon addressing arrays of Rb atoms in a quantum information processor, *Appl. Opt.* **47**, 1816 (2008).
- [40] K. Barnes, P. Battaglini, B. J. Bloom, K. Cassella, R. Coxe, N. Crisosto, J. P. King, S. S. Kondov, K. Kotru, S. C. Larsen, *et al.*, Assembly and coherent control of a register of nuclear spin qubits, *Nat. Commun.* **13**, 2779 (2022).
- [41] Z. Z. Yan, B. M. Spar, M. L. Prichard, S. Chi, H.-T. Wei, E. Ibarra-García-Padilla, K. R. A. Hazzard, and W. S. Bakr, Two-dimensional programmable tweezer arrays of fermions, *Phys. Rev. Lett.* **129**, 123201 (2022).
- [42] A. Radnaev *et al.*, Universal neutral-atom quantum computer with individual optical addressing and nondestructive readout, *PRX Quantum* **6**, 030334 (2025).
- [43] S. Yang, G. Masella, V. Moeini, A. Bellahsene, C. Li, T. Bienaimé, and S. Whitlock, Compact arbitrary optical waveform modulator with digital feedback, *Phys. Rev. Appl.* **23**, 054009 (2025).
- [44] N. Defenu, T. Donner, T. Macrì, G. Pagano, S. Ruffò, and A. Trombettoni, Long-range interacting quantum systems, *Rev. Mod. Phys.* **95**, 035002 (2023).
- [45] M. Morgado and S. Whitlock, Quantum simulation and computing with Rydberg-interacting qubits, *AVS Quantum Sci.* **3**, 023501 (2021).
- [46] S. L. Cornish, M. R. Tarbutt, and K. R. Hazzard, Quantum computation and quantum simulation with ultracold molecules, *Nat. Phys.* **20**, 730 (2024).
- [47] L. Pause, T. Preuschoff, D. Schäffner, M. Schlosser, and G. Birkel, Reservoir-based deterministic loading of single-atom tweezer arrays, *Phys. Rev. Res.* **5**, L032009 (2023).
- [48] N.-C. Chiu, E. C. Trapp, J. Guo, M. H. Abobeih, L. M. Stewart, S. Hollerith, P. Stroganov, M. Kalinowski, A. A. Geim, S. J. Evered, S. H. Li, L. M. Peters, D. Bluvstein, T. T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin, Continuous operation of a coherent 3,000-qubit system, *Nature* **646**, 1075 (2025).
- [49] B. W. Reichardt *et al.*, Logical computation demonstrated with a neutral atom quantum processor, *arXiv:2411.11822*.
- [50] M. J. Bedalov *et al.*, Fault-tolerant operation and materials science with neutral atom logical qubits, *arXiv:2412.07670*.
- [51] P. S. Rodriguez *et al.*, Experimental demonstration of logical magic state distillation, *Nature* **645**, 620 (2025).
- [52] K.-N. Schymik, V. Lienhard, D. Barredo, P. Scholl, H. Williams, A. Browaeys, and T. Lahaye, Enhanced atom-by-atom assembly of arbitrary tweezer arrays, *Phys. Rev. A* **102**, 063107 (2020).
- [53] W. Tian, W. J. Wee, A. Qu, B. J. M. Lim, P. R. Datla, V. P. W. Koh, and H. Loh, Parallel assembly of arbitrary defect-free atom arrays with a multitweezer algorithm, *Phys. Rev. Appl.* **19**, 034048 (2023).
- [54] M. A. Norcia *et al.*, Midcircuit qubit measurement and rearrangement in a ^{171}Yb atomic array, *Phys. Rev. X* **13**, 041034 (2023).
- [55] S. Wang, W. Zhang, T. Zhang, S. Mei, Y. Wang, J. Hu, and W. Chen, Accelerating the assembly of defect-free atomic arrays with maximum parallelisms, *Phys. Rev. Appl.* **19**, 054032 (2023).
- [56] X. Guo, J. Winklmann, D. Stober, A. Elsharkawy, and M. Schulz, Design of an FPGA-based neutral atom rearrangement accelerator for quantum computing, in *2025 Design, Automation & Test in Europe Conference* (IEEE, New York, 2025), pp. 1–6.
- [57] J. Tu and S. Subhankar, Fast real-time arbitrary waveform generation using graphic processing units, *IEEE Trans. Signal Process.* **73**, 2368 (2025).
- [58] A. H. Dadpour, T. Khayrullin, F. Afiouni, R. E. Sabeh, A. E. Mouawad, I. E. Hajj, and A. Cooper, Low-latency control system for feedback experiments with optical-tweezer arrays, *Phys. Rev. Appl.* **24**, 024062 (2025).
- [59] A. Korpel, Acousto-optics—a review of fundamentals, *Proc. IEEE* **69**, 48 (1981).
- [60] S. Antonov and V. Kotelnikov, A review of physical principles and applications of acousto-optic deflectors on the basis paratellurite, *Phys. Astron. Int. J.* **3**, 235 (2019).
- [61] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, Atom-by-atom assembly of defect-free one-dimensional cold atom arrays, *Science* **354**, 1024 (2016).
- [62] N. B. Carvalho and W. Jang, in *Digital Front-End in Wireless Communications and Broadcasting: Circuits and Signal Processing*, edited by F.-L. Luo (Cambridge University Press, Cambridge, 2011), pp. 309–330.
- [63] P. Guillaume, J. Schoukens, R. Pintelon, and I. Kollar, Crest-factor minimization using nonlinear Chebyshev approximation methods, *IEEE Trans. Instrum. Meas.* **40**, 982 (1991).
- [64] E. Van der Ouderaa, J. Schoukens, and J. Renneboog, Peak factor minimization using a time-frequency domain swapping algorithm, *IEEE Trans. Instrum. Meas.* **37**, 145 (1988).

- [65] A. Van Den Bos, A new method for synthesis of low-peak-factor signals, *IEEE Trans. Acoust. Speech Signal Process.* **35**, 120 (1987).
- [66] H. Kitayoshi, S. Sumiea, K. Siirakawa, and S. Takeshita, in *Digest of Papers-International Test Conference* (IEEE, New York, 1985), pp. 825–834.
- [67] D. Newman, An L^1 extremal problem for polynomials, *Proc. Am. Math. Soc.* **16**, 1287 (1965).
- [68] M. Schroeder, Synthesis of low-peak-factor signals and binary sequences with low autocorrelation (Corresp.), *IEEE Trans. Inf. Theory* **16**, 85 (1970).
- [69] S. Narahashi and T. Nojima, New phasing scheme of N -multiple carriers for reducing peak-to-average power ratio, *Electron. Lett.* **30**, 1382 (1994).
- [70] Y. Shibasaki, K. Asami, R. Aoki, A. Hatta, A. Kuwana, and H. Kobayashi, in *2020 IEEE 29th Asian Test Symposium (ATS)* (IEEE, New York, 2020), pp. 1–6.
- [71] T. Preuschoff, M. Schlosser, and G. Birkel, Digital laser frequency and intensity stabilization based on the STEM-lab platform (originally Red Pitaya), *Rev. Sci. Instrum.* **91**, 083001 (2020).
- [72] A. Cooper, J. P. Covey, I. S. Madjarov, S. G. Porsev, M. S. Safronova, and M. Endres, Alkaline-earth atoms in optical tweezers, *Phys. Rev. X* **8**, 041055 (2018).
- [73] E. Biglieri, S. Barberis, and M. Catena, Analysis and compensation of nonlinearities in digital transmission systems, *IEEE J. Sel. Areas Commun.* **6**, 42 (1988).
- [74] D. R. Morgan, Z. Ma, J. Kim, M. G. Zierdt, and J. Pastalan, A generalized memory polynomial model for digital predistortion of rf power amplifiers, *IEEE Trans. Signal Process.* **54**, 3852 (2006).
- [75] R. Lin, H.-S. Zhong, Y. Li, Z.-R. Zhao, L.-T. Zheng, T.-R. Hu, H.-M. Wu, Z. Wu, W.-J. Ma, Y. Gao, Y.-K. Zhu, Z.-F. Su, W.-L. Ouyang, Y.-C. Zhang, J. Rui, M.-C. Chen, C.-Y. Lu, and J.-W. Pan, AI-enabled parallel assembly of thousands of defect-free neutral atom arrays, *Phys. Rev. Lett.* **135**, 060602 (2025).
- [76] F. Afiouni, R. El Sabeh, N. Nishimura, I. El Hajj, A. E. Mouawad, and A. Cooper, Efficient algorithms to solve atom reconfiguration problems. III. The bird and batching algorithms and other parallel implementations on GPUs, *Phys. Rev. A* **112**, 023109 (2025).
- [77] K. Singh, S. Anand, A. Pocklington, J. T. Kemp, and H. Bernien, Dual-element, two-dimensional atom array with continuous-mode operation, *Phys. Rev. X* **12**, 011040 (2022).
- [78] X. Yan, C. He, K. Wen, Z. Ren, P. T. F. Wong, E. Hagiye, and G.-B. Jo, Multi-reservoir enhanced loading of tweezer atom arrays, *Phys. Rev. Res.* **7**, 013031 (2025).
- [79] D. B. Tan, D. Bluvstein, M. D. Lukin, and J. Cong, Compiling quantum circuits for dynamically field-programmable neutral atoms array processors, *Quantum* **8**, 1281 (2024).
- [80] G. Pichard, D. Lim, E. Bloch, J. Vaneeckloo, L. Bourachot, G.-J. Both, G. Mériaux, S. Dutartre, R. Hostein, J. Paris, B. Ximenez, A. Signoles, A. Browaeys, T. Lahaye, and D. Dreon, Rearrangement of individual atoms in a 2000-site optical-tweezer array at cryogenic temperatures, *Phys. Rev. Appl.* **22**, 024073 (2024).