

An atomic boson sampler

<https://doi.org/10.1038/s41586-024-07304-4>

Received: 17 July 2023

Accepted: 13 March 2024

Published online: 8 May 2024



Aaron W. Young^{1✉}, Shawn Geller^{2,3}, William J. Eckner¹, Nathan Schine^{1,4}, Scott Glancy², Emanuel Knill^{2,5} & Adam M. Kaufman^{1✉}

A boson sampler implements a restricted model of quantum computing. It is defined by the ability to sample from the distribution resulting from the interference of identical bosons propagating according to programmable, non-interacting dynamics¹. An efficient exact classical simulation of boson sampling is not believed to exist, which has motivated ground-breaking boson sampling experiments in photonics with increasingly many photons^{2–12}. However, it is difficult to generate and reliably evolve specific numbers of photons with low loss, and thus probabilistic techniques for postselection⁷ or marked changes to standard boson sampling^{10–12} are generally used. Here, we address the above challenges by implementing boson sampling using ultracold atoms^{13,14} in a two-dimensional, tunnel-coupled optical lattice. This demonstration is enabled by a previously unrealized combination of tools involving high-fidelity optical cooling and imaging of atoms in a lattice, as well as programmable control of those atoms using optical tweezers. When extended to interacting systems, our work demonstrates the core abilities required to directly assemble ground and excited states in simulations of various Hubbard models^{15,16}.

There is a rich history of studying quantum optics with atoms instead of photons¹⁷, including demonstrations of two-atom interference^{14,18–20}. Scaling up these demonstrations requires reliable preparation of chosen patterns of many identical atoms, evolution under non-interacting dynamics that can exchange the positions of the atoms and high-fidelity detection of the atom positions after their evolution. The latter two requirements can be met with quantum gas microscopy, in which it is possible to prepare and detect thousands of atoms in an optical lattice in which those atoms can tunnel and interfere^{21,22}. Although in principle the first requirement can be met by additionally using sophisticated optical control techniques to isolate a subset of atoms from a larger many-body state^{20,23–27}, in practice, the typical state fidelities and cycle times of up to tens of seconds make it challenging to perform a large-scale boson sampling demonstration. To explore alternative routes to improve the speed and quality of state preparation, we turn to tools that have been developed to rapidly image²⁸, optically cool^{29–32} and deterministically rearrange individual atoms trapped in optical tweezers^{33,34}. Using the tweezers to implant atoms in a tunnel-coupled optical lattice^{35,36} allows for both fast state preparation (in hundreds of milliseconds) and the required dynamics for boson sampling. Moreover, the tweezers can be used to modify the lattice potential³⁶ to implement different non-interacting dynamics with low loss.

Combining the aforementioned tools for state preparation, evolution and detection in a single apparatus enables large-scale demonstrations of boson sampling with Fock states that were not feasible before this work. In particular, we study specific instances of boson sampling involving up to 180 atoms spread over about 1,015 sites in the lattice. Importantly, direct verification of boson sampling is infeasible for these system sizes. To build confidence that these boson sampling experiments sample from a very large state space and are not feasible to simulate classically, we develop and implement broadly applicable

tests to examine and quantify the performance of our system for boson sampling. These tests include (1) stringent tests of indistinguishability with up to eight atoms; (2) characterizations of the atom evolution from data with one and two atoms; (3) evidence of expected bunching features as a result of interference for a range of atom numbers and effective particle statistics. This diversity of tests is enabled by the programmability of the input states available in our system, the high repetition rate of about 1 Hz and the family of different unitaries that can be realized using Hamiltonian evolution on the lattice.

We begin with a description of our experimental setup and its features as they relate to the tests and demonstrations performed (Fig. 1). To prepare patterns of identical ⁸⁸Sr atoms in an optical lattice, we load a thermal cloud of atoms into a tweezer array (typically with dimensions of 16 × 24) with random filling of 50–75% (refs. 37,38), and implant these atoms in the lattice^{36,39}. We image the implanted atoms with a typical fidelity of 99.8(1)% per lattice site, and, based on these images, rearrange the atoms into the desired target patterns (Fig. 1a). The rearrangement can be performed with a per-atom success probability as high as 99.5%, but 98% is typical for the data presented in this work (Methods). A second image is used to check if the atoms have been rearranged correctly, allowing us to postselect for perfect atom rearrangement (up to imaging errors). After the atoms are appropriately positioned in the lattice, they are cooled to their three-dimensional (3D) motional ground state by resolved sideband cooling using the narrow line ¹S₀ ↔ ³P₁ transition³⁶ (Methods). Note that similar performance for state preparation and imaging is achieved across a 48 × 48-site region in the lattice (Supplementary Information section VI).

Once prepared, each atom undergoes quantum walk dynamics described by an $m \times m$ single-particle unitary $U = e^{-iHt/\hbar}$, where m is the number of sites in the lattice, H is the lattice Hamiltonian (Methods) and t is the evolution time. During the quantum walk dynamics, 5.0(2)% of

¹JILA, University of Colorado and National Institute of Standards and Technology and Department of Physics, University of Colorado, Boulder, CO, USA. ²National Institute of Standards and Technology, Boulder, CO, USA. ³Department of Physics, University of Colorado, Boulder, CO, USA. ⁴Joint Quantum Institute, University of Maryland Department of Physics and National Institute of Standards and Technology, College Park, MD, USA. ⁵Center for Theory of Quantum Matter, University of Colorado, Boulder, CO, USA. ✉e-mail: aaron.young.w@gmail.com; adam.kaufman@colorado.edu




Fig. 1 | Assembling Fock states of bosonic atoms in a tunnel-coupled optical lattice. **a**, States containing up to 180 bosonic atoms are prepared and measured on-demand in an optical lattice (grey grid) using site- and atom-resolved imaging (pictured images are single-shot experimental data), parallelized rearrangement with optical tweezers (green cones) and high-fidelity laser cooling. The evolution of these atoms in the lattice is described by factorially many (in particle number) different interfering multiparticle trajectories (light and dark blue lines show two possible interfering trajectories for a subset of atoms). **b**, The interference between atoms can be studied by measuring correlations in the occupation of specific sites in the lattice (red). Partitioning the lattice into arbitrary subsets of sites κ (grey) can allow one to more efficiently (in terms of the number of experimental trials) study the effects of interference. Owing to the separability of quantum walk dynamics on an ideal square lattice, it is particularly useful to partition the lattice into columns indexed by their x coordinates (blue). **c**, Specifically, by binning two-dimensional quantum walk data over the y -axis, we can study one-dimensional quantum walks along x , in which the initial coordinate in the y -axis acts as an additional ‘hidden’ DOF that can modify the distinguishability of the atoms in the remaining ‘visible’ x axis. For example, we can prepare two atoms with the same (unlabelled) or different (labelled) y -coordinates (sites initialized with an atom are marked in red). The atom number density on each site after the evolution of these atoms is shown here at an evolution time of $t = t_{\text{HOM}} := 0.96$ ms. **d**, After binning to one dimension, each two-particle output is uniquely labelled by the x coordinates of the two atoms (x_1, x_2), with $x_2 \leq x_1$. The measured probabilities P of observing a given two-particle output are in agreement with theory, in which the unlabelled atoms (top left) agree with the expectation for perfectly indistinguishable bosons (bottom left) and labelled atoms (top right) with distinguishable particles (bottom right). In these and subsequent probability distributions in this work, the input state before the atom evolution is highlighted in red.

the atoms are lost, which we refer to as single-particle loss (Methods). Importantly, single-particle loss is not strongly dependent on evolution time (Supplementary Information section VI) and is primarily because of state preparation errors in which the atoms occupying excited in-plane bands leave the analysis region and are unlikely to return. After the quantum walk, a final image is used to measure the positions of the atoms. This measurement is not number resolving, and instead detects atom number parity on each site because of the effect of light-assisted collisions (we refer to this process as parity projection)^{21,22,40}. However, for most experiments in this work, we can

operate in a regime in which the probability that more than one atom occupies a given site is low, minimizing the effect of parity projection¹. Moreover, except for experiments that involve large atom numbers, we postselect our measurements on observing no lost or extra atoms after the quantum walk dynamics to account for any lingering effects of loss, image infidelity and parity projection (Methods).

Our use of ground-state ⁸⁸Sr atoms ensures that on-site elastic and inelastic interactions are weak^{41,42} (Supplementary Information section VII), and that the many-atom dynamics are, to a good approximation, non-interacting. For n non-interacting bosonic atoms, the many-atom evolution is related to the permanent of an $n \times n$ submatrix of U (Methods). Computing the permanent of an arbitrary matrix is in the complexity class #P-hard⁴³, and so even sampling from the probability distribution resulting from the above evolution is believed to be intractable for more than around 50 atoms using classical techniques⁴⁴. By contrast, for distinguishable atoms the evolution is related to the permanent of an $n \times n$ submatrix of $|U|^2$, where $|\cdot|$ denotes the element-wise norm (Methods). Approximating the permanent of a non-negative real matrix can be performed in polynomial time, and the corresponding sampling task can efficiently be accomplished classically⁴⁵. Different degrees of atom distinguishability result in behaviours that lie in between these two scenarios^{45,46}.

Although our atoms are fundamentally bosonic composite particles, they may not behave bosonically on the lattice. This is because the single-atom state space includes degrees of freedom (DOFs) other than just the location of the atom in the lattice. For example, the atom can be in different electronic states or in different motional states in the direction that is normal to the lattice. The single-atom state space is thus a tensor product of the Hilbert spaces $\mathcal{H}_v$ (‘visible’ DOFs) spanned by $|i\rangle$ (sites) and $\mathcal{H}_h$ (‘hidden’ DOFs) spanned by $|h\rangle$ (labels). In our system, the overall unitary dynamics is non-interacting for the atoms, and the single-atom Hamiltonian acts independently on the hidden and the visible DOFs, which means that it is of the form $H_v \otimes \mathbb{1}_h + \mathbb{1}_v \otimes H_h$ (where the subscripts indicate the subspace that each operator acts on and H is a Hamiltonian acting on the indicated subspace). Owing to this independence, the hidden DOFs can affect the visible behaviour of the atoms only by changing their effective particle statistics. For example, if we prepare some number of atoms in specific visible sites but each atom has a different hidden label, then the visible behaviour of the atoms is that of perfectly distinguishable particles. Other types of visible particle statistic are also possible, including fermionic statistics if the multiparticle wavefunction is antisymmetric in the hidden DOFs (ref. 3). Our tests focus on determining bounds on the deviation from bosonic behaviour because of the hidden DOFs. We also characterize the specific particle statistics exhibited in the visible behaviour in experiments with only two and three particles. Note that errors in the visible DOFs are unlikely because of our procedure for state preparation and postselection (Methods), which rules out certain exotic error models⁴⁷.

The indistinguishability of two atoms is defined to be the Hilbert–Schmidt inner product of the density matrices of the two single-atom wavefunctions on the hidden DOFs. It can be inferred from a Hong–Ou–Mandel (HOM) experiment by comparing the coincidence probability of the atoms P_{H}^B to the corresponding coincidence probability for perfectly distinguishable atoms P_{H}^D (Methods). We measure HOM interference of atoms^{14,18,19} by studying quantum walks in one dimension using the binning procedure described in the Methods (Fig. 2). In these quantum walks, an evolution time of $t_{\text{HOM}} := 0.96$ ms approximates a balanced beam splitter between adjacent sites (Fig. 2b,c and Methods). To measure the behaviour of distinguishable atoms, we can control the hidden (after binning) initial y -coordinates of the atoms by rearrangement (we refer to this as position labelling; see Fig. 1c,d). Alternatively, we can perform a pair of experiments in which only one atom is prepared at a time, and then combine the data in subsequent analysis (we refer to this as time labelling; see Methods). Both approaches agree to within statistical errors.




Fig. 2 | Multiparticle quantum walks in one dimension. **a**, Quantum walks in one dimension, pictured here for two particles at an evolution time of $t = 4.23$ ms for different input states and particle statistics. The atoms are either prepared in neighbouring sites in the lattice (nearest-neighbour, NN) or separated by one site (next-nearest-neighbour, NNN). The theoretical predictions are for perfectly distinguishable particles (Dist.) and indistinguishable bosons (Indist.). The associated measurements either introduce (Labelled) or do not introduce (Unlabelled) a time label to distinguish between the different atoms. In subsequent figures, we use Dist. or Indist. to refer to both data and theory for consistency, with the understanding that the experimental measurements may not correspond to perfectly indistinguishable bosons. **b**, We consider three ways of coarse-graining multiautom distributions: P_{11} refers to the probability of coincident detection of one atom on each of two input sites (black output), P_{bunch} refers to a coincident detection of all atoms on the same site (summing over red outputs) and P_{cloud} refers to all atoms appearing on the same half of the array (summing over blue outputs). Example trajectories that fulfil these conditions are shown on the left. **c**, We plot the ratio P_{11}^D/P_{11}^B , where the superscripts D and B, respectively, refer to experiments with and without introducing an additional label for distinguishability. At an evolution time of $t = t_{\text{HOM}}$ (grey dotted line, callout in **b**), this is analogous to measurements of

HOM interference using a balanced beam splitter. Note that although the HOM dip can have unity visibility for identical bosons in an ideal lattice with only nearest-neighbour tunnelling (grey theory curve), higher order tunnelling terms in our lattice³⁶ result in imperfect visibility even for identical bosons (black theory curve). **d**, Measurements of P_{bunch} and P_{cloud} for two atoms as a function of evolution time, and for both NN and NNN inputs, are in good agreement with theory. **e**, Measurements of P_{bunch} and P_{cloud} for NN input patterns (with up to $n = 5$ and 8 atoms, respectively) are also in good agreement with theory. For comparison, we show a prediction for when each trial of the experiment contains one randomly selected atom that is distinguishable from the rest (1 Dist.). The inset shows the measured atom density in two dimensions with eight prepared atoms, illustrating how at the chosen evolution time of $t = (n-1)t_{\text{HOM}}$ for the measurements in **e**, all n input sites are approximately uniformly coupled to each other. The theory predictions appearing throughout this figure are for error-free preparations of atoms with the indicated particle statistics. The theory curves in **e** are performed using Monte Carlo methods, and thus the shaded regions denote $\pm 1\sigma$ confidence intervals that include both sampling errors and systematic errors relating to fluctuations in the applied unitary evolution (Supplementary Information section II).

Using the above techniques, we measure the contrast of the HOM dip to be $97.1_{-1.5}^{+1.0}\%$, which serves as a lower bound on the atom indistinguishability. We estimate the indistinguishability to be $99.5_{-1.6}^{+0.5}\%$ given additional modelling of our lattice potential (Methods and Supplementary Information section VI). We suspect that the dominant source of distinguishability is imperfect cooling in the direction normal to the lattice (see Methods and Supplementary Information section VIII). We also performed a separate measurement of the indistinguishability that does not rely on either postselection or binning and found that the two measurements were consistent (Supplementary Information section I). Our HOM measurements are averaged over three different regions in the lattice, with similar performance attained across a region that contains all input sites used in this work (Supplementary Information sections VI and IX).

One way to test the visible particle statistics of more than two atoms is to measure full bunching⁴⁸, which is the probability P_{bunch} that all atoms occupy the same visible output site. If the visible and hidden dynamics are perfectly separable, full bunching is uniquely maximized by bosonic particle statistics. Note that to account for the effect of parity projection, we normalize the bunching probability of distinguishable particles by the probability of full survival for bosonic particles (Supplementary Information section III). The results of our full bunching measurements are shown in Fig. 2d,e, and the results are in good agreement with theoretical predictions. Although the lattice Hamiltonian in our experiments is not perfectly separable, this agreement provides strong evidence for the indistinguishability of the atoms.

The full bunching probability decreases very rapidly with atom number and is infeasible for us to measure for more than five atoms. To test the indistinguishability of larger numbers of atoms, it is beneficial to look for signals that are sensitive to interference but converge more quickly. Two quantities of interest are clouding and modified generalized bunching, defined in the Methods. Both of these quantities can be thought of as a way to quantify a general tendency for the atoms to end up on the same site, or (in the case of clouding) close together. Unlike full bunching, interpreting measurements of clouding and modified generalized bunching requires precise knowledge of the atom evolution, namely, of the single-particle unitary U .

For all theory predictions appearing in this work, we perform spectroscopic measurements that use the atoms as local probes of the lattice depth to generate a model of the lattice Hamiltonian and thus of U (Supplementary Information section VI). However, more direct measurements of U are important for future studies that attempt to program U using local, possibly time-varying potentials imposed by the optical tweezers^{13,36}. As a proof of principle, we show that U can be inferred directly by measuring the interference resulting from different preparations of Fock states⁴⁹ (Fig. 3). Specifically, we perform a maximum likelihood fit of the parameters of interest in the unitary using data involving preparations of one and two atoms in the lattice, which results in better precision than previous approaches⁴⁹ (Methods). Using this fitting procedure, we perform characterizations of the terms in U corresponding to four input sites and five output sites at an evolution time of $t = 1.46$ ms (Fig. 3). We found that the maximum likelihood fit is within statistical variation of the spectroscopic characterization of U




Fig. 3 | Characterizing the single-particle unitary. **a**, We directly characterize the evolution of atoms from the input sites circled in red to the output sites marked by the green crosses. This evolution is determined by a 5×4 submatrix M of the single-particle unitary U . **b**, A series of different one- and two-particle quantum walks are used to infer M . Density plots associated with measurements of these different quantum walks are shown in the top panel, with input sites that are populated in a given preparation circled in dark red (unpopulated input sites in light red). M is estimated using the maximum likelihood estimation (MLE) procedure described in the text. Schematically, different preparations provide information about different elements in M , as indicated by the grey lines⁴⁹. The colour scale is shared across all parts of this figure.

(Extended Data Fig. 2e). In principle, these fits allow one to characterize all parameters in U at constant precision with a number of experimental trials that scales polynomially in m . However, the number of parameters in U that we can accurately infer using quantum walk data is currently limited by the cycle time of our experiment and drifts in the experiment that modify U on long timescales.

Given the above calibrations of U , we can compare measurements of clouding with up to $n = 8$ atoms at an evolution time of $t = (n - 1)t_{\text{HOM}}$ with theory (Fig. 2d,e). The behaviour of the atoms is in line with the prediction for ideal bosons and is separated from both measurements with time-labelled (and thus distinguishable) atoms and from theoretical predictions for partially distinguishable atoms.

To go beyond eight atoms, we do not perform binning, and instead study the modified generalized bunching probability $\bar{P}_k$ (Methods) as a function of atom number for square $\sqrt{n} \times \sqrt{n}$ input patterns with next-nearest-neighbour spacing and at a fixed evolution time of $t = 6.45$ ms. Measurements of $\bar{P}_k$ show a clear separation between the distinguishable and bosonic visible behaviours (Fig. 4a). We expect that the main source of distinguishability in our experiment is thermal motional excitations normal to the lattice (Methods and Supplementary Information section VIII). This motional DOF is well-approximated by a harmonic oscillator with motional quantum numbers n_{ax} . Atoms that possess different values of n_{ax} are distinguishable. Our measurements of $\bar{P}_k$ are consistent with a thermal occupation of $\langle n_{\text{ax}} \rangle = 0$, corresponding to the fully indistinguishable case, and inconsistent with significantly higher thermal occupation ($\langle n_{\text{ax}} \rangle \geq 0.167$).

Because of parity projection, $\bar{P}_k$ is closely related to the probability an atom is detected after the quantum walk dynamics and is akin to standard measurements performed in the photonics community using click detectors^{10,12}. It is, therefore, important to calibrate any errors that lead to a modification of the observed atom number, including single-particle loss and certain kinds of detection errors, and include them in our simulations (Methods and Supplementary Information section V). This is in contrast to the measurements in Fig. 2 and the Extended Data, which, due to postselection, are insensitive to the above errors. Despite being sensitive to calibration errors, the agreement between our measurements of $\bar{P}_k$ and low-temperature simulations suggests that the indistinguishability measured using few-particle calibrations is not noticeably degraded when scaling our experiments




Fig. 4 | Interference of large bosonic Fock states. **a**, Measurements of the modified generalized bunching probability $\bar{P}_k$ as a function of atom number for square $\sqrt{n} \times \sqrt{n}$ input patterns at next-nearest-neighbour spacing and at a fixed evolution time of $t = 6.45$ ms with (triangles) and without (circles) introducing time labels that make the atoms fully distinguishable. These measurements can be compared with a model for partial distinguishability (coloured lines), which is described by a harmonic oscillator with an expected thermal occupation of $\langle n_{\text{ax}} \rangle$ (Methods). The fully distinguishable case (dashed line) corresponds to infinite temperature. Note that simulations for low temperatures and large atom numbers are absent because of computational overhead. Insets in **a** and **b** show the atom density after their evolution and share a colour bar. **b**, Measurements of $\bar{P}_k$ with input patterns containing 180 atoms as a function of distinguishability (circles). The distinguishability is controlled by partitioning the input state into n_{labels} sub-ensembles, such that only atoms within a sub-ensemble share a time label and thus can interfere. For particles that are sufficiently distinguishable to simulate, we can compare these measurements with theory (solid line). **c**, The distribution in the observed fraction of surviving atoms on each shot of the experiment is also sensitive to the effects of interference due to parity projection. For the case of nine labels (second from bottom panel), simulations capture the measured distribution of atom survival probabilities. Measurements with different numbers of labels (and thus atom distinguishability) are resolved both from each other and from a simulation of the fully distinguishable case with 180 labels (bottom panel). The vertical lines denote the mean of each distribution (dotted lines are measurements and solid lines are theory). Insets in **c** denote the relevant assignment of labels for each dataset in both **b** and **c**, with each colour corresponding to a unique time label for a subset of input sites. Shaded regions about all theory curves denote $\pm 1\sigma$ confidence intervals, including systematic errors relating to fluctuations in the applied unitary.

up to more particles. This motivates experiments with large ensembles of atoms whose behaviour we are unable to simulate, which we now discuss.

The largest input patterns we study in this work contain 180 atoms (Fig. 4b,c). Although these patterns can be prepared with no defects, in subsequent measurements, we no longer enforce perfect rearrangement to avoid incurring substantial overhead in the number of required experimental trials. However, because we image the atoms after rearrangement and before their evolution, we can identify the locations of any defects. This results in a version of scattershot boson

sampling⁵⁰, but with much less variation in input states than is typical⁷. Based on the few-particle characterizations of atom indistinguishability performed across relevant regions in the lattice (Supplementary Information section VI), we expect the on-demand success rate for preparing a single ground-state atom, evolving it with no loss, and detecting its position in an arbitrary site to be about 92%. The dominant source of deviation from perfect boson sampling is from 5.0(2)% atom loss due to imperfect cooling in the in-plane directions, with further contributions from imaging, rearrangement and distinguishability (Methods and Supplementary Information section VIII). Given these calibrations, we expect that, on an average run of the experiment, around 166 of the 180 input sites are populated with identical bosons that evolve under approximately Haar-random unitary dynamics (Methods) to about 1,015 output sites in the lattice with no loss or detection errors.

The above observations suggest that the experiment is performing a difficult sampling task, but it is not feasible to directly verify that the collected samples are drawn from the correct distribution. To compare with simulations, we control the distinguishability of the atoms by introducing additional time labels. For a distinguishability that is high enough to allow for classical simulation, we find good agreement between theory and measurements of both $\overline{P}_k$ (Fig. 4b) and the distribution of atom survival probabilities (Fig. 4c). We observe the expected qualitative behaviour in which reduced distinguishability leads to an increase in $\overline{P}_k$ and reduced atom survival. The above measurements serve as indirect evidence that the experiment with no additional time labels is behaving in line with expectations for bosonic particle statistics.

It is instructive to compare our approach with pioneering experiments that study boson sampling using photons (Extended Data Table 1). In our experiment, we benefit from low loss (Supplementary Information section VI) that is not strongly dependent on evolution time (equivalently, the depth of the applied linear optical network)⁵¹, high-state preparation and detection fidelity, and many lattice sites (equivalently, many output modes). However, similar to previous large-scale demonstrations of boson sampling^{10–12}, we apply only a restricted family of unitaries. These unitaries possess additional structures that could, in principle, be taken advantage of in efficient classical simulations. Haar-averaging of the unitary would remove the possibility of such simulations and provide access to additional tests of boson sampling that rely on random matrix theory⁵². Based on previous demonstrations, universal control over the single-particle unitary⁵³ can, in principle, be implemented using the same optical tweezers we use for atom rearrangement without introducing additional loss³⁶. However, an important open problem is to improve the efficiency of protocols for directly calibrating the applied unitary for large systems, especially because errors in the applied unitary can have marked consequences on the resulting interference (Methods and Supplementary Information section II).

In this work, we have demonstrated an approach to performing large-scale boson sampling that is enabled by a unique combination of tools for the rapid assembly, evolution and detection of individual atoms in a tunnel-coupled optical lattice, as well as techniques for benchmarking the quality of state preparation and evolution in such a system. In the future, we expect that flexible programmability of the single-particle unitary using optical tweezers will enable stronger certifications of the high-order interference at the core of boson sampling^{47,54}, as well as studies of dynamical phase transitions in sample complexity⁵⁵. More broadly, these tools could be combined with other types of atoms, and with the controllable interactions that are readily available in atomic systems, to rapidly assemble and study interacting Hubbard models^{16,26} for simulations of condensed matter physics¹⁵, to perform tests of complexity in the presence of interactions^{27,56} and to realize new approaches to computing⁵⁷, including in hardware-efficient fermionic architectures⁵⁸.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-024-07304-4>.

1. Aaronson, S. & Arkhipov, A. The computational complexity of linear optics. In *Proc. Forty-Third Annual ACM Symposium on Theory of Computing* 333–342 (Association for Computing Machinery, 2011).
2. Peruzzo, A. et al. Quantum walks of correlated photons. *Science* **329**, 1500–1503 (2010).
3. Sansoni, L. et al. Two-particle bosonic-fermionic quantum walk via integrated photonics. *Phys. Rev. Lett.* **108**, 010502 (2012).
4. Broome, M. A. et al. Photonic boson sampling in a tunable circuit. *Science* **339**, 794–798 (2013).
5. Spring, J. B. et al. Boson sampling on a photonic chip. *Science* **339**, 798–801 (2013).
6. Tillmann, M. et al. Experimental boson sampling. *Nat. Photon.* **7**, 540–544 (2013).
7. Bentivegna, M. et al. Experimental scattershot boson sampling. *Sci. Adv.* **1**, e1400255 (2015).
8. Carolan, J. et al. Universal linear optics. *Science* **349**, 711–716 (2015).
9. Wang, H. et al. Boson sampling with 20 input photons and a 60-mode interferometer in a 10^{14} -dimensional Hilbert space. *Phys. Rev. Lett.* **123**, 250503 (2019).
10. Zhong, H.-S. et al. Quantum computational advantage using photons. *Science* **370**, 1460–1463 (2020).
11. Madsen, L. S. et al. Quantum computational advantage with a programmable photonic processor. *Nature* **606**, 75–81 (2022).
12. Deng, Y.-H. et al. Gaussian boson sampling with pseudo-photon-number-resolving detectors and quantum computational advantage. *Phys. Rev. Lett.* **131**, 150601 (2023).
13. Muralidharan, G., Miyake, A. & Deutsch, I. H. Quantum computational supremacy in the sampling of bosonic random walkers on a one-dimensional lattice. *New J. Phys.* **21**, 055003 (2019).
14. Robens, C. et al. Boson sampling with ultracold atoms. Preprint at arxiv.org/abs/2208.12253 (2022).
15. Lee, P. A., Nagaosa, N. & Wen, X.-G. Doping a Mott insulator: physics of high-temperature superconductivity. *Rev. Mod. Phys.* **78**, 17–85 (2006).
16. Bloch, I., Dalibard, J. & Zwierger, W. Many-body physics with ultracold gases. *Rev. Mod. Phys.* **80**, 885–964 (2008).
17. Cronin, A. D., Schmiedmayer, J. & Pritchard, D. E. Optics and interferometry with atoms and molecules. *Rev. Mod. Phys.* **81**, 1051–1129 (2009).
18. Kaufman, A. M. et al. Two-particle quantum interference in tunnel-coupled optical tweezers. *Science* **345**, 306–309 (2014).
19. Lopes, R. et al. Atomic Hong–Ou–Mandel experiment. *Nature* **520**, 66–68 (2015).
20. Islam, R. et al. Measuring entanglement entropy in a quantum many-body system. *Nature* **528**, 77–83 (2015).
21. Bakr, W. S., Gillen, J. I., Peng, A., Fölling, S. & Greiner, M. A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice. *Nature* **462**, 74–77 (2009).
22. Sherson, J. F. et al. Single-atom-resolved fluorescence imaging of an atomic Mott insulator. *Nature* **467**, 68–72 (2010).
23. Weitenberg, C. et al. Single-spin addressing in an atomic Mott insulator. *Nature* **471**, 319–324 (2011).
24. Murmann, S. et al. Two fermions in a double well: exploring a fundamental building block of the Hubbard model. *Phys. Rev. Lett.* **114**, 080402 (2015).
25. Kaufman, A. M. et al. Quantum thermalization through entanglement in an isolated many-body system. *Science* **353**, 794–800 (2016).
26. Yan, Z. Z. et al. Two-dimensional programmable tweezer arrays of fermions. *Phys. Rev. Lett.* **129**, 123201 (2022).
27. Zheng, Y.-G. et al. Efficiently extracting multi-point correlations of a Floquet thermalized system. Preprint at arxiv.org/abs/2210.08556 (2022).
28. Covey, J. P., Madjarov, I. S., Cooper, A. & Endres, M. 2000-times repeated imaging of strontium atoms in clock-magic tweezer arrays. *Phys. Rev. Lett.* **122**, 173201 (2019).
29. Kaufman, A. M., Lester, B. J. & Regal, C. A. Cooling a single atom in an optical tweezer to its quantum ground state. *Phys. Rev. X* **2**, 041014 (2012).
30. Thompson, J. D., Tiecke, T. G., Zibrov, A. S., Vuletić, V. & Lukin, M. D. Coherence and Raman sideband cooling of a single atom in an optical tweezer. *Phys. Rev. Lett.* **110**, 133001 (2013).
31. Norcia, M. A., Young, A. W. & Kaufman, A. M. Microscopic control and detection of ultracold strontium in optical-tweezer arrays. *Phys. Rev. X* **8**, 041054 (2018).
32. Cooper, A. et al. Alkaline-earth atoms in optical tweezers. *Phys. Rev. X* **8**, 041055 (2018).
33. Endres, M. et al. Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science* **354**, 1024–1027 (2016).
34. Barredo, D., de Léséleuc, S., Lienhard, V., Lahaye, T. & Browaeys, A. An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science* **354**, 1021–1023 (2016).
35. Trisnadi, J., Zhang, M., Weiss, L. & Chin, C. Design and construction of a quantum matter synthesizer. *Rev. Sci. Instrum.* **93**, 083203 (2022).
36. Young, A. W., Eckner, W. J., Schine, N., Childs, A. M. & Kaufman, A. M. Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice. *Science* **377**, 885–889 (2022).
37. Young, A. W. et al. Half-minute-scale atomic coherence and high relative stability in a tweezer clock. *Nature* **588**, 408–413 (2020).
38. Jenkins, A., Lis, J. W., Senoo, A., McGrew, W. F. & Kaufman, A. M. Ytterbium nuclear-spin qubits in an optical tweezer array. *Phys. Rev. X* **12**, 021027 (2022).
39. Schine, N., Young, A. W., Eckner, W. J., Martin, M. J. & Kaufman, A. M. Long-lived Bell states in an array of optical clock qubits. *Nat. Phys.* **18**, 1067–1073 (2022).
40. Schlosser, N., Raymond, G., Protchenko, I. & Grangier, P. Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature* **411**, 1024–1027 (2001).

41. Martínez de Escobar, Y. N. et al. Two-photon photoassociative spectroscopy of ultracold ^{88}Sr . *Phys. Rev. A* **78**, 062708 (2008).
42. Goban, A. et al. Emergence of multi-body interactions in a fermionic lattice clock. *Nature* **563**, 369–373 (2018).
43. Valiant, L. The complexity of computing the permanent. *Theor. Comput. Sci.* **8**, 189–201 (1979).
44. Clifford, P. & Clifford, R. The classical complexity of boson sampling. In *Proc. 2018 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 146–155 (Society for Industrial and Applied Mathematics, 2018).
45. Tichy, M. C. Sampling of partially distinguishable bosons and the relation to the multidimensional permanent. *Phys. Rev. A* **91**, 022316 (2015).
46. Dufour, G., Brünner, T., Rodríguez, A. & Buchleitner, A. Many-body interference in bosonic dynamics. *New J. Phys.* **22**, 103006 (2020).
47. Tichy, M. C., Mayer, K., Buchleitner, A. & Mølmer, K. Stringent and efficient assessment of boson-sampling devices. *Phys. Rev. Lett.* **113**, 020502 (2014).
48. Spagnolo, N. et al. General rules for bosonic bunching in multimode interferometers. *Phys. Rev. Lett.* **111**, 130503 (2013).
49. Laing, A. & O'Brien, J. L. Super-stable tomography of any linear optical device. Preprint at arxiv.org/abs/1208.2868 (2012).
50. Lund, A. P. et al. Boson sampling from a Gaussian state. *Phys. Rev. Lett.* **113**, 100502 (2014).
51. García-Patrón, R., Renema, J. J. & Shchesnovich, V. Simulating boson sampling in lossy architectures. *Quantum* **3**, 169 (2019).
52. Boixo, S. et al. Characterizing quantum supremacy in near-term devices. *Nat. Phys.* **14**, 595–600 (2018).
53. Reck, M., Zeilinger, A., Bernstein, H. J. & Bertani, P. Experimental realization of any discrete unitary operator. *Phys. Rev. Lett.* **73**, 58–61 (1994).
54. Dittel, C. et al. Totally destructive many-particle interference. *Phys. Rev. Lett.* **120**, 240404 (2018).
55. Deshpande, A., Fefferman, B., Tran, M. C., Foss-Feig, M. & Gorshkov, A. V. Dynamical phase transitions in sampling complexity. *Phys. Rev. Lett.* **121**, 030501 (2018).
56. Maskara, N. et al. Complexity phase diagram for interacting and long-range bosonic Hamiltonians. *Phys. Rev. Lett.* **129**, 150604 (2022).
57. Childs, A. M., Gosset, D. & Webb, Z. Universal computation by multiparticle quantum walk. *Science* **339**, 791–794 (2013).
58. González-Cuadra, D. et al. Fermionic quantum processing with programmable neutral atom arrays. *Proc. Natl Acad. Sci. USA* **120**, e2304294120 (2023).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2024

Methods

State preparation

State preparation in our experiment involves atom rearrangement and high-fidelity optical cooling. The rearrangement must balance several conflicting requirements associated with working in a tightly spaced optical lattice that is compatible with strong tunnel coupling. In particular, our tweezers have a beam waist of 480(20) nm in comparison to the lattice spacing of about $813/\sqrt{2}$ nm. This means that tweezers directed at a specific lattice site still have substantial overlap with adjacent sites, and so atoms trapped in those adjacent sites can experience higher loss rates, especially when the tweezers are moving. Moreover, overlapping tweezers will beat against each other because of our use of crossed acousto-optic deflectors to project the tweezer array³¹. Tweezers on adjacent lattice sites are separated in frequency by 1.95 MHz, far higher than the relevant trap frequencies of 50–200 kHz, and so one-dimensional (1D) tweezer arrays with equal spacing to the lattice do not cause much heating. However, because the lattice is slightly rectangular, placing a tweezer on every lattice site in two dimensions can result in lower beat frequencies that cause increased heating. Our algorithm for atom rearrangement seeks to balance these concerns against the desire to minimize the total distance travelled by the atoms and the total duration of the rearrangement procedure.

The basic premise is to perform most operations in parallel with 1D arrays of tweezers^{59,60} and to perform as many operations as possible on a sublattice that does not lead to undesirable beat frequencies between overlapping tweezers. We choose to work with either nearest- or next-nearest-neighbour spacing along one axis and third-nearest-neighbour spacing along the other axis. Atoms are stochastically loaded into tweezers on this sublattice with 50–75% filling (depending on whether or not enhanced loading is used³⁸) and implanted into the lattice³⁶. Rearrangement proceeds in four stages (Extended Data Fig. 1a):

1. Pre-sorting: If a column is loaded with more atoms than are required in the target pattern, any excess atoms are pushed to the neighbouring column in a single step using a 1D tweezer array. Similarly, if the column has too few atoms, the missing atoms are pulled from the adjacent column. In the rare case that this step fails, the experiment is terminated and we simply load a new ensemble of atoms.
2. 1D rearrangement: Each column is then rearranged in one dimension, with excess atoms pushed to the edges of the array.
3. Filtering: A subset of atoms is transferred back into a two-dimensional (2D) tweezer array that addresses only the correct sublattice, and the lattice potential is extinguished. Laser light that addresses the $^1S_0 \leftrightarrow ^3P_1$ transition is applied to resonantly blow away the excess atoms while optically cooling the tweezer-trapped atoms. The remaining atoms are subsequently transferred back into the lattice.
4. Compression: Optionally, the columns of rearranged atoms can be translated closer together one column at a time.

The number of steps in the resulting algorithm scales as $O(\sqrt{n})$, where n is the number of atoms in the target pattern. For target patterns with similar density to the loaded sublattice, the mean distance travelled by each atom is $O(1)$, and the runtime of the algorithm typically scales as $O(\sqrt{n})$. For dense target patterns, the worst-case scaling of the mean distance travelled is $O(\sqrt{n})$ leading to a runtime of $O(n)$.

Throughout rearrangement, the lattices are held at a constant depth of $U_{ax}/\hbar = 2\pi \times 2.4$ MHz for the axial lattice, and $U_{2D}/\hbar = 2\pi \times 1.7$ MHz for the 2D lattice. Resolved sideband cooling is continuously applied to the lattice-trapped atoms. To move a given set of atoms, tweezers are ramped on and off over 60 μ s to a depth of $U_{tw}/\hbar \approx 2\pi \times 30$ MHz and moved at a constant speed of 26 μ m ms⁻¹. Rearrangement occurs under magic conditions in the lattice for the $^1S_0 \leftrightarrow ^3P_1$ transition³⁶, and so the tweezer-trapped atoms are shifted out of resonance from the cooling light while they are being translated. Because the appropriate

moves in our algorithm can be computed efficiently, and because of the architecture of our tweezer control system³⁷, analysing the atom images and programming the control system takes less than 5 ms for target patterns containing up to 270 atoms. However, in our current experimental system, there is a technical delay of 110 ms between taking the image of the stochastically loaded atoms and initiating rearrangement because of the time it takes to extract the image data from the camera. This could be addressed in the future by performing image processing on the same field-programmable gate array hardware used for atom rearrangement⁶¹.

For the 180 atom patterns in this work, rearrangement is performed in about 30 ms. The per-atom success rate for this rearrangement can be as high as 99.5% and is primarily limited by imaging fidelity and loss. However, owing to drifts in the experiment, the typical per-atom success rate is 98% for the experiments in this work.

Once the success of this rearrangement is verified using a second image, the atoms are optically cooled to near their 3D motional ground state. This cooling is composed of 120 pulses that alternate between cooling pulses with a duration of 200 μ s on two nearly-orthogonal radial axes in the plane of the 2D lattice separated by 400 μ s cooling pulses on the axial out-of-plane axis. Each cooling pulse is separated by a delay of 200 μ s leading to an overall cooling sequence that is 60 ms in duration. This sequence is optimized for high-fidelity cooling of the axial direction at the cost of slightly worse radial cooling. We choose to make this trade-off because we are able to postselect for perfect cooling in the radial directions. Specifically, for our experimental parameters, atoms that occupy higher bands in the 2D lattice are lost during tunnelling³⁶. On the basis of a master equation calculation (Extended Data Fig. 1b), we expect this cooling to result in a 3D motional ground-state occupation of 96.97%, which is consistent with the expectation based on an approximate analytical calculation⁶². The expected axial motional ground-state occupation of 99.58% leads to an expected indistinguishability of 99.27%, in agreement with our experimental estimate of $99.5^{+0.5}_{-1.6}\%$. The expected combined ground-state occupation of 97.37% in the radial directions leads to 2.63% loss, which explains a large fraction of our average observed loss of 5.0(2)%. The remaining loss is probably because of the imperfect adiabaticity of ramps in the lattice depth as well as the imperfect overlap between the Wannier functions of the atoms in the lattice before and after the lattice is quenched to the conditions used for tunnelling. Note that because both cooling and imaging involve the spontaneous emission of photons from atoms occupying different lattice sites, our state preparation is expected to be perfectly dephasing. This rules out error models for boson sampling that rely on coherences between different input modes, like the mean field sampler⁴⁷.

Image analysis and feedback

Similar to previous works, our image analysis involves multiplying raw images by a discrete set of masks corresponding to the known locations of sites in the lattice and thresholding the results to identify the presence or absence of an atom in a given site³⁶. However, to improve the fidelity of this procedure, we optimized the applied masks by training a single-layer neural network on simulated data corresponding to 30% random filling of the lattice. The resulting masks are similar to the ones we previously used based on the measured point spread function of our imaging system but with a slight negative bias on adjacent lattice sites that reduces errors due to leakage light from one lattice site to another. Although a deep neural network can result in better performance⁶³, these performance gains are marginal in our setup because of the already high imaging resolution in comparison with the lattice spacing. This being the case, we opt to use a single-layer neural network to ease the characterization of errors and image fidelity.

The resulting analysis yields a combined imaging loss and infidelity of $P_{10} = 0.002(1)$ for a calibration dataset, however, P_{10} can fluctuate day to day by around 0.002. P_{10} can be interpreted as the probability of a false negative—that is, a site that contains an atom is incorrectly

Article

identified as being empty. The probability of a false positive (where an empty site is incorrectly identified as containing an atom) is much lower, with a value of $P_{01} \approx 10^{-5}$. For experiments with very low density (for example, when one atom occupies an analysis region containing about 1,015 sites), the effect of false positives can become comparable to or larger than the effect of false negatives when trying to correctly identify the presence and position of the atom after the quantum walk dynamics.

All imaging errors are quadratically suppressed when postselecting for perfect rearrangement and no lost or extra observed atoms, and so measurements involving this postselection are negligibly affected by imaging errors. For experiments in which we do not postselect on the number of atoms remaining after the quantum walk dynamics (namely, the experiments shown in Fig. 4), imaging errors can play a part. Specifically, for the atom and mode numbers relevant to Fig. 4, the combined effects of false positives and negatives lead to overall error rates of approximately 1% in our ability to correctly identify the presence and location of a given atom. In experiments without postselection on the final atom number, we calibrate the effect of these imaging errors and include them in our simulations (Supplementary Information section V).

An important consideration in our experiment is the possibility that the position of the lattice drifts relative to the tweezer array and to the masks used in image analysis, which can occur on a timescale of about 30 min. To correct this, we use the images taken before and after rearrangement (but not after the atoms are allowed to propagate through the lattice) to identify the positions of the lattice-trapped atoms relative to the imaging system, and thus also relative to the optical tweezers. This information can be used to correct any drifts by adjusting the tweezer and mask positions on subsequent runs of the experiment. Because the repetition rate of the experiment is about 1 Hz, these corrections can be made much faster than drifts can occur. The tweezer positions can drift relative to the imaging system on a timescale of several hours, but this is readily corrected by taking images of atoms trapped in the tweezers³¹ and adjusting their positions to match the lattice.

For datasets in which we simulate n partially or fully distinguishable particles by combining multiple runs with fewer than n particles, we compensate for parity projection during imaging by summing the resulting processed images and taking the result mod(2).

Computing atom indistinguishability

The goal of our HOM experiments is to measure the indistinguishability of a pair of atoms. To do so, we measure the probability that two nominally indistinguishable particles arrive in disjoint subsets of sites and compare this with the corresponding probability for distinguishable particles. We now show that the ratio of these two quantities gives a lower bound on the indistinguishability of the atoms.

Our model is that we have linear optical evolution on atoms with an extra DOF that evolves independently of the visible DOF. Then, the probability that two partially distinguishable atoms start in the sites k, l with $k \neq l$ and end in sites i, j is

$$P^{\text{partial}}(ij|kl) = |U_{i,k}|^2 |U_{j,l}|^2 + |U_{i,l}|^2 |U_{j,k}|^2 + 2\mathcal{J} \text{Re}(U_{i,k} U_{j,l} U_{j,k}^* U_{i,l}^*) \quad (1)$$

$$\begin{aligned} &\text{if } i \neq j \\ P^{\text{partial}}(ii|kl) &= (1 + \mathcal{J}) |U_{i,k}|^2 |U_{i,l}|^2 \\ &\text{otherwise} \end{aligned} \quad (2)$$

where the indistinguishability $\mathcal{J}$ is the Hilbert–Schmidt inner product of the two single-particle density matrices on the hidden DOFs and U is the single-particle unitary. Correspondingly, distinguishable atoms have the distribution

$$P^{\text{dist}}(ij|kl) = |U_{i,k}|^2 |U_{j,l}|^2 + |U_{i,l}|^2 |U_{j,k}|^2 \quad \text{if } i \neq j \quad (3)$$

$$P^{\text{dist}}(ii|kl) = |U_{i,k}|^2 |U_{i,l}|^2 \quad \text{otherwise} \quad (4)$$

In the experiment, we expect that the probability of coincidence of atoms on any two particular sites is small, and therefore estimation of this probability is difficult. Thus, it is useful to be able to bundle many sites together.

So, let S_1 and S_2 be disjoint sets of sites. The total probability that we prepare partially distinguishable atoms and see one atom each in S_1, S_2 is

$$P_{S_1, S_2}^{\text{partial}} = \sum_{i \in S_1} \sum_{j \in S_2} P^{\text{partial}}(ij|kl) \quad (5)$$

and the corresponding probability for distinguishable atoms is

$$P_{S_1, S_2}^{\text{dist}} = \sum_{i \in S_1} \sum_{j \in S_2} P^{\text{dist}}(ij|kl) \quad (6)$$

so their ratio is

$$Q_{S_1, S_2}^{\text{HOM}} := 1 + \mathcal{J} \frac{\sum_{i \in S_1} \sum_{j \in S_2} \text{Re}(U_{i,k} U_{j,l} U_{j,k}^* U_{i,l}^*)}{\sum_{i \in S_1} \sum_{j \in S_2} (|U_{i,k}|^2 |U_{j,l}|^2 + |U_{i,l}|^2 |U_{j,k}|^2)} \quad (7)$$

$\underbrace{\hspace{10em}}_{:= -\tau(S_1, S_2)}$

Note that the ratio $-\tau(S_1, S_2) \geq -1$. We thus have the inequality

$$Q_{S_1, S_2}^{\text{HOM}} \geq 1 - \mathcal{J} \quad (8)$$

When equality holds so $\tau(S_1, S_2) = 1$, we say that S_1 and S_2 satisfy the balanced condition, in analogy with the case of a balanced beam splitter.

HOM with loss and parity projection. In our experiment, we are interested in the case that $S_1 = C_k$ and $S_2 = C_l$, where C_k is the column of lattice sites that includes site k , and similarly for C_l . It remains to construct an estimator of $Q_{C_k, C_l}^{\text{HOM}}$. We will construct it from separate estimates of $P_{C_k, C_l}^{\text{dist}}$ and $P_{C_k, C_l}^{\text{partial}}$.

To estimate $P_{C_k, C_l}^{\text{dist}}$ (called P_{11}^{D} in the main text), we can directly use the single-particle data. Denoting the event that the particle is not lost as $\neg\lambda$, we can write the coincidence probability for distinguishable particles as

$$P_{C_k, C_l}^{\text{dist}} = P(C_k|k, \neg\lambda) P(C_l|l, \neg\lambda) + P(C_k|l, \neg\lambda) P(C_l|k, \neg\lambda) \quad (9)$$

As the right-hand side is a multilinear polynomial of single-particle probabilities, the plug-in estimator is unbiased.

Now we construct an estimator of $P_{C_k, C_l}^{\text{partial}}$ (called P_{11}^{B} in the main text). The probability that the prepared nominally indistinguishable particles end up on the same columns as the ones they start on is

$$P_{C_k, C_l}^{\text{obs}} = P(v_2|kl) P_{C_k, C_l}^{\text{partial}} \quad (10)$$

where the prepared sites are named k and l , and the event v_2 is that neither of the two particles is lost because of a single-particle loss event. Owing to the effect of parity projection, we cannot measure $P(v_2|kl)$ without further assumptions. In particular, we assume that the loss acts independently and identically on each site. Then defining β to be the event that we observe one particle in the output, the probability of β occurring is

$$P(\beta|kl) = 2(1 - P_\lambda) P_\lambda \quad (11)$$

where P_λ is the probability that a single particle is lost because of a single-particle loss event. So we can solve for the loss probability,

$$P_\lambda = \frac{1 - \sqrt{1 - 2P(\beta|kl)}}{2} \quad (12)$$

and compute the probability that neither particle is lost because of a single-particle loss event:

$$P(v_2|kl) = (1 - P_\lambda)^2 \quad (13)$$

Then we can construct a plug-in estimate of $Q_{C_k, C_l}^{\text{HOM}}$ from those of $P_{C_k, C_l}^{\text{partial}}$ and $P_{C_k, C_l}^{\text{dist}}$, and apply the delta method⁶⁴ to obtain a first-order unbiased estimate of $Q_{C_k, C_l}^{\text{HOM}}$. To construct a confidence interval, we construct 1,000 bootstrap estimates of $Q_{C_k, C_l}^{\text{HOM}}$ and apply the bias-corrected percentile method⁶⁵.

From a calibration of the unitary (Supplementary Information section VI), we can calculate $\tau(C_k, C_l)$. Then we can extract the indistinguishability from

$$\mathcal{J} = \frac{1}{\tau(C_k, C_l)} (1 - Q_{C_k, C_l}^{\text{HOM}}) \quad (14)$$

We expect that our uncertainty in our calibration in τ is negligible compared with the statistical fluctuation in our estimate of $Q_{C_k, C_l}^{\text{HOM}}$ (Supplementary Information section II), so we ignored this effect when calculating a confidence interval for $\mathcal{J}$. The confidence interval was constructed through the bias-corrected percentile method, using 1,000 bootstraps. When constructing our point estimate and our confidence interval, the estimate may go above 1 because $\tau \leq 1$. To account for this, after constructing the bootstrap confidence interval, we clip the upper end to be equal to 1 if it is larger than 1. Similarly, we apply the same procedure to the point estimate.

Binning

There is flexibility in how to partition between visible and hidden DOFs (Fig. 1b), which we take advantage of in tests of indistinguishability involving up to eight atoms. In an ideal square lattice, H takes the form $H_x \otimes \mathbb{1}_y + \mathbb{1}_x \otimes H_y$, where x and y denote the spatial coordinates of the lattice sites, and in this case, we can consider one of the spatial coordinates as hidden. For example, if we ignore the y -coordinate of the atoms, the resulting visible behaviour is that of a 1D multiparticle quantum walk along x (Fig. 1c,d). We refer to this as binning, because this is accomplished by summing the observed numbers of atoms along the y -axis in the final image. Binning is convenient because it allows us to operate in a regime in which it is rare that there is more than one atom on a given lattice site, minimizing the effect of parity projection and providing effectively number-resolved measurements of the visible sites. Note that the above form of H is weakly violated in our lattice because of higher order tunnelling terms (for example, diagonal tunnelling) and small non-factorizable variations in the nearest-neighbour tunnelling terms. These violations are accounted for in our simulations, and do not markedly affect our results.

Clouding

The clouding probability is defined as the probability that all atoms end up on the same half of the array for a 1D quantum walk⁶⁶ (Fig. 2b,d,e). Similar to bunching, we normalize the clouding probability of distinguishable particles by the probability of full survival for bosonic particles to account for the effect of parity projection (Supplementary Information section III). Note that, unlike full bunching, the enhancement of clouding for identical bosons in comparison to distinguishable particles is strongly dependent on both the system evolution and the specific input state. For example, preparing atoms at next-nearest-neighbour spacing in the lattice can make the difference

in clouding for bosonic and distinguishable atoms almost zero, as confirmed by measurements of two and three atoms (Fig. 2a,d and Extended Data Fig. 3).

Generalized bunching

The generalized bunching probability P_κ that all n atoms appear in an arbitrary subset κ of sites can serve as a useful quantity for benchmarking the performance of a boson sampler^{67,68}. For most unitaries, including those applied in this work, numerical calculations indicate that P_κ is maximized by bosonic particle statistics in comparison with other particle statistics. This is not true in certain fine-tuned cases^{69,70}, which would be interesting to explore in future experiments with more control over the applied unitary.

The above generalization of bunching helps to unify our earlier measurements, in which binning columns of sites and measuring coincidences, full bunching or clouding simply correspond to specific choices of κ (Fig. 1b). For appropriate selections of the size of κ , the difference in generalized bunching probabilities of bosonic behaviour in comparison with other behaviours is expected to converge in a number of measurements that is polynomial in n (refs. 67,71). Specifically, we choose $|\kappa| = \lfloor m - m/n \rfloor$ (where $\lfloor \cdot \rfloor$ denotes rounding to the nearest integer), with $m = 500$ for Fig. 4a and $m = 1,015$ for Fig. 4b.

We cannot directly measure P_κ , and instead we measure the probability P'_κ that all observed atoms on a given run of the experiment appear in the set of sites κ . P'_κ differs from P_κ because even-numbered occupancy of a site not contained in κ contributes to a successful event in which all remaining atoms after parity projection appear within κ . Although we do not claim that P'_κ is maximized by perfectly bosonic visible behaviour for any input state, κ and U , we find in numerics that P'_κ still serves as a useful observable: it converges in a reasonable number of measurements and distinguishes between a family of experimentally relevant models for the hidden DOFs (discussed in the main text). To avoid concerns of biasing in our choice of κ , we average P'_κ over all choices of a given size $k := |\kappa|$ to compute the quantity $\overline{P'_\kappa}$, which we refer to as the modified generalized bunching probability in the main text. Note that the main contribution to $\overline{P'_\kappa}$ is bunching and the resulting loss as a result of parity projection. As a result, we do not postselect on the number of surviving atoms after the quantum walk dynamics in our measurements of $\overline{P'_\kappa}$.

Although there are many selections of κ that must be averaged over, $\overline{P'_\kappa}$ can be estimated efficiently from the observations using the following combinatorial argument: Let $[m] = \{1, \dots, m\}$ be shorthand for the set of output sites. Let $S \subseteq_k [m]$ denote a subset of sites of size k . Let G be the random variable denoting the site occupation of the output, with g being a single sample of that random variable. Then, the average probability that all n particles arrive in a set of size k is

$$\overline{P'_\kappa} = \binom{m}{k}^{-1} \sum_{S \subseteq_k [m]} \Pr(G \subseteq S) \quad (15)$$

$$= \binom{m}{k}^{-1} \sum_g \Pr(g) \sum_{S \subseteq_k [m]} \mathbb{I}(g \subseteq S) \quad (16)$$

$$= \binom{m}{k}^{-1} \sum_g \Pr(g) \binom{m - \#(g)}{k - \#(g)} \quad (17)$$

where $\mathbb{I}$ is the indicator function that is 1 when its argument is true, and 0 otherwise, and $\#(g)$ is the number of nonzero entries of g , and the sum over g ranges over all possible mode occupations of the output. As the expression in equation (17) is a linear combination of probabilities of outcomes, we can weight the corresponding frequencies by the coefficients appearing in the sum to estimate the quantity $\overline{P'_\kappa}$.

Quantum walk dynamics

The single-particle Hamiltonian governing the evolution of atoms in the lattice is

$$H = - \sum_{\langle i,j \rangle} J_{ij} (|i\rangle\langle j| + |j\rangle\langle i|) - \sum_i V_i |i\rangle\langle i|, \quad (18)$$

where $|i\rangle$ denotes the occupation of the lattice site i , and $\langle i,j \rangle$ denotes all pairs of sites in the lattice. V_i denotes a position-dependent potential that captures the harmonic confinement imposed by the lattice beams and, in principle, can be adjusted using the optical tweezers³⁶. J_{ij} is the strength of the tunnel coupling between sites i and j . Nearest-neighbour tunnelling dominates with an energy of $J_{ij}/\hbar \simeq J/\hbar = 2\pi \times 119$ Hz; however, small contributions from diagonal and next-nearest-neighbour tunnelling are also included in our simulations³⁶ (Supplementary Information section VI). Evolution under H results in a quantum walk described by an $m \times m$ single-particle unitary $U = e^{-iHt/\hbar}$, where m is the number of sites in the lattice.

As mentioned in the text, some loss occurs during the evolution as a result of imperfect state preparation, but we do not observe this loss to depend on t (although we expect additional loss effects to play a part on timescales that are very long in comparison to the dynamics; Supplementary Information section VI). This contrasts with previous demonstrations of boson sampling in which the survival probability decays exponentially with evolution time (or equivalently the depth of the linear optical circuit), which can be exploited in classical simulation methods^{51,72–74}.

For many of our measurements with 1D nearest-neighbour input patterns, we use an evolution time of $t = (n - 1)t_{\text{HOM}}$ such that all input sites are approximately uniformly coupled to each other after the quantum walk dynamics¹³. At an evolution time of 6.45 ms, as applies to the data shown in Fig. 4, U does not appreciably couple all 180 input sites to all 1,015 output sites. This is primarily because of the finite size of the lattice beams and resulting harmonic confinement (Extended Data Fig. 2a,b and Supplementary Information section VI). This harmonic confinement is negligible in a 15×15 site region near the centre of the lattice, in which the distribution of the norm-square of the elements in U is well-captured by the Porter–Thomas distribution for 385 outputs, indicating behaviour that shares features with a Haar-random unitary (Extended Data Fig. 2c,d). Over the full 1,015-site region considered in this work, U couples each site to an average of 83 of the 180 input sites with an amplitude of at least 10^{-3} . Some sites are coupled by at least 10^{-3} to as many as 156 input sites.

From single- to many-particle dynamics

Adopting the first-quantized treatment typically used in discussions of boson sampling¹, consider n indistinguishable, bosonic and non-interacting atoms occupying input sites $\mathbf{j} = (j_1, \dots, j_n)$, with $j_1 \leq \dots \leq j_n$, where $j_i \in \{1, \dots, m\}$, and m is the total number of available sites in the lattice. Each atom undergoes a quantum walk described by an $m \times m$ single-particle unitary $U = e^{-iHt/\hbar}$, where t is the evolution time. To calculate the probability $P^B(\mathbf{k}|\mathbf{j}, U)$ of observing a specific output pattern $\mathbf{k} = (k_1, \dots, k_n)$ after this evolution, we must sum over all permutations of atom labels on the input going to atom labels on the output. This treatment is equivalent to properly symmetrizing the state and evolution associated with these atoms, and yields¹

$$P^B(\mathbf{k}|\mathbf{j}, U) = \frac{1}{\mathbf{k}!} |\text{Perm}(U_{\mathbf{k},\mathbf{j}})|^2 \quad (19)$$

where $U_{\mathbf{k},\mathbf{j}}$ is the $n \times n$ submatrix of U that contains only the rows corresponding to sites $\mathbf{k}$ (including any duplicates) and columns corresponding to $\mathbf{j}$ (which contains no duplicates in our experiments because the lattice is initialized with at most one atom per site). The

normalization constant $\mathbf{k}!$ is more conveniently expressed in the site occupation basis, $\mathbf{k}' = (k'_1, \dots, k'_m)$, where k'_i counts the number of atoms occupying a given site $i \in \{1, \dots, m\}$. Then we define $\mathbf{k}! = \prod_{i=1}^m k'_i!$. The equivalent calculation for distinguishable atoms yields $P^D(\mathbf{k}|\mathbf{j}, U) = \text{Perm}(|U_{\mathbf{k},\mathbf{j}}|^2)/\mathbf{k}!$. Note that because the single-particle distributions are given by $|U_{\mathbf{k},\mathbf{j}}|^2$, we can simulate the distinguishable situation by combining separate single-particle measurements (as applies in the case of time-labelled distinguishable data throughout this work).

Simulations

For up to three particles, the simulations in this work involve exactly solving for the full output distribution by evaluating the permanent in equation (19) (or the corresponding expression for distinguishable particles) for all possible outputs using Glynn's formula. For larger particle numbers, we follow the approach in ref. 44 to sparsely sample from the full output distribution.

The simulations involving thermal occupation of the motional DOF normal to the lattice in Fig. 4a assume that the evolution of this hidden DOF is independent of the visible evolution of the atoms and dephased by our state preparation. As a result, the evolution of atoms with a mixed motional state can be simulated by assigning a specific motional state drawn from the appropriate thermal distribution to each atom in a given simulated sample. For each subset of atoms that share a motional state, we draw a sample using the approach in ref. 44 and combine these samples into a single sample of partially distinguishable atoms. The effect of loss and detection errors are simulated incoherently and applied after these samples are generated, as is parity projection (Supplementary Information section V). The above simulations indicate that the thermal model for distinguishability leads to expected measurements of modified generalized bunching that monotonically interpolate from the distinguishable to the indistinguishable bosonic case as the temperature is reduced.

Characterizing the single-particle unitary

As discussed in the main text, we characterize the dynamics of the atoms in two ways. The first is the spectroscopic characterization (Supplementary Information section VI) and the second is a maximum likelihood procedure to fit the single-particle unitary that determines the dynamics, from one- and two-particle data. Here we describe the maximum likelihood procedure and our method of determining its performance.

First, we describe the model used to compute the likelihoods. Our model describes atoms that are subjected to single-particle loss, undergo tunnelling dynamics and then are measured by parity projection. As we are only inferring some of the parameters of the unitary, it suffices to use a restricted model that uses only the entries of the single-particle unitary U that describes scattering from the input sites I (highlighted in red in Fig. 3) to the output sites S (green crosses in Fig. 3). The model for the single-particle distribution is

$$P_{U,P_\lambda}(s|i) = (1 - P_\lambda) |U_{si}|^2 \quad \text{if } s \in S \quad (20)$$

$$P_{U,P_\lambda}(\tau|i) = (1 - P_\lambda) \left(1 - \sum_{s \in S} |U_{si}|^2 \right) \quad (21)$$

$$P_{U,P_\lambda}(\emptyset|i) = P_\lambda. \quad (22)$$

Here τ is the event that the atom arrived outside of S , $\emptyset$ is the event that the particle was lost, $i \in I$ is the initial site, P_λ is a parameter describing the single-particle loss, and U_{si} is the parameter describing the amplitude for one particle to start in i and end in s . The model for the two-particle distribution is

$$P_{U,P_\lambda,\mathcal{J}}(s,s'|i,j) = (1 - P_\lambda)^2 P_{U,\mathcal{J}}^{\text{partial}}(s,s'|i,j) \quad (23)$$

for $\{s,s'\} \in P_2(S)$

$$P_{U,P_\lambda,\mathcal{J}}(s|i,j) = P_\lambda(1 - P_\lambda)(|U_{s,i}|^2 + |U_{s,j}|^2) \quad (24)$$

for $s \in S$

$$P_{U,P_\lambda,\mathcal{J}}(\zeta|i,j) = 1 - \sum_{\{s,s'\} \in P_2(S)} P(s,s'|i,j) - \sum_{s \in S} P(s|i,j) \quad (25)$$

Here the set $P_2(S)$ is of sets of pairs of elements of S , the event ζ is the event that it was not the case that all surviving particles arrived in S and $i,j \in I$ are the initial sites. Finally, the probability that lossless, partially distinguishable atoms start in sites i,j and arrive at distinct sites $s,s' \in S$ is

$$P_{U,\mathcal{J}}^{\text{partial}}(s,s'|i,j) = |U_{s,i}|^2 |U_{s',j}|^2 + |U_{s,j}|^2 |U_{s',i}|^2 \quad (26)$$

$$+ 2\mathcal{J} \text{Re}(U_{s,i} U_{s',j} U_{s,j}^* U_{s',i}^*) \quad (27)$$

where $\mathcal{J}$ is the indistinguishability of the two atoms.

The parameters P_λ and U are the parameters that we wish to infer, whereas the parameter $\mathcal{J}$ is obtained from separate calibration data, as discussed in the section ‘Computing atom indistinguishability’. To simplify the inference procedure, we first infer P_λ using only the single-particle data, then with P_λ fixed, we run the quasi-Newton L-BFGS optimizer as implemented in PYTORCH to maximize the log-likelihood of the data with respect to U . It remains to specify a parameterization of U .

Let M be the submatrix of U that is the intersection of the columns specified by I and the rows specified by S . Then as M is only $|S| \times |I|$, we wish to find a parameterization of it that does not require too many more parameters, while at the same time respecting the constraint that it is a submatrix of a unitary. To accomplish this, we specify M as the $|S| \times |I|$ top-left submatrix of a $(|S| + |I|) \times (|S| + |I|)$ unitary matrix V . We parameterize the matrix V by specifying its anti-Hermitian matrix logarithm iH , so $V = e^{iH}$. We parameterize H by writing it as a (real) linear combination of the generalized Gell–Mann matrices. To define the $d \times d$ Gell–Mann matrices, first let $k, l, q, j \in \{1, \dots, d\}$. Then the k, l th Gell–Mann matrix B_{kl} is defined by

$$(B_{kl})_{qj} = \frac{1}{\sqrt{2}} (\delta_{qk} \delta_{jl} + \delta_{ql} \delta_{jk}) \quad \text{if } k < l \quad (28)$$

$$(B_{kl})_{qj} = \frac{1}{\sqrt{2}} (i\delta_{qk} \delta_{jl} - i\delta_{ql} \delta_{jk}) \quad \text{if } k > l \quad (29)$$

and for $k = l < d$, we have

$$(B_{kk})_{qj} = \frac{1}{\sqrt{k(k+1)}} \delta_{qj} \quad \text{for } q, j \leq k \quad (30)$$

$$(B_{kk})_{(k+1),(k+1)} = \frac{-k}{\sqrt{k(k+1)}} \quad (31)$$

$$(B_{kk})_{qj} = 0 \quad \text{else} \quad (32)$$

and finally, we have

$$B_{d,d} = \frac{1}{\sqrt{d}} \mathbb{I}_d \quad (33)$$

The Gell–Mann matrices $\{B_{ij}\}$ are a basis of Hermitian matrices that are orthonormal with respect to the Hilbert–Schmidt inner product.

Thus, we can specify H in terms of its coefficients c_{ij} on the basis of $(|S| + |I|) \times (|S| + |I|)$ Gell–Mann matrices, so $H = \sum_{ij} c_{ij} B_{ij}$. This gives us a $(|S| + |I|)^2$ dimensional parameter space without a boundary.

To specify the initial point of the algorithm, we start with a model M_0 of M as computed from the spectroscopic calibration (Supplementary Information section VI), compute an isometric completion W_0 of it, expressed in block form by $W_0^\dagger = (M_0, \sqrt{1 - M_0 M_0^\dagger})$, then compute a unitary completion $V_0^\dagger$ of W_0 by appending an orthonormal basis of the nullspace of $W_0 W_0^\dagger$ as columns. Then, if the eigenvalues of V_0 are $e^{i\phi_k}$, where each $\phi_k \in (-\pi, \pi]$ for $k \in \{1, \dots, |I| + |S|\}$, and Q diagonalizes V_0 , we construct an initial Hermitian matrix H_0 from $H_0 = Q \text{diag}(\phi_1, \dots, \phi_{n+m}) Q^\dagger$. As the Gell–Mann matrices are orthonormal, we can then extract the initial coefficients $(c_0)_{ij}$ from $(c_0)_{ij} = \text{Tr}(H_0 B_{ij})$.

Having specified the initial point of the L-BFGS optimizer, we can run it to maximize the log-likelihood of the data. The data have seven measurement settings, consisting of four single-particle settings and three two-particle settings, as shown in Fig. 3, and the numbers of experiments performed in each setting are given in Supplementary Information section IX. Unfortunately during the optimization, sometimes the parameters run off to very large values, leading to numerical instability. In these cases, we simply restart the algorithm with slightly adjusted parameters, by adding shifts s_{ij} to the initial parameters $(c_0)_{ij}$ drawn from independent Gaussians of mean zero and standard deviation 0.1.

We would like to get a sense of the performance of this inference procedure, and in particular whether our calibrated model deviated from the maximum likelihood estimate more than would be expected from statistical fluctuation. The calibrated model $P_{M_0, (P_\lambda)_0, \mathcal{J}_0}$ is specified by the evolution parameters M_0 , the loss parameter $(P_\lambda)_0$ and the indistinguishability parameter $\mathcal{J}_0$. The evolution parameters M_0 are computed from the spectroscopic characterization. The loss parameter $(P_\lambda)_0$ and indistinguishability parameter $\mathcal{J}_0$ are computed from the HOM data that are used in the main text. The loss $(P_\lambda)_0$ is the frequency that no particles survived in the one-particle preparations of the HOM data, the indistinguishability $\mathcal{J}_0$ was computed using the method described in the section ‘Computing atom indistinguishability’.

To quantify the deviation of the maximum likelihood fit $(M^*, (P_\lambda)^*)$ to the calibrated model, we use the total variation distance of the implied distributions. Specifically, we compute the total variation distances between $P_{M_0, (P_\lambda)_0, \mathcal{J}_0}(\cdot|a)$ and $P_{M^*, (P_\lambda)^*, \mathcal{J}_0}(\cdot|a)$ for each one- and two-particle input a and take the maximum of the results. We call this the maximum total variation distance (max TVD) $d(P_{M_0, (P_\lambda)_0, \mathcal{J}_0}, P_{M^*, (P_\lambda)^*, \mathcal{J}_0})$. The max TVD has the following operational interpretation: suppose that we are allowed to choose among the 7 measurement settings to perform a single experiment, and our task is to decide whether the parameters that describe the system are $(M_0, (P_\lambda)_0, \mathcal{J}_0)$ or $(M^*, (P_\lambda)^*, \mathcal{J}_0)$. Then $\frac{1}{2} + \frac{1}{2} d(P_{M_0, (P_\lambda)_0, \mathcal{J}_0}, P_{M^*, (P_\lambda)^*, \mathcal{J}_0})$ is the optimal probability with which we could guess correctly.

We would like to capture the statistical variation in the max TVD. To do so, we perform bootstrap resamples of the HOM data to obtain a bootstrap estimate $\mathcal{J}_i^b$ of the indistinguishability, which is then used to perform MLE on bootstrap resampled data, to obtain the bootstrap fit parameters $(M_i^b, (P_\lambda)_i^b)$. We then compute the histogram of values $d(P_{M_i^b, (P_\lambda)_i^b, \mathcal{J}_i^b}, P_{M^*, (P_\lambda)^*, \mathcal{J}_0})$.

We would like to compare the value of $d_{0*} := d(P_{M_0, (P_\lambda)_0, \mathcal{J}_0}, P_{M^*, (P_\lambda)^*, \mathcal{J}_0})$ to the resulting histogram. In Extended Data Fig. 2e, we show a bootstrap histogram that shows the max TVD from the point estimate to the bootstrap MLEs. We can see that d_{0*} is slightly larger than the mean of the bootstrap distribution. This is the expected behaviour because statistical fluctuations in the calibrated model also contribute to d_{0*} , but a more thorough characterization of the statistical fluctuations in the calibrated model would be required to confirm this. Also shown in Extended Data Fig. 2e are the max TVDs from $P_{M_0, (P_\lambda)_0, \mathcal{J}_0}$ and $P_{M^*, (P_\lambda)^*, \mathcal{J}_0}$ to the frequencies of the data.

Validating few-particle quantum walks

To validate the quality of the quantum walks performed in this work, we determine the total variation distances (TVDs) between the experimental probability distributions and those expected from the calibrated model. In the limit of many samples, we can estimate this by determining the empirical TVD between the experimentally observed frequencies and those expected according to the model. If, as in our experiment, the number of trials is relatively small, the empirical TVD is expected to be biased high. We, therefore, compare the empirical TVD to that expected if we were to sample from an experimental probability distribution that is equal to the model. The latter distribution is estimated by Monte Carlo sampling from the model distribution.

In Extended Data Fig. 4, we show the binned experimental frequencies and the probability distributions of the model for two- and three-atom quantum walks at evolution times of t_{HOM} and $2t_{\text{HOM}}$, respectively. The corresponding empirical TVDs are 0.046 and 0.109. The expected empirical TVDs if the experimental distributions are the same as the models are 0.047 and 0.099, respectively. The standard deviations of the distributions of these empirical TVDs observed by Monte Carlo sampling are 0.009 and 0.006 respectively. This comparison shows that although the empirical TVDs are conservatively estimated upper bounds on the true TVDs, they are consistent with having TVDs of zero. Useful estimates require many more experimental trials, which is challenging with our experimental cycle time of about 1 s.

Instead of estimating the TVD directly, we can estimate the contributions of known discrepancies between the model and the experimental implementation. Apart from the distinguishability and loss errors discussed in the main text, there are two kinds of errors arising from evolution under different unitaries than that of the calibrated model: calibration error, which leads to a systematic discrepancy between the model and the experimental unitaries, and shot-to-shot errors due to fluctuations in the experimental unitary. Although calibration errors contribute to the TVD of the sampled distribution to the desired boson sampling distribution, they can, in principle, be accounted for by performing better calibration. Shot-to-shot errors cannot be overcome in this way and instead effectively result in sampling from a state that is mixed, which can make the associated sampling task easier to accomplish classically⁷⁴.

As discussed in the next section, and in Supplementary Information section VI, we expect the calibration errors in our system to be larger than those in state-of-the-art photonics experiments. The precision in our calibration is fundamentally limited by the number of experimental trials that we can take before drifts in the Hamiltonian parameters change the unitary applied (although we do not currently saturate this limit). The limitation in the precision of our calibration can be addressed by improving the stability of the experiment or by improving the methods of inference of the unitary. We have taken a first step in the direction of improving the methods of inference by introducing the maximum likelihood method described in the section ‘Characterizing the single-particle unitary’.

As discussed in Supplementary Information section II, we expect that shot-to-shot errors in our experiment are dominated by fluctuations in the lattice depth. We do not expect these errors to significantly limit our experiment: if we ignore the other sources of errors, we estimate that shot-to-shot errors lead to a lower bound on the fidelity of $F \geq 0.3$ after the quantum walk dynamics for the 180 atom measurements performed in this work (Supplementary Information section II).

We estimate the combined contribution of shot-to-shot and calibration errors to the TVD to be below about 0.01 for the two- and three-atom measurements considered in this section. These estimates take into account our model for shot-to-shot errors (Supplementary Information section II), the uncertainty in the spectroscopic calibration procedure (Supplementary Information section VI) and the expected

contribution to calibration errors because of finite interaction strength (Supplementary Information section VII).

Our experimental distributions also differ from boson sampling with number-resolving readout because of parity projection. We include parity projection in our models and do not consider this to be a true discrepancy between the experimental probability distribution and the models. However, we determined that the TVD between parity-projected frequencies and number-resolved ones is about 0.05 for both measurements in this section.

Comparing atomic and photonic implementations of boson sampling

We compare some relevant figures of merit of the state-of-the-art implementations of (different variants of) boson sampling in Extended Data Table 1. In photonics experiments, it can be challenging to generate and interfere large Fock states of photons because of transmission losses and the probabilistic techniques that are often^{2–8,75} (but not always^{9,76}) used for single-photon generation. To overcome this difficulty, recent experiments have performed modified versions of boson sampling that take advantage of more easily accessible non-classical states of light^{710–12,50} at the cost of requiring additional assumptions to support the claim that classical simulation of the resulting sampling problem is computationally hard^{77,78}. Moreover, recent theoretical results suggest that the presence of sufficient amounts of loss allows us to accomplish the implemented sampling tasks classically⁷⁹.

The low loss and high-state preparation and detection fidelities presented in this work enable studies of boson sampling that require fewer assumptions for the hardness of classical simulation¹. Furthermore, we do not expect the shot-to-shot errors in our experiment to exceed those in photonics experiments because of optical path length fluctuations (either in the interferometer or, in the case of Gaussian boson sampling, between the photon source and the interferometer)^{10,11}. However, it is more challenging to characterize the applied unitary U to high precision in our experiment than in photonics experiments, probably leading to larger calibration errors. In photonics, characterizations of U can be performed efficiently and with very low noise by taking advantage of bright coherent states containing a macroscopic number of photons^{10,11}. We cannot prepare equivalent states of atoms and instead rely on characterizations based on few-particle quantum walks, which have substantial experimental overhead, or on the indirect characterizations described in Supplementary Information section VI.

Units and confidence intervals

Unless otherwise noted, all error bars and uncertainties in this Article and its Supplementary Information are provided as $\pm 1\sigma$ confidence intervals.

Data availability

Experimental data used in this work are available on Zenodo at <https://doi.org/10.5281/zenodo.10453016> (ref. 80).

Code availability

Codes used in this work are available on Zenodo at <https://doi.org/10.5281/zenodo.10453016> (ref. 80).

59. Ebadi, S. et al. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature* **595**, 227–232 (2021).
60. Tian, W. et al. Parallel assembly of arbitrary defect-free atom arrays with a multitweezer algorithm. *Phys. Rev. Appl.* **19**, 034048 (2023).
61. Wang, S. et al. Accelerating the assembly of defect-free atomic arrays with maximum parallelisms. *Phys. Rev. Appl.* **19**, 054032 (2023).
62. Eschner, J., Morigi, G., Schmidt-Kaler, F. & Blatt, R. Laser cooling of trapped ions. *J. Opt. Soc. Am. B* **20**, 1003–1015 (2003).
63. Imperio, A. et al. An unsupervised deep learning algorithm for single-site reconstruction in quantum gas microscopes. *Commun. Phys.* **6**, 166 (2023).

64. Shao, J. *Mathematical Statistics* (Springer, 2003).
65. Efron, B. & Tibshirani, R. *An Introduction to the Bootstrap* (Chapman and Hall/CRC, 1994).
66. Carolan, J. et al. On the experimental verification of quantum complexity in linear optics. *Nat. Photon.* **8**, 621–626 (2014).
67. Shchesnovich, V. S. Universality of generalized bunching and efficient assessment of boson sampling. *Phys. Rev. Lett.* **116**, 123601 (2016).
68. Seron, B., Novo, L., Arkhipov, A. & Cerf, N. J. Efficient validation of boson sampling from binned photon-number distributions. Preprint at arxiv.org/abs/2212.09643 (2022).
69. Seron, B., Novo, L. & Cerf, N. J. Boson bunching is not maximized by indistinguishable particles. *Nat. Photon.* **17**, 702–709 (2023).
70. Pioge, L., Seron, B., Novo, L. & Cerf, N. J. Enhanced bunching of nearly indistinguishable bosons. Preprint at arxiv.org/abs/2308.12226 (2023).
71. Shchesnovich, V. Distinguishing noisy boson sampling from classical simulations. *Quantum* **5**, 423 (2021).
72. Oszmaniec, M. & Brod, D. J. Classical simulation of photonic linear optics with lost particles. *New J. Phys.* **20**, 092002 (2018).
73. Oh, C., Noh, K., Fefferman, B. & Jiang, L. Classical simulation of lossy boson sampling using matrix product operators. *Phys. Rev. A* **104**, 022407 (2021).
74. Oh, C., Jiang, L. & Fefferman, B. On classical simulation algorithms for noisy boson sampling. Preprint at arxiv.org/abs/2301.11532 (2023).
75. Owens, J. O. et al. Two-photon quantum walks in an elliptical direct-write waveguide array. *New J. Phys.* **13**, 075003 (2011).
76. Loredó, J. C. et al. Boson sampling with single-photon Fock states from a bright solid-state source. *Phys. Rev. Lett.* **118**, 130503 (2017).
77. Hamilton, C. S. et al. Gaussian boson sampling. *Phys. Rev. Lett.* **119**, 170501 (2017).
78. Quesada, N., Arrazola, J. M. & Killoran, N. Gaussian boson sampling using threshold detectors. *Phys. Rev. A* **98**, 062322 (2018).
79. Oh, C., Liu, M., Alexeev, Y., Fefferman, B. & Jiang, L. Classical algorithm for simulating experimental Gaussian boson sampling. Preprint at arxiv.org/abs/2306.03709 (2023).
80. Young, A. W. et al. An atomic boson sampler. *Zenodo* <https://doi.org/10.5281/zenodo.10453016> (2024).

Acknowledgements We thank S. Aaronson, N. Darkwah Oppong, I. H. Deutsch, J. P. D’Incao, A. V. Gorshkov and K. Mølmer for their discussions on the original manuscript and N. Cerf, L. Novo and B. Seron for their discussions on recent advancements in generalized bunching. We further thank N. Darkwah Oppong, N. E. Frattini, L. R. Liu, K. Mølmer and S. Sun for close readings of the paper. This work includes contributions from the National Institute of Standards and Technology, which are not subject to US copyright. The use of trade, product and software names is for informational purposes only and does not imply endorsement or recommendation by the US government. This work was supported by the AFOSR (FA95501910079), ARO (W911NF1910223), the National Science Foundation Physics Frontier Center at JILA (1734006) and NIST. S. Geller acknowledges support from the Professional Research Experience Program (PREP) operated jointly by NIST and the University of Colorado.

Author contributions A.W.Y., W.J.E., N.S. and A.M.K. contributed to developing the experiments. A.W.Y. and S. Geller performed the analysis of the results. All authors contributed to interpreting the results and preparing the paper.

Competing interests S. Glancy works as a consultant for Xanadu Quantum Technologies. All other authors declare no competing interests.


Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-024-07304-4>.

Correspondence and requests for materials should be addressed to Aaron W. Young or Adam M. Kaufman.


Peer review information *Nature* thanks Loïc Henriët and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>.




Extended Data Fig. 1 | State preparation. a, Rearrangement of atoms (blue circles) in an optical lattice (grey circles denote sites in the lattice) using optical tweezers (green) must balance several conflicting requirements, leading to the multistep algorithm described in the Methods. **b,** To optically cool lattice-trapped atoms with high fidelity, we use a pulsed cooling sequence involving 0.4 ms axial cooling pulses, and 0.2 ms radial cooling pulses (timing diagram pictured

in lower panel). We compute the expected average thermal occupation $\bar{n}$ as a function of time in each of three nearly-orthogonal axes of a given site in the lattice via a master equation calculation, yielding reasonable agreement with measured values in the experiment. Note that we optimize for high-fidelity cooling of the axial direction at the cost of slightly worse cooling in the radial directions.




Extended Data Fig. 2 | Properties of the single particle unitary. **a, b,** The single particle unitary U is depicted here for an evolution time of $t = 6.45$ ms, as is relevant to the measurements in Fig. 4. The finite waist of the optical lattice beams and resulting harmonic confinement means that U does not appreciably couple all sites to each other, and thus is not Haar random. **c, d,** U exhibits features of a Haar random matrix when considering only a 15×15 -site region near the center of the lattice. In this region, the distribution of the norm-square of the amplitudes in U are well-captured by the Porter-Thomas distribution for 385 outputs (black line in **c**). The distribution of the phases in U is well-captured by the uniform distribution (black line in **d**). **e,** We perform maximum likelihood inference of a submatrix of the single-particle unitary based on one- and two-particle data (see Fig. 3), and compare the point estimate to maximum likelihood estimates of bootstrap resamples of the data, and to the spectroscopic calibration (see Supplementary Information section VI). To quantify this

comparison, we compute the one- and two-particle distributions generated by the inferred parameters, and compute the total variation distances (TVDs) of these distributions, then take the maximum of the TVDs over the prepared input patterns. We call this quantity the max TVD between two sets of distributions. The depicted histogram is the max TVD between the point estimate and the maximum likelihood estimates of the bootstrap resampled data. Shown also are the max TVD between the frequencies of the data (Freq.) and the point estimate (Pt.), and that between the point estimate and the spectroscopic model (Model). The bootstrap histogram gives a sense of the size of the statistical fluctuation of the max TVD between the point estimate and the truth. The max TVD between the spectroscopic model and the point estimate is large compared to the bulk of the histogram, which is the expected behavior because statistical fluctuations in the model add to the statistical fluctuations in the point estimate.



Extended Data Fig. 3 | Three particle quantum walks in one dimension. The output distributions resulting from three particle quantum walks at evolution times of **a**, 1.97 ms and **b**, 4.23 ms are in good agreement with theory. Similar to the two particle case, each three particle output can be uniquely labelled by the coordinates of the three particles (x_1, x_2, x_3) , with $x_3 \leq x_2 \leq x_1$. The probability p of measuring an output state (x_1, x_2, x_3) is indicated by both the size and color of the circle at the corresponding coordinates. The prepared input states are marked by the red disks, and include patterns with nearest-neighbor (NN) and

next-nearest-neighbor (NNN) spacing. For NN input patterns, indistinguishable bosons (Indist.) exhibit enhanced probability to lie near the leading edge of the distribution along the main diagonal ($x_1 = x_2 = x_3$) in comparison to distinguishable particles (Dist.). This tendency disappears for NNN input patterns. **c, d**, Like in the two particle case, we can coarse-grain the three particle distributions by measuring bunching and clouding, and find good agreement with theory as a function of evolution time. All theory predictions in this figure correspond to error-free preparations of atoms with the appropriate particle statistics.



Extended Data Fig. 4 | Validation of two and three particle quantum walks in one dimension. The full output distribution after binning for **a**, two and **b**, three particles initialized at nearest-neighbor spacing, at an evolution time of t_{HOM} and $2t_{\text{HOM}}$ respectively. The grey bars are theory for error-free state preparation, evolution, and detection with parity projection, and the black

points are data. The upper row corresponds to distinguishable (Dist.) atoms, and the bottom row to unlabelled, nominally indistinguishable (Indist.) atoms. The outputs are grouped by the number of collisions (1, 2, or 3 atoms on the same site) that occur after binning, indicated by the inset cartoons.

Extended Data Table 1 | Comparison of large-scale boson sampling demonstrations

	n	$\mathcal{P}$	$\mathcal{J}$	r	m	Loss	Detection	Input	Evolution
[9]	20	0.975	0.93-0.954(1)	-	60	26 %	60-82 %	Click	Fixed
[10]	50*	0.938	-	1.34-1.84	100	$\sim$ 45 %	73-92 % [‡]	Click	Fixed
[11]	216*	-	-	$\sim$ 1.1	216	$\sim$ 67 %	95 %	Counting	Tunable
[12]	50*	0.962	-	1.2-1.6	144	57 % [§]	-	PPNRD	Phase
This work	180	[†]	0.995 ⁺⁵ ₋₁₆	-	$\sim$ 1015	5.0(2)%	99.8(1) %	Parity	Pattern

For works involving Fock state boson sampling, n denotes both the particle number and the number of input modes. Works involving Gaussian boson sampling are marked with a *, in which case n corresponds only to the number of input modes. $\mathcal{P} = 1 - g^{(2)}(0)$ is typically referred to as the “purity” in photonics experiments, and is measured via second order correlations in Hanbury-Brown-Twiss-like experiments. [†]To the extent that these measurements characterize the single-particle nature of the input field⁹, in our experiments $\mathcal{P} \simeq 1$ and is lower-bounded by our imaging fidelity of 0.998(1). However, our state purity is primarily limited by thermal motional excitations normal to the lattice, and can be estimated using the measured particle indistinguishability of $\mathcal{J} = 0.995^{+5}_{-16}$, which is an estimate of the purity assuming that the single-particle density matrices in the out-of-plane motional DOF are equal. To characterize state preparation, we list $\mathcal{J}$ for Fock state boson sampling results and the squeezing parameter r for Gaussian boson sampling results. m denotes the number of output modes in the linear optical network. “Loss” denotes the fraction of particles lost during evolution, including incoupling from the source to the linear-optical network, loss in the network, and outcoupling to the detectors. Detection is characterized by the detection efficiency, and the type of measurement performed on each output mode. “Click” refers to detecting the presence or absence of particles, “parity” to detecting particle number parity, “PPNRD” to pseudo-photon-number-resolving detection, and “counting” to full particle number-resolved readout. The work marked with * includes fiber coupling loss in the estimate of detection efficiency, and the work marked with [‡] includes detection efficiency in the quoted value for loss. [§]The listed value for our work is a detection fidelity rather than an efficiency, and includes contributions from both particle loss and infidelity. Converting the other listed values to detection fidelities would involve including the effects of leakage light and dark counts, resulting in slightly lower values. “Input” refers to the class of states that can be prepared as inputs to the linear optical network, with “phase” referring to tunability of the phases of the prepared squeezed states, and “pattern” to nearly arbitrary Fock states with occupations of 0 or 1 on each input mode (see Methods). “Evolution” refers to the family of linear optical networks that can be applied in a given system, with “Hamiltonian” referring to unitary evolution for variable time under a fixed Hamiltonian. For both “input” and “evolution”, “fixed” refers to a single instance, and “tunable” to flexible, but not universal, programmability. The numbers appearing in this table are representative values for approximate comparison only, please refer to the original publications for details.