Time reversed optical waves by arbitrary vector spatiotemporal field generation

Mickael Mounaix[1], Nicolas K. Fontaine[2], David T. Neilson[2], Roland Ryf[2], Haoshuo Chen[2],

Juan Carlos Alvarado-Zacarias[2] and Joel Carpenter[1]

[1] School of Information Technology and Electrical Engineering, The University of Queensland, Brisbane, QLD, 4072, Australia.

[2] Nokia Bell Labs, 791 Holmdel Rd., Holmdel, NJ 07722, USA

Supplementary Note 1 : Videos

Two videos are included which complement the content of the document itself. The first is a 5 minute summary of this work aimed at a more general audience, and the second, a 75 minute detailed technical summary of all work presented in the main document and Supplementary Information.

Video 1, *General Audience Summary (5 mins)* : https://youtu.be/1WcIejZd__w

Video 2, *Detailed Technical Summary (75 mins)* : https://youtu.be/9hVEJvfWjRQ

Technical Video Table of Contents {minutes:seconds}

Supplementary Methods

Experimental apparatus

Summary

(video timestamp : 24m14s)

The experimental apparatus can be thought of in two parts; the spatially and polarisation resolved spectral pulse shaper itself, and the swept-wavelength digital holography system used to characterise it. Both are shown in Fig. 2 of the main document which for convenience is repeated here as Supplementary Figure 1.
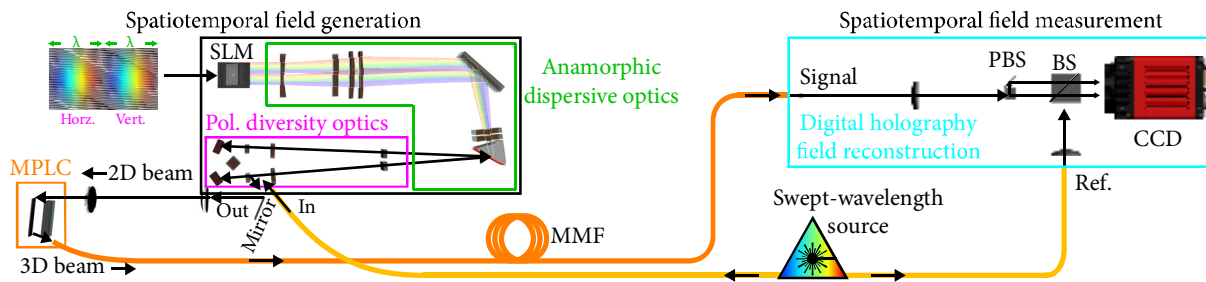
Supplementary Figure 1 | Schematic of optical time reversal system and characterisation apparatus. A polarisation and spatially resolved spectral pulse shaper for generating arbitrary vector spatiotemporal states, in conjunction with a swept-wavelength digital holography system for characterisation.
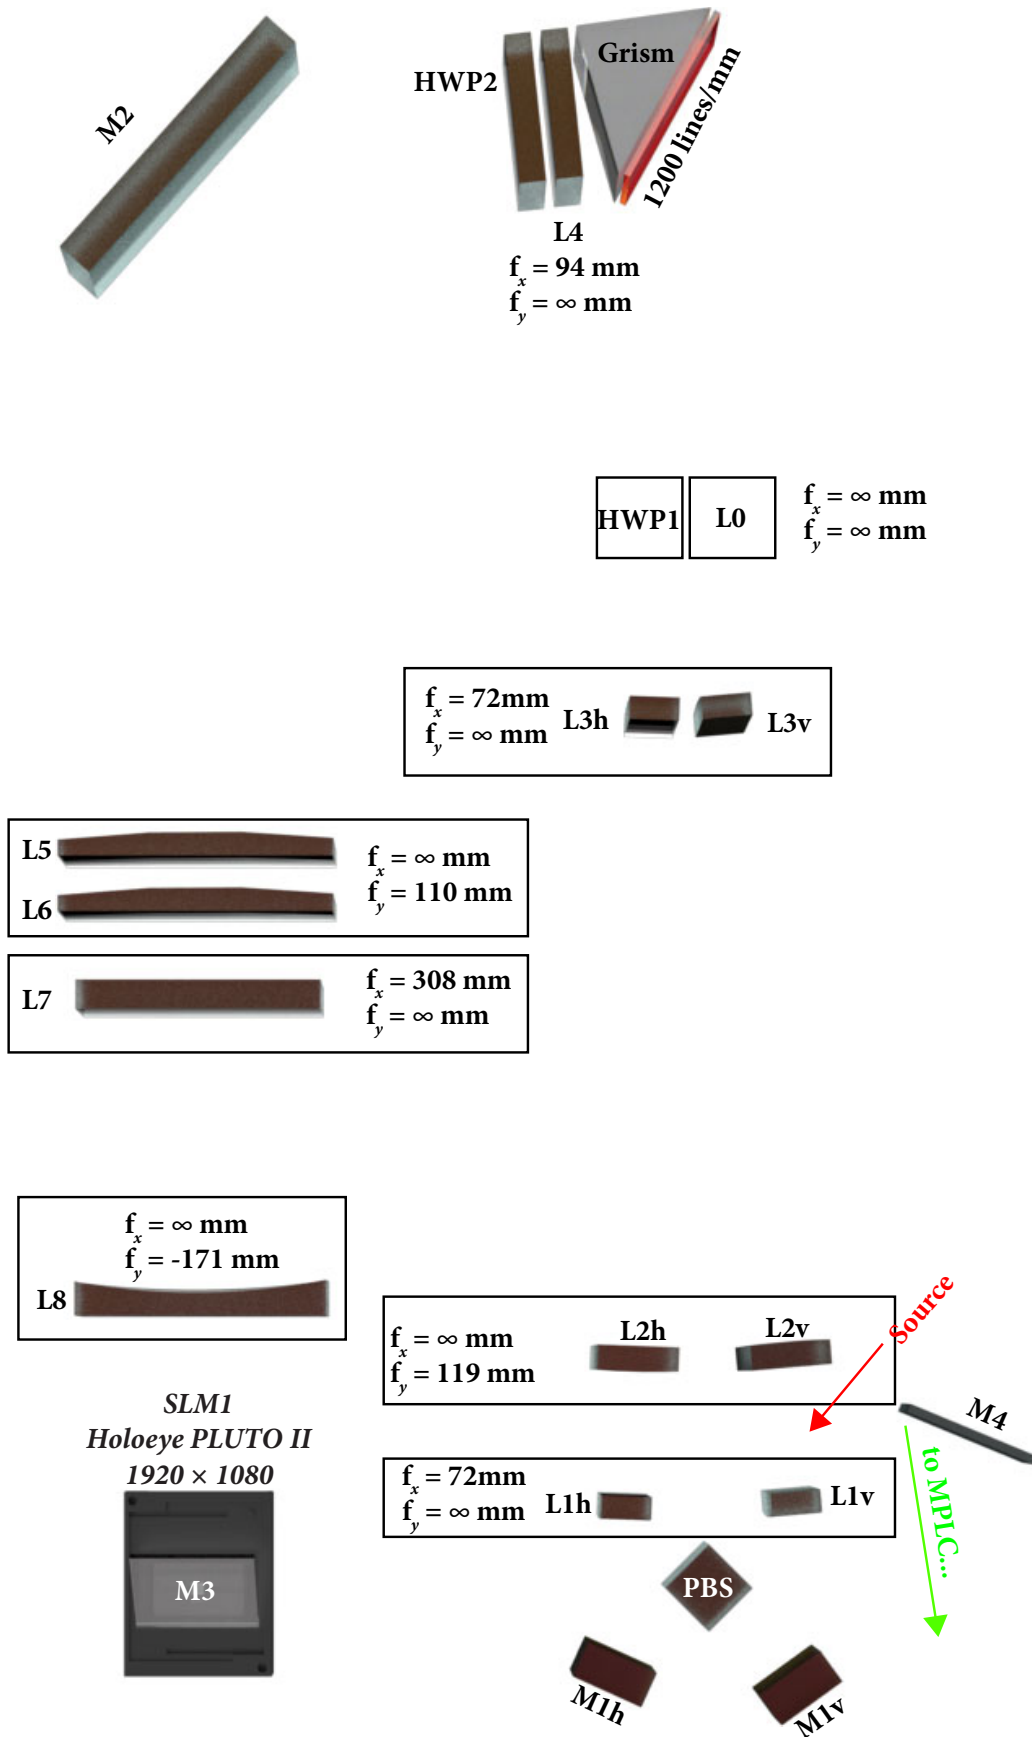
The spatially and polarisation resolved spectral pulse shaper, is illustrated on the left-side of Supplementary Figure 1 and consists of two optical sub-systems. First, the polarisation-resolved multi-port spectral pulse shaper sub-system. This consists of polarisation diversity optics (pink in Supplementary Figure 1), dispersive and beam resizing optics (green), and the spatial light modulator (SLM). The operation principles of this portion of the device are similar to that of other spectral pulse shapers[1] or wavelength-selective switches[2], albeit with the unusual inclusion of polarisation-diverse processing and a relatively large number of addressable single-mode Gaussian output ports. The output ports in this case are a 1D array of 45 Gaussian spots, with each spot being addressable by steering light in the direction in/out of the page as of Supplementary Figure 1 by applying the appropriate phase ramp across the vertical direction of the SLM. Light from the spectral pulse shaper section is imaged through a 4*f* telescope (*f*=10cm) onto the multi-plane light conversion (MPLC) device[3]. This MPLC sub-system maps the 1D array of Gaussian spots onto a 2D set of HG modes through 14 planes of phase manipulation. For a given 2D phase mask programmed onto the SLM, a corresponding 2D field is generated in the Fourier plane, approximately at the grism. This 2D field in turn becomes a 3D field through the 1D-to-2D transformation of the MPLC device. This field is coupled into the multimode optical fibre (MMF)[4] at the output of the MPLC. As light propagates through the MMF it becomes redistributed spatially and temporally due to mode coupling and mode dispersion along the length. At the output of the MMF this field propagates into the swept-wavelength digital holography portion of the setup. This portion is used to characterise the spatial, polarisation and spectral/temporal properties of the beams[5,6]. Both to measure the optical transfer function[7] of the device and the fibre under test, as well as to characterise the spatiotemporal and spatiospectral beams being generated out the output of the fibre.

The swept-wavelength laser used is a NewSpec TLM-8700-L-CL. The CCD camera is an Allied Vision Goldeye CL-008. An external Mach-Zehnder interferometer samples light from the TLM-8700 as the laser is swept and through additional circuitry generates camera triggers at fixed frequency spacing. Which for the results presented here is 10.0675 GHz spaced frequency samples from 195.510 THz to 191.000 THz. The swept-wavelength digital holography characterisation procedure is similar to that explained in previous work[5,6].
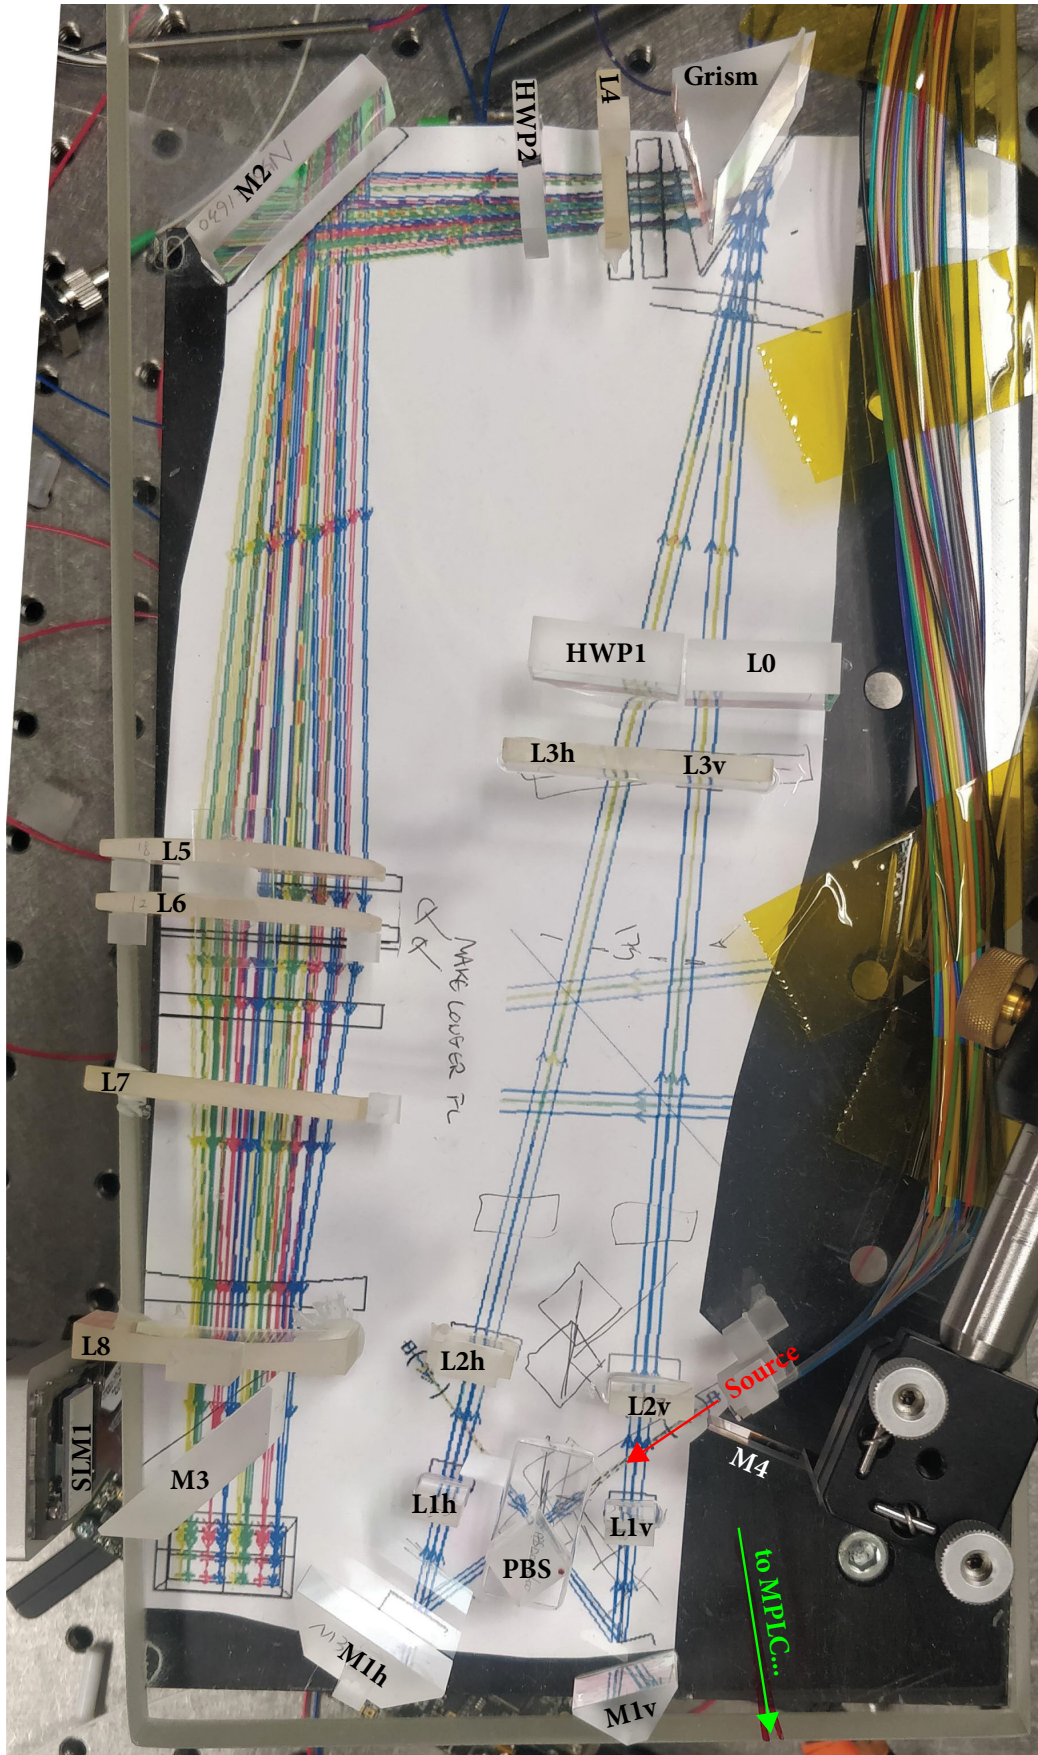
Spectral pulse shaper sub-system

The spectral pulse shaper sub-system design schematic is illustrated in Supplementary Figure 2, in addition to a photograph of the sub-system in Supplementary Figure 3. The source light enters from the TLM-8700 swept-wavelength laser source as a single-mode optical fibre (SMF-28) attached to a microlens such that the source beam enters the optical system with a mode-field diameter (MFD) of 60 μm. The same beam diameter as the spots in the 1D array of the MPLC device which are optically in the same plane as the source. The input light has been aligned by a fibre polarisation controller attached to the TLM-8700 such that approximately equal power is in the horizontal and vertical polarisation components with respect to the optics of the pulse shaper at the output of the MMF. Hence input power will exit both ports of the polarising beam splitter (PBS) and propagate through spatially separated polarisation diversity optics for each polarisation component. M1h and M1v are simply dielectric mirrors. L1h/v and L3h/v are a set of identical cylindrical lenses which image the input source spot onto the grism in the y-direction (in/out of page, spatial axis of the device). L2h/v are cylindrical lenses along the orthogonal direction, collimating the input source spot onto the grism in the x-direction (plane of the page, spectral/temporal axis of the device). That is, the spot at the grating is small in the spatial direction, and wide in the temporal direction. When Fourier transformed onto the SLM for each spectral component, this becomes a large beam in the spatial direction and narrow in the spectral direction, corresponding with the ability to address a large number of spatial ports and large delays. HWP1 is a half-wave plate to align one of the incoming polarisation components (H) to that of the grating. For the other polarisation component (V), which is already aligned optimally for the grating, L0 is simply a flat block of glass to maintain optical path length matching between the polarisation components. Light is then incident on the grism and starts to disperse along the spectral/temporal axis (plane of the page). Together L4, L5, L6, L7 and L8 perform some resizing of the beams in the x and y direction, but ultimately are setup to place the SLM in the Fourier plane of the grism. M2 and M3 are simply dielectric mirrors and HWP2 aligns the incoming polarisation from the grism to that required by the SLM.

Programming a phase mask onto the SLM will cause light to be redistributed amongst the spatial and spectral/temporal axes of the device. Steering light to different positions along the spatial axis (y-axis, in/out of the page) of the grating correspond with different Hermite-Gaussian (HG) modes at the output of the MPLC. Different positions along the temporal axis (x-axis, plane of the page) correspond with longer or shorter path lengths traced back through the outer or inner edges of the grism. Light retraces its path back through HWP1/L0, L3, L2, L1, M1, and the PBS towards the input source. However due to a slight deliberate tilt, light through the return path is not coupled back into the laser source, but instead, is picked off with a mirror (M4) and routed towards the MPLC sub-system.

M2

HWP2

Grism

1200 lines/mm

L4
$f_x = 94$ mm
$f_y = \infty$ mm

HWP1   L0 $f_x = \infty$ mm
$f_y = \infty$ mm

$f_x = 72$mm
$f_y = \infty$ mm   L3h   L3v

L5 $f_x = \infty$ mm
L6 $f_y = 110$ mm

L7 $f_x = 308$ mm
$f_y = \infty$ mm

$f_x = \infty$ mm
$f_y = -171$ mm
L8

SLM1
Holoeye PLUTO II
1920 × 1080

M3

$f_x = \infty$ mm
$f_y = 119$ mm   L2h   L2v

Source

M4

to MPLC...

$f_x = 72$mm
$f_y = \infty$ mm   L1h   L1v

PBS

M1h   M1v

Supplementary Figure 2 | Design schematic of the multi-port polarisation-resolved spectral pulse shaper sub-system
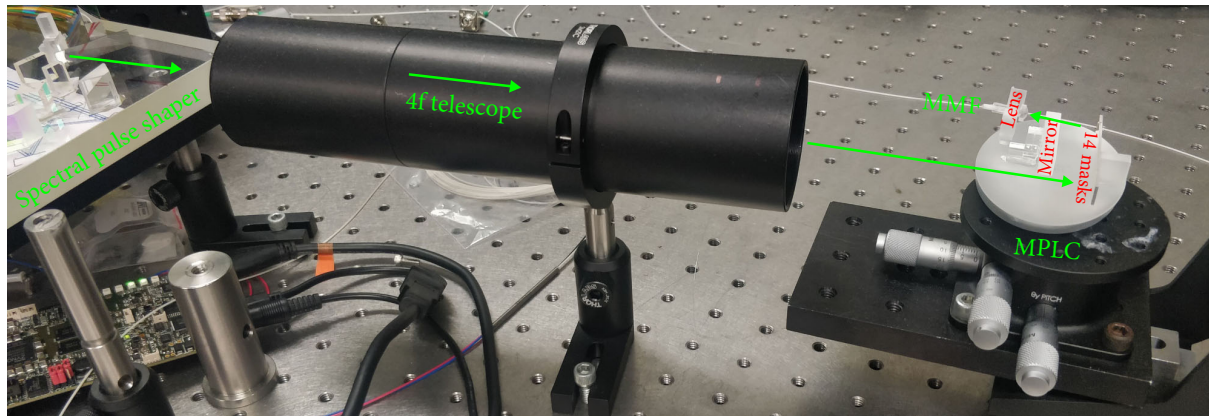
Supplementary Figure 3 | Photograph of the multi-port polarisation-resolved spectral pulse shaper sub-system

Multi-plane light conversion sub-system

As shown in Supplementary Figure 4, light is routed from the spectral pulse shaper sub-system to the MPLC sub-system by way of a 4*f* telescope (*f*=10cm). The MPLC device itself[3,8] consists of 14 lithographically etched phase planes which map a 1D array of 45 Gaussian spots (60 μm MFD, 127 μm pitch), to 45 Hermite-Gaussian modes of MFD 592 μm. Which are in turn coupled to the 50 μm diameter core graded index multimode fibre[4] through a *f*=4.51 mm lens. Further details on the MPLC design available here[3], as well as

https://www.youtube.com/watch?v=iXjFQj15Xzg&t=21m47s

Supplementary Figure 4 | Photograph of the lens relay imaging light from the spectral pulse shaper sub-system to the MPLC sub-system

Digital holography sub-system

The digital holography sub-system is the characterisation apparatus used to measure the optical fields exiting the multimode fibre at the distal end. In off-axis digital holography, the intensity of the interference between the beam being characterised (image of the MMF core facet) and a reference beam (tilted quasi-plane wave) are measured. From this interference pattern the full complex optical field can be recovered[9]. A similar setup has been used previously[5] and is explained in more detail here; https://www.youtube.com/watch?v=TqAr4UFtoqw&t=38m33s
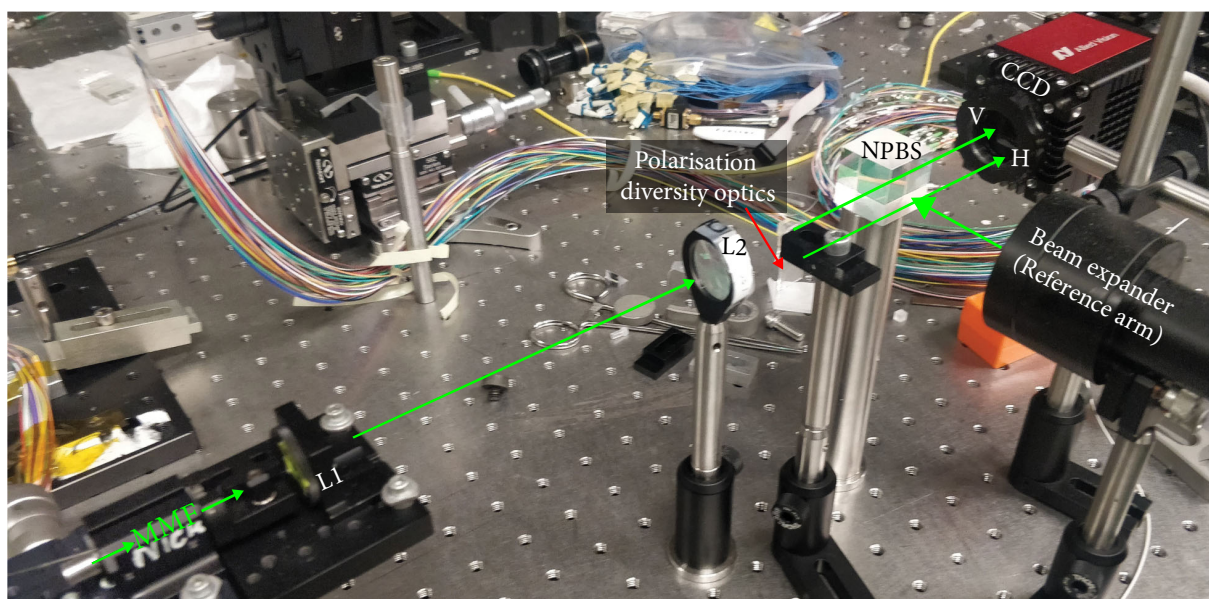
Supplementary Figure 5 | Photograph of off-axis digital holography beam characterisation sub-system.

The beam exiting the multimode fibre is magnified by a factor of 50 through L1 (*f*=3 mm) and L2 (*f*= 150 mm). The two polarisation components of the beam are displaced from one another using polarisation diversity optics before

travelling through a non-polarising beamsplitter (NPBS) where these beams are interfered with the reference wave. The reference wave is a single-mode fibre attached to a beam expander coming from the same TLM-8700 laser source which enters the spectral pulse shaper. The interference between the reference wave and the horizontal and vertical beam components are recorded on the camera and then processed to recover the fields. The CCD camera is triggered by a k-clock. That is, an external Mach-Zehnder interferometer connected to additional circuitry which signals the camera to record images every 10.0675 GHz as the wavelength of the laser is swept. The measured optical fields in the basis of camera pixels are then overlapped with the first 55 Hermite-Gaussian modes to recover the complex amplitude of each spatial mode component in each polarisation. Ultimately the fibre and spectral pulse shaper only support 45 spatial modes per polarisation, however additional modes are decomposed for the purposes of alignment and checking the validity of results. The resulting data is then processed to remove the polarisation mode dispersion of the digital holography setup itself as well as removing residual group delay (6.05 ps) and chromatic dispersion (0.0127 ps^2) of the reference beam.

Measurement of the optical transfer function (25m44s)

The optical transfer function is a complete linear description of the system as a whole. In this case, it consists of a set of frequency-dependent complex transfer matrices which map any spatiotemporal input to any spatiotemporal output. It consists of 449 matrices (10.0675 GHz spaced samples across 4.5 THz). Each matrix maps the linear relationship between 90 spatial/polarisation modes as excited by the device itself (spectral pulse shaper in combination with MPLC and attached MMF) and the 110 spatial/polarisation modes measured on the digital holography system.

With regards to programming the SLM, each input mode is ultimately specified by a single number; the position of the corresponding spot in the 1D Gaussian array of the MPLC as seen from the plane of the SLM. As the SLM is in the Fourier plane with respect to the 1D array, each output spot corresponds with a particular angle of incidence on the SLM and the phase of the wavefront corresponding with this angle is what must be programmed onto the SLM to steer light to that position. Remembering that wavelengths are dispersed across the horizontal axis of the SLM, so the phase front for each spectral component will be slightly different, although the physical location (angle or position in the output array) is common to all wavelengths. Additional corrections such as aberrations and misalignments are also applied for best performance. For example, defocus, and the mapping of wavelength to pixel position on the SLM. This includes slight misalignments of the SLM axis with respect to the grating axis, as well as chromatic aberrations associated with the grism. A consistent basis for both the measurement of the transfer matrices, as well as the calculation of the required phase masks are required to successfully generate the desired output spatiotemporal and spatiospectral states of interest. The system is calibrated as a whole, by a single set of matrices which encapsulates the spectral pulse shaper, MPLC, MMF and digital holography sub-systems. As mentioned above, the polarisation mode dispersion, delay and chromatic dispersion of the digital holography apparatus is removed from the measurements.

Supplementary Figure 6 contains measured examples of the time reversal operator[10]. That is, the transfer matrix (T) of the system multiplied by its conjugate transpose (T^*). Supplementary Figure 6a contains the worst performing frequency (195.409 THz) and Supplementary Figure 6b corresponds with the best performing frequency (193.819 THz). Performance in this case is measured in terms of mode-dependent loss, which is the ratio of the maximum and minimum singular value, also known as the condition number. For an ideal device, the transfer matrix, T, would be unitary and the matrix, TT^*, would be the identity matrix, ignoring some overall loss which is common to all modes supported by the system. That is, an ideal device would have a transfer matrix where all the singular values are all the same, the inverse of T is equal to T^* and the matrices of Supplementary Figure 6 would be diagonal, with all diagonal elements having the same value. Any deviation from this scenario means at least some degradation in the ability of the system to generate some fields that contain superpositions of spatial/polarisation components which the system cannot generate efficiently, if at all.
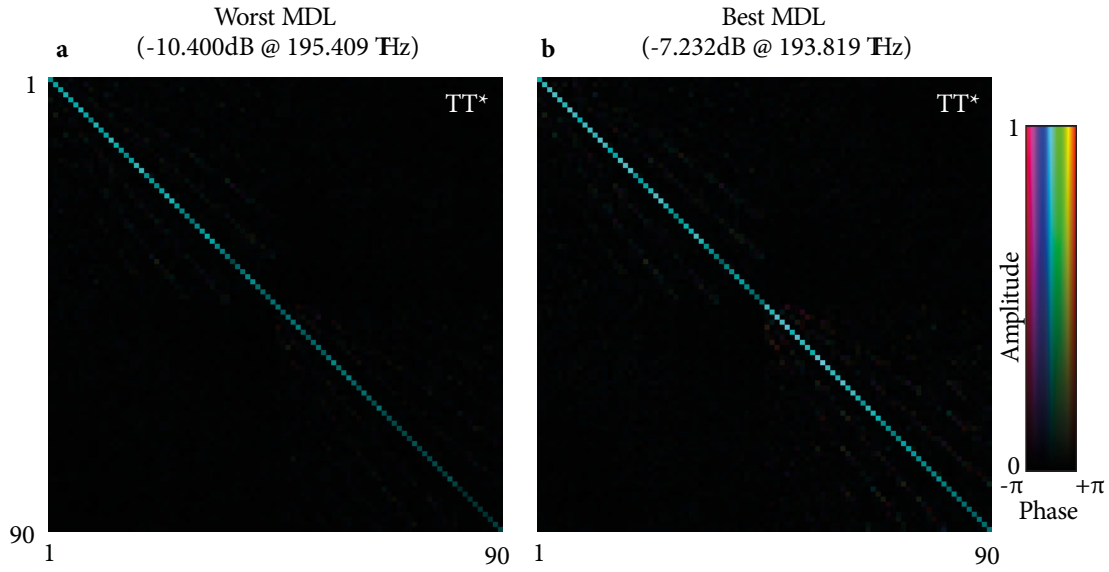
Supplementary Figure 6 | Measured time reversal operators for the system. a, Worst MDL frequency and **b,** Best MDL frequency.

Calculation of the SLM hologram (31m18s)

Generation of the desired spatiotemporal or spatiospectral field requires calculating a phase mask to display on the SLM which excites the correct amplitude and phase for each spatial and polarisation mode as a function of wavelength. How this 2D phase mask distributes light in 2D in the Fourier plane, approximately at the grating, defines how light from the source is ultimately distributed in space and time. The task of calculating the phase mask can be approached either in the spectral domain, or the temporal domain. We calculate the masks using a modified Gerchberg-Saxton (GS) algorithm[11,12]. The traditional GS algorithm is structured as a phase retrieval problem. It consists of two planes separated by a Fourier transform. In each plane, the intensity is a fixed goal and it is the phase in each plane that are the free parameters which the algorithm will attempt to optimise. In our application, the free parameters are arranged differently. The phase in the plane of the SLM is a free parameter, however the phase in the plane of the 1D array is not. All spots in the array must be excited with a specific amplitude and phase. Hence the free parameter in the plane of the 1D array is the region of space outside the array. Any light scattered to this region, will not be coupled into the MMF and hence not affect the beam quality, at the unavoidable expense of loss. For a single plane of phase manipulation it is impossible to guarantee lossless conversion. This is a property of all single plane holograms, not specific to this application. The reason for this can be understood by imagining the desired output spatiotemporal state back propagating through the system from the camera plane to the SLM plane. If a particular spatiotemporal state does not have an amplitude on the SLM that matches the amplitude of the Gaussian source beam coming from the laser in the opposite direction, there is nothing that a single phase mask can do to match the two beams (62m29s). At least two masks would be required, which can be imagined as one mask to redistribute the amplitude onto the next mask such that this second plane can match the phase (73m12s).

The modified GS algorithm iterates through the loop as follows;

a) Set the amplitude in the plane of the SLM to the amplitude of the Gaussian beam from the laser source. The phase in this plane remains unchanged from the previous iteration.
b) Fourier transform the plane of the SLM into the plane of the output 1D array.
c) Check the normalised overlap (beam quality) of the generated field with the goal field. Check the unnormalised overlap (loss). We want to know both how accurately the mask is exciting the modes with the correct amplitude and phase, but also how much power we are discarding to achieve this.

d) The currently generated field in the region of interest nearer the 1D array is then replaced with the ideal target field with some weighting. The relative weighting between the ideal target being applied inside the region of interest, and the background scattered light which is allowed to remain, defines the total loss of the hologram. If the background scatter is weighted too high, the overall loss will be high. If the background scatter is weighted too low, the algorithm does little on every iteration converging to a low quality result.

e) The algorithm is continually iterated until a desired beam quality is achieved. Typically 95%.

For the experimental results shown here, we have calculated the masks as a set of N 1D holograms corresponding to N spectral components. In order to promote smooth 2D masks, the result from the 1D hologram for one wavelength is used as the starting point for the calculation of the next 1D hologram. An alternative approach of similar effect is to calculate the holograms in 2D using a target distribution specified in the time domain, with a spatial filter in the iteration loop to promote low spatial resolution solutions.

An example illustration is shown in Supplementary Figure 7. In this example, two focused spots are being created at the output of the fibre. One spot is horizontally polarised at 0 ps and the other is a vertically polarised spot at a different spatial location at a delay of -26 ps. Supplementary Figure 6a is the literal greyscale bitmap programmed onto the SLM which includes all corrections and calibrations. Supplementary Figure 7b are the numerically calculated phase masks with the illuminating beam coming from the source. These numerically calculated masks are calibrated and transformed onto the surface of the SLM to become Supplementary Figure 7a. For the numerically calculated masks, the x-axis represents frequency, and the y-axis represents Fourier space (k_y) with respect to the 1D input array of the MPLC. For the results shown in this work, the masks are calculated in the spectral domain, and the resulting field are shown in Supplementary Figure 7c. Each spectral component of Supplementary Figure 7b has been Fourier transformed to reveal the distribution of amplitude and phase in the spectral domain at the input to the MPLC. This plane is divided into two regions, a region of interest, corresponding with the positions of the 45 inputs of the MPLC device, and a "don't care" region, where light must be discarded in order to achieve the desired beam quality as mentioned above. The plane of Supplementary Figure 7c does not physically exist in the actual device, it is simply the plane in which the target field is computed. However for an alternate implementation of this device based on two planes of phase manipulation using two SLMs, this plane would correspond with a second SLM. Supplementary Figure 7d is the 2D Fourier transform of Supplementary Figure 6b and represents a physical plane in the device at the grating or at the input to the MPLC.

a Horizontal polarisation   Vertical polarisation

Space (ky)

Greyscale phase levels as displayed on SLM

Frequency   Frequency

b Pixel 1

Space (ky)

Pixel 1080   Plane of SLM

c +Max/2

Space (y)

"Don't care region" (discard region)

45 ports of MPLC (region of interest)

"Don't care region" (discard region)

Plane of GS goal (computational, not physical)

-Max/2

195.5   Frequency (THz)   191.0   195.5   Frequency (THz)   191.0

d +Max/2

Space (y)

-26 ps   0 ps   -26 ps   0 ps

"Don't care region" (discard region)

45 ports of MPLC (region of interest)

"Don't care region" (discard region)

Plane of grating

-Max/2

-50   time (ps)   +50   -50   time (ps)   +50

Supplementary Figure 7 | Example phase masks and corresponding fields in the spectral and temporal domains for the generation of two delayed spots at the end of the multimode optical fibre.

Supplementary Note 2 : Characterisation of spatially-dependent loss (47m35s)

From the measured optical transfer function of the device it is possible to extract the spatial and frequency-dependence of properties such as loss. Another advantage of the full linear description obtained through a measurement of the frequency-dependent transfer matrix of the device is that all possible linear properties can be extracted. Properties such as the loss of each individual mode/port as addressed in the MPLC array, but also absolute worst-case performance metrics can be extracted such as mode-dependent loss (MDL) and averages over all possible combinations of modes such as insertion loss (IL). MDL, also known as the condition number of a matrix, is a measure of how "invertible" the matrix is. It is the ratio between the maximum and minimum singular value of the matrix and represents the maximum possible difference in loss between any two superpositions of states through the device. A detailed discussion and explanation of the singular value decomposition and the concepts of insertion loss and mode dependent loss as they relate to optical devices such as these have been previously discussed[6,13] (https://www.youtube.com/watch?v=TqAr4UFtoqw&t=45m50s). The insertion loss of the device can be seen in Supplementary Figure 8a. At the centre wavelength there is 17 dB of insertion loss. Which consists of approximately 3 dB of polarisation-dependent loss due to the static splitting of the horizontal and vertical components of the beam onto separate sides of the SLM. As loss is defined here, when a given port of the MPLC is being selected by the spectral pulse shaper in a given polarisation, the power in the orthogonal polarisation is being considered loss. This is a worst-case assumption, as in practice most input states contain components of both horizontal and vertical components, or the input polarisation could be aligned to the dominant axis to minimise loss. An alternate design for this device would address all 90 spatial/polarisation modes as a single 1D array as seen by the SLM. That is, no polarisation splitting onto the SLM surface, only polarisation recombining near the MPLC. The advantage of such an approach would be a reduction in worst-case loss, at the expense of higher requirements on the spatial resolution of the SLM and its associated phase masks. The MPLC device has approximately 5 dB of loss. This consists of approximately 2 dB due to the reflectivity of gold over 14 reflections, 1 dB of theoretical loss, and a remaining 2 dB of coupling loss and other imperfections. The 45-port spectral pulse shaper incurs the remainder of the loss at approximately 9 dB. For context, commercially available multi-port spectral pulse shapers such as a Finisar Waveshaper 16000A (19-port) are currently specified as 5 dB of insertion loss.

The measured mode dependent loss for the device is relatively good at between 7.2 dB for 193.819 THz and 10.4 dB at 195.409 THz. Whilst minimising insertion loss is important for applications that are power-limited, MDL is the more important measure with respect to the beam quality and the number of spatial channels the device can truly address. A high MDL means the transfer matrix of the system is difficult or impossible to invert and some spatiotemporal states may not be accessible, as the required spatial states are effectively missing. As opposed to MDL, which is calculated using the singular value decomposition and corresponds with superpositions of the modes the transfer matrix was measured in, Supplementary Figure 8b illustrates the losses in terms of the ports of the MPLC addressed by the spectral pulse shaper, a total of 90 ports/modes corresponding with 45 spatial modes in both polarisations. The differential port loss for a given frequency is fairly consistent at approximately 5 dB, with a roll-off of approximately 4 dB across the frequency band, as can also be seen in the insertion loss plot of Supplementary Figure 8a. Taken together, the lowest loss port at the lowest loss frequency has approximately 9 dB better loss than the worst port at the worst frequency. Similar information is presented in Supplementary Figure 8c as the total loss of each port at each frequency. The 45 modes located at the top of Supplementary Figure 8c correspond with the 45 ports in the horizontal polarisation, and the bottom 45 modes correspond with the same 45 ports in the vertical polarisation. For example, mode 1 and mode 46 are addressing the same port of the MPLC, from the left and right sides of the SLM respectively, corresponding with the two orthogonal polarisation components. The ports are enumerated in order of the Hermite-Gaussian modes they address at the output of the MPLC, rather than their 1D positions at the input to the MPLC. That is, ports 1 and 2 are not physically adjacent spots at the input to the MPLC, they are adjacent modes in Hermite-Gaussian space. Mode 37 and the corresponding mode in the orthogonal polarisation (mode 82) are the worst ports of the MPLC device, particularly at the lower frequencies, which can also be seen as the light blue and red traces of Supplementary Figure 8b.
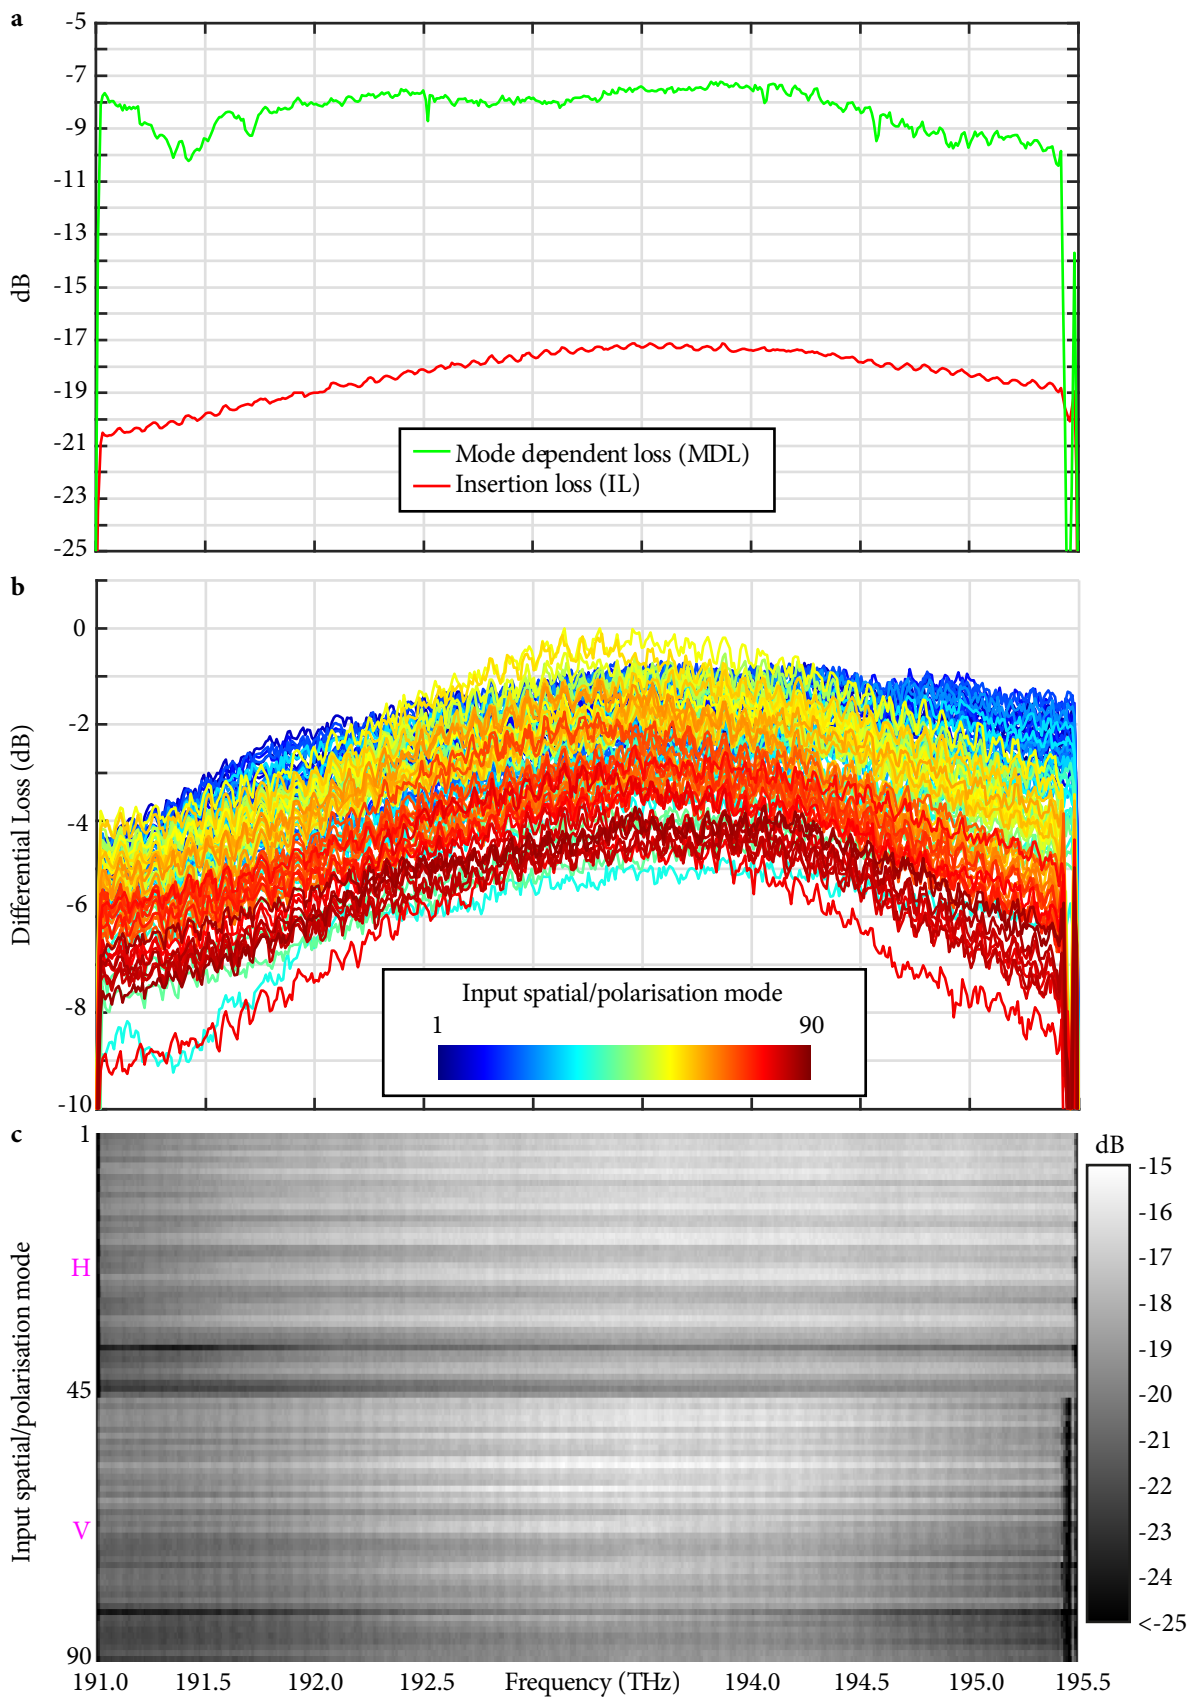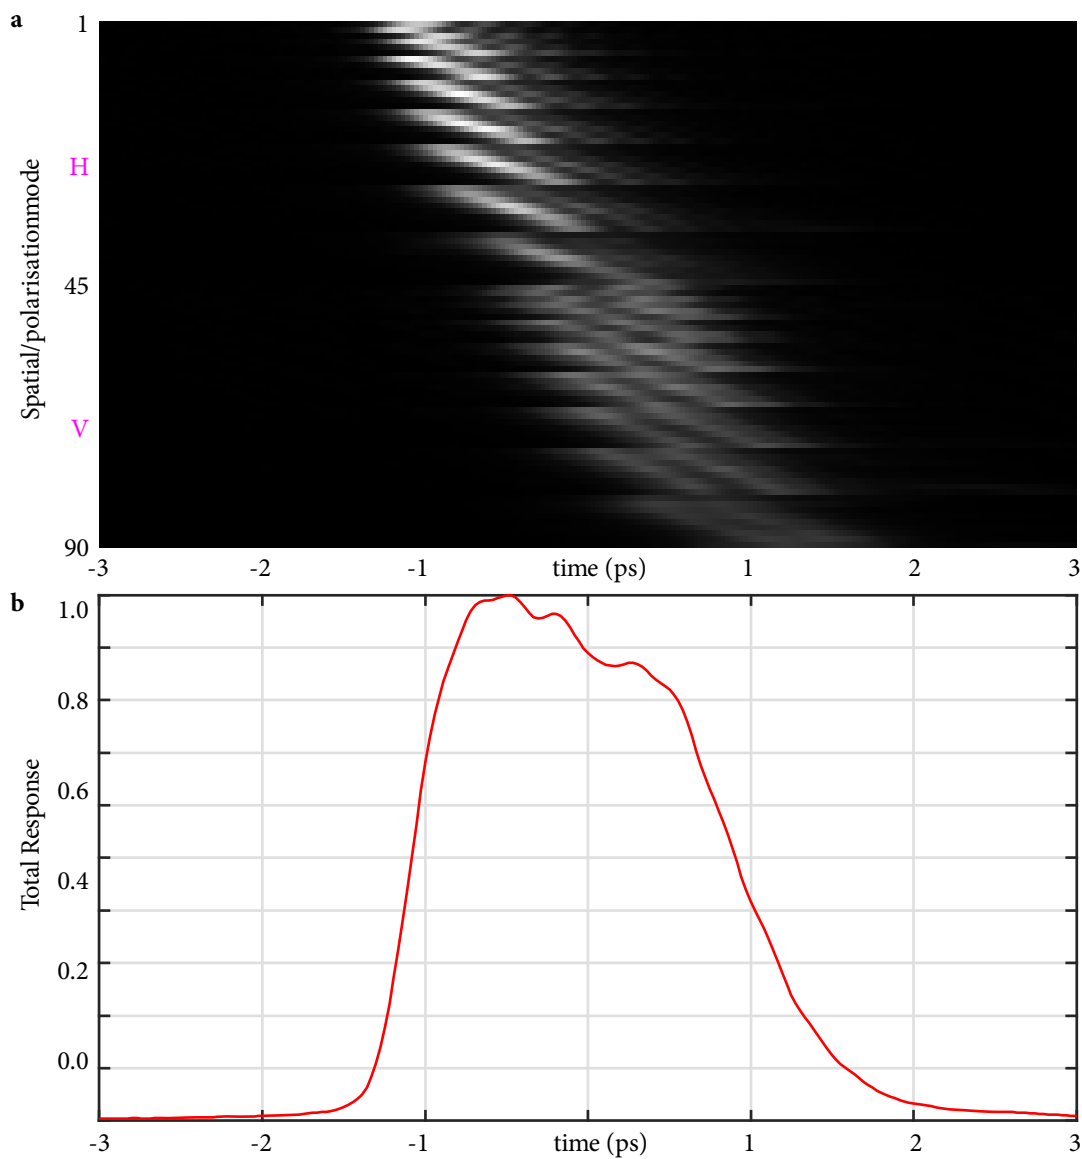
Supplementary Figure 8 | Characterisations of the spatial and frequency dependent losses of the device. a,

Insertion loss and mode-dependent loss. **b,** Differential loss between HG modes excited at the input of the fibre as seen from the output of the fibre. **c,** Total loss between HG modes excited at the input of the fibre as seen from the output of the fibre.

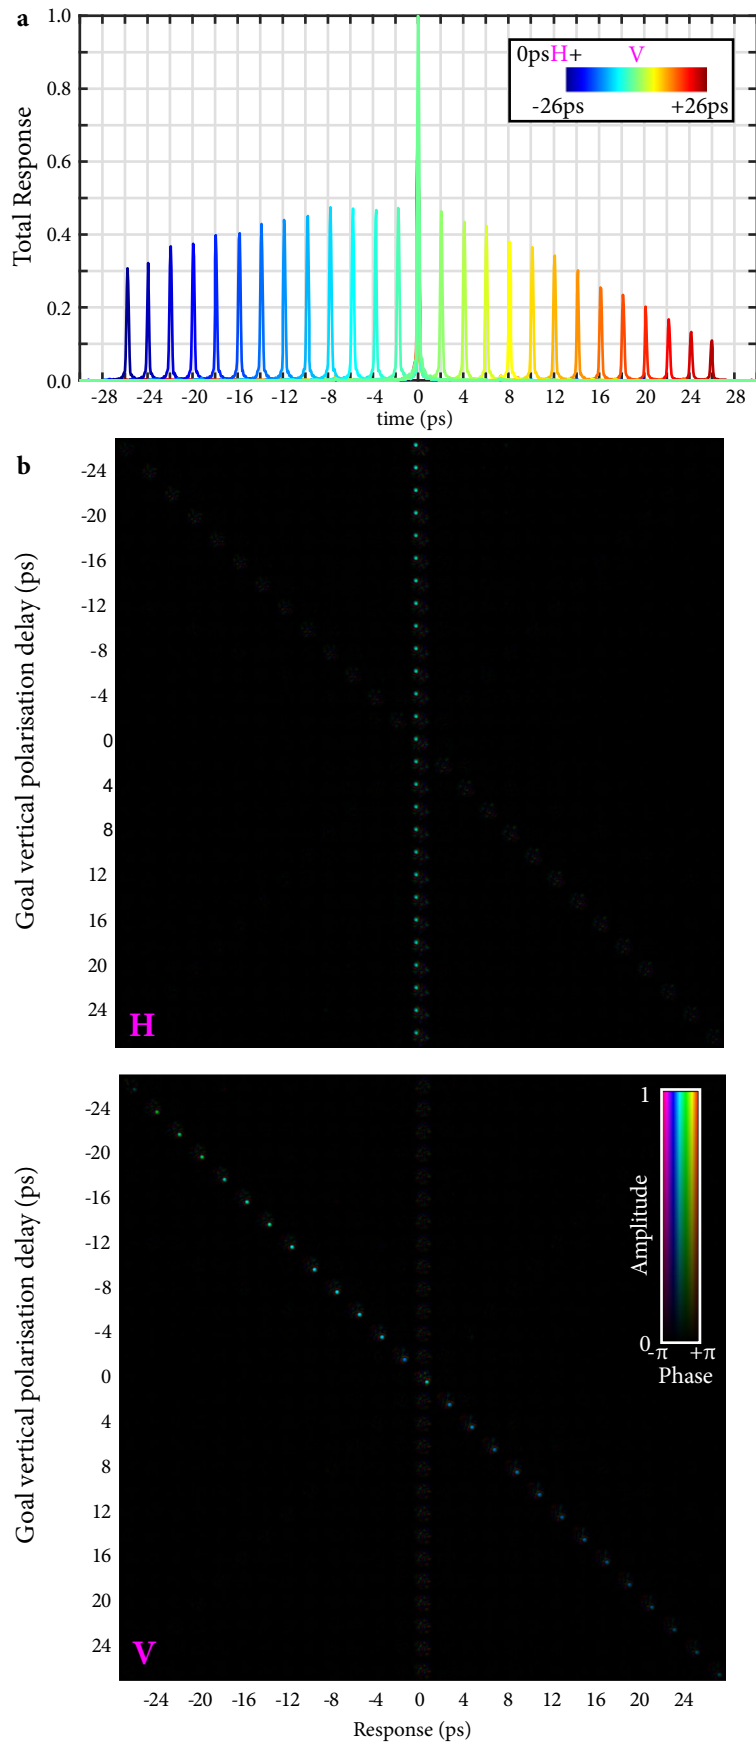## Supplementary Note 3 : Characterisation of impulse response (55m40s)

Supplementary Figure 9 illustrates aspects of the temporal performance of the device, including the attached MMF. Supplementary Figure 9a is the impulse response obtained from the measured transfer matrices. For each spatial HG mode in each polarisation as addressed at the input of the MMF, the corresponding impulse response is shown. The total impulse response averaged over all modes in both polarisations is illustrated in Supplementary Figure 9b. This is the spatiotemporal response of the system as a whole and includes contributions both from the MMF, as well as the spectral pulse shaper sub-system. The differential group delay and chromatic dispersion within the spectral pulse shaper is primarily between the two orthogonal polarisation components. The MPLC sub-system should have negligible modal dispersion. In total there is approximately 2 ps of differential delay with approximately equal amounts of delay introduced within the spectral pulse shaper itself, as well as the MMF.

Supplementary Figure 9 | Characterisation of the impulse response and delay-dependent loss of the device. a, Measured impulse response for each input HG mode in each polarisation. **b,** Total impulse response over all spatial/polarisation modes, this is a characterisation of the uncompensated optical system.

Supplementary Note 4 : Characterisation of delay-dependent loss (56m52s)

Supplementary Figure 10 illustrates a characterisation of the delay-dependent loss of the device. There are multiple mechanisms by which delay-dependent loss are introduced within such a device. First, there is the delay-dependent tilt of the beam, which reduces coupling into the MPLC. A delayed beam path will be spatially offset in the plane of the grating, which results in a corresponding angle of incidence when Fourier transformed into the plane of the MPLC input. Second, there is a delay-dependent defocus of the beam associated with the fact that it has physically travelled a longer path. Third, and the most significant contribution here, is the diffraction efficiency of the SLM. Creating light at large delays involves the steering of light to large angles across the spectral axis of the SLM. The diffraction efficiency of an SLM decreases as the diffraction angle is increased. This is partially due to quantisation and pixelisation effects, but largely due to the non-ideal performance of liquid crystals[14]. In the test of Supplementary Figure 10, a horizontally polarised focused spot is created at the output of the MMF at a delay of 0 ps. In the vertical polarisation at the MMF output another focused spot is created at a different spatial location. The delay of this vertically polarised spot is stepped from -26 ps to 26 ps in 2 ps steps. The measured fields in both polarisation components at the output of the fibre for each trial can be seen in Supplementary Figure 10. The x-axis of these images represent the spatial field at a particular delay for a given phase mask on the SLM, in the same manner as examples such as Fig. 3 of the main document. The y-axis represents different phase masks designed to generate a vertically polarised spot at a desired delay. The results of Supplementary Figure 10a are normalised with respect to the case where the spots in both polarisations arrive at the same delay of 0 ps. Hence for all other delays, the ideal case would be a response of 0.5, as naturally the intensity in a given time interval is halved when the spots are temporally separated. It can be seen that there is asymmetry in the delay-dependent loss, with negative delays suffering less loss than positive delays of the same magnitude. The exact mechanisms that contribute to this asymmetry cannot be definitely measured, however one aspect could be the differing diffraction efficiency of the SLM when steering in one direction versus the other[14]. Other mechanisms could be related to issues of optical alignment such as defocus into the MPLC input or clipping of the beam.

**Supplementary Figure 10 | Measured fields for characterisation of delay-dependent loss.**

Supplementary Note 5 : Spatiospectral stress test (59m26s)

The tests of Supplementary Figure 11 to Supplementary Figure 15 are specifically designed to be challenging for the device to implement. Tests which require all spatial and polarisation modes to be excited with equal power by the SLM across the entire operating frequency bandwidth of the device, with rapid changes in mode superposition from one frequency component to the next. That is, tests which require the SLM to address dense combinations of the spatial and temporal degrees of freedom simultaneously with a single phase mask. The spatiotemporal states these represent are not likely to be of practical use, or likely to occur in any realistic setting. They are contrived as a means of testing the limits of the device in worst-case scenarios. The tests consist of implementing columns of the discrete Fourier transform (DFT) matrix in spectral bands of between 10.0675 GHz and 50.3375 GHz. Corresponding with approximately 2 to 10 pixel widths across the spectral axis of the SLM. The significance of the DFT matrices here is that they require the SLM to generate all 90 spatial and polarisation modes with equal amplitude across the entire operating frequency bandwidth of the device (4.4 THz), with each frequency band (10.0675 to 50.3375 GHz) being orthogonal to all adjacent and nearby bands. This is a particularly challenging scenario which requires phase masks that have high spatial frequencies in both the spatial and spectral directions across the entire clear aperture of the device. This state is being generated by the SLM in some plane roughly corresponding with the input of the MMF. However it is the corresponding output state that is experimentally accessible and will be measured and tested against. Attempting to generate the same DFT spatiospectral state at the output of the MMF would in fact be easier, as the scattering in the fibre would make it very unlikely that all HG modes would need to be equally excited at all frequencies. The DFT matrix is implemented in the basis of the Hermite-Gaussian modes (output of the MPLC). Implementing in the basis of the ports of the 1D array facing the spectral pulse shaper (input of the MPLC) would be even more difficult, as the mismatch between the Fourier transform of the 1D array and the source beam coming from the laser on the SLM would be always be at its worst, across all frequency components, requiring a large amount of loss. This worst-case scenario also occurs in the HG basis for frequency bands where all modes are required to be generated in-phase. A scenario where the DFT is being implemented in both the 1D basis of the spots and the 2D basis of the HG modes simultaneously.

Supplementary Figure 11a illustrates the theoretical spatiospectral state being generated by the SLM. For the case of 50 GHz bands, there are coincidently as many spatial/polarisation modes as there are frequency bands across the 4.5 THz of the device. For all other bandwidths, illustrated in Supplementary Figure 12 to 15, the DFT matrix repeats across the frequency axis. Supplementary Figure 11b are the theoretical coefficients for the corresponding phase mask calculated as per the description above. Each spectral band is calculated to have a beam quality of 95%, but no attempt is made to equalise the overall loss of each frequency band. The differences in loss are inherent for a single plane hologram as discussed above. The worst-case can be seen at the far right of Supplementary Figure 11b, where all modes must be generated in phase. This scenario corresponds with the Fourier transform of the 1D spot array focusing to a single small spot on the SLM. This small spot is much smaller than the source beam coming from the laser. Theoretically due to this size mismatch, the maximum possible coupling between the two would be 14%. Supplementary Figure 11c is the corresponding output spatiospectral state that would be expected if we propagated the ideal state of Supplementary Figure 11a through the measured transfer matrices of the device and MMF. Supplementary Figure 11d is the measured output spatiospectral state. The complicated nature of the output states makes comparison by eye of Supplementary Figure 11c and 11d difficult. However it is summarised in Supplementary Figure 11e, as the normalised field correlation squared between the expected and measured spatial/polarisation output states as a function of frequency. That is, the proportion of the power that is measured to be in the expected spatial state at each wavelength. Ideally this would be 1.0 for all frequencies, or more realistically, 0.95 for all frequencies as that was the specified quality for convergence in the Gerchberg-Saxton algorithm used to calculate the phase masks. The green series of Supplementary Figure 11e follows the centre frequencies of each band and is relatively flat at approximately 0.9 from 191.0 THz to 195.4 THz. For this particular test, the worst-case in-phase spectral band is just outside the operating bandwidth of the device. The red series is the full measurement at the resolution of the swept-wavelength digital holography system (10.0675 GHz). The dips in response are to be expected, as these correspond with the spatial/polarisation state transitioning between two orthogonal states from one frequency band to the next. This transition must occur over some frequency bandwidth due to the physical size of the beam on the SLM along the

spectral axis, as well as the finite spatial resolution of the SLM itself. Supplementary Figure 11f is the total impulse response of the ideal and measured output states corresponding with Supplementary Figure 11d and 11e. There is good correlation between the two, especially at the smaller magnitude delays, consistent with the delay-dependent loss characterisation of Supplementary Figure 9c.

In Supplementary Figure 12 to 15 it can be seen that the correlation between the ideal and measured states in either the spectral or temporal domains, are well correlated for the 40 and 30 GHz bands at mostly in excess of 80%. Some degradation starts to be observed at 20 GHz, and quality is poor for 10 GHz. This is to be expected as the size of the beam on the SLM is approximately 15 GHz in the spectral direction. Looking closely at the spectral correlations, it can also be seen that there are dips in correlation for the green series at frequencies which correspond with the worst-case of in-phase mode generation. Testing narrower bands in the frequency domain, corresponds with testing longer delays in the time domain. This can be seen in the temporal responses of Supplementary Figure 11f to 15f, as an increasingly long impulse response as the spectral bands become narrower. In the time domain plots of Supplementary Figure 11f to 15f, there is a high correlation between the ideal and measured responses within the delay interval of approximately -20ps to +10ps. For delays outside this range, performance degrades as the finite spectral resolution in the plane of the SLM, makes these larger delays difficult to obtain accurately. This is consistent with the delay-dependent loss characterisation of Supplementary Figure 9c.
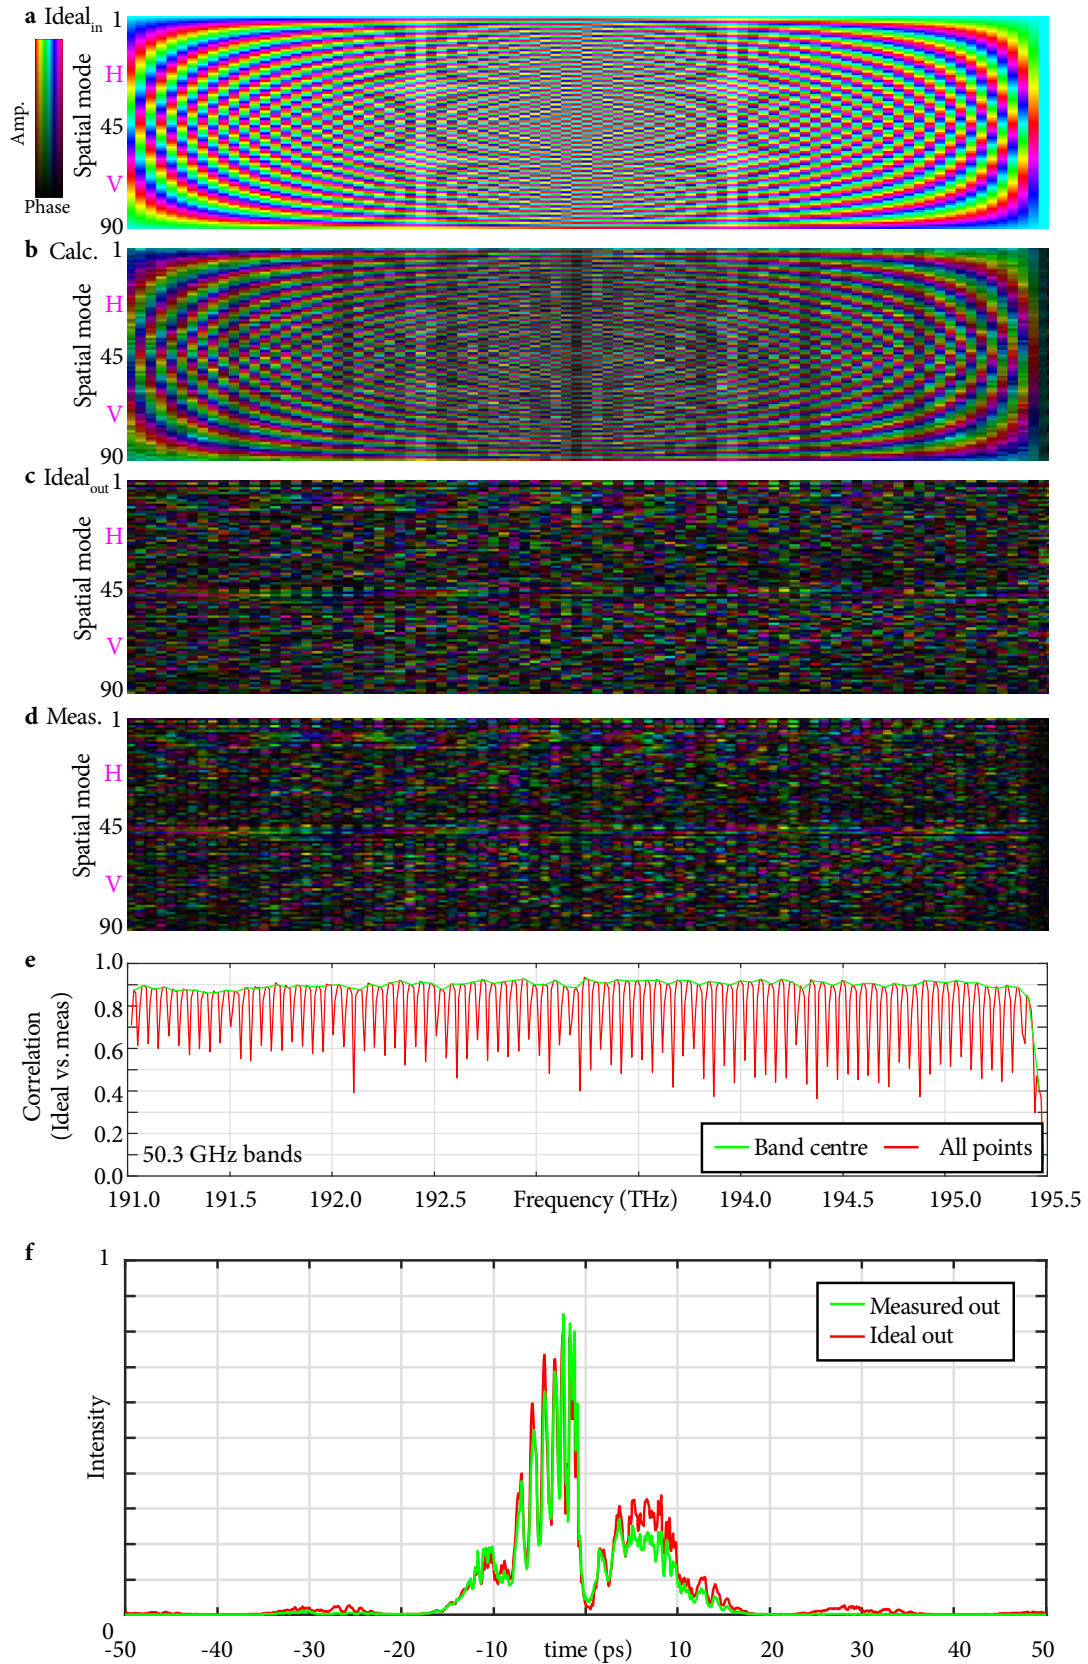
Supplementary Figure 11 | Characterisation of discrete Fourier transform spatiospectral states generated by the SLM for 50 GHz spectral bands.
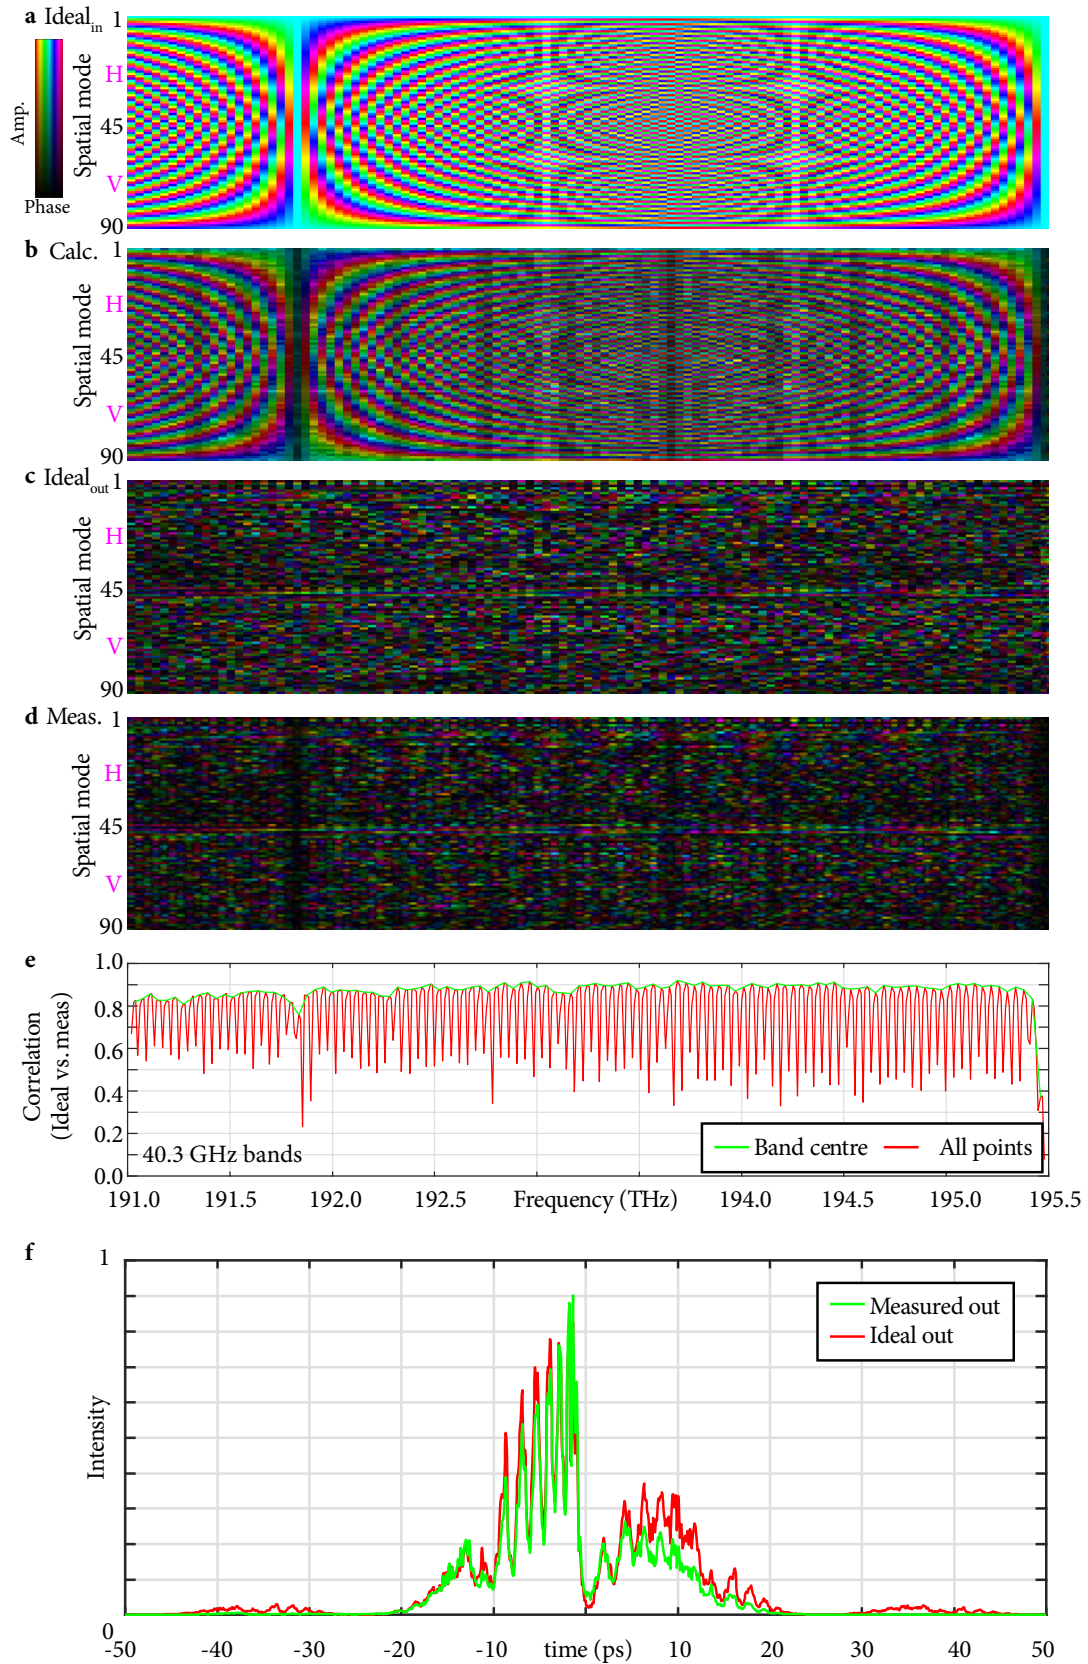
Supplementary Figure 12 | Characterisation of discrete Fourier transform spatiospectral states generated by the SLM for 40 GHz spectral bands.
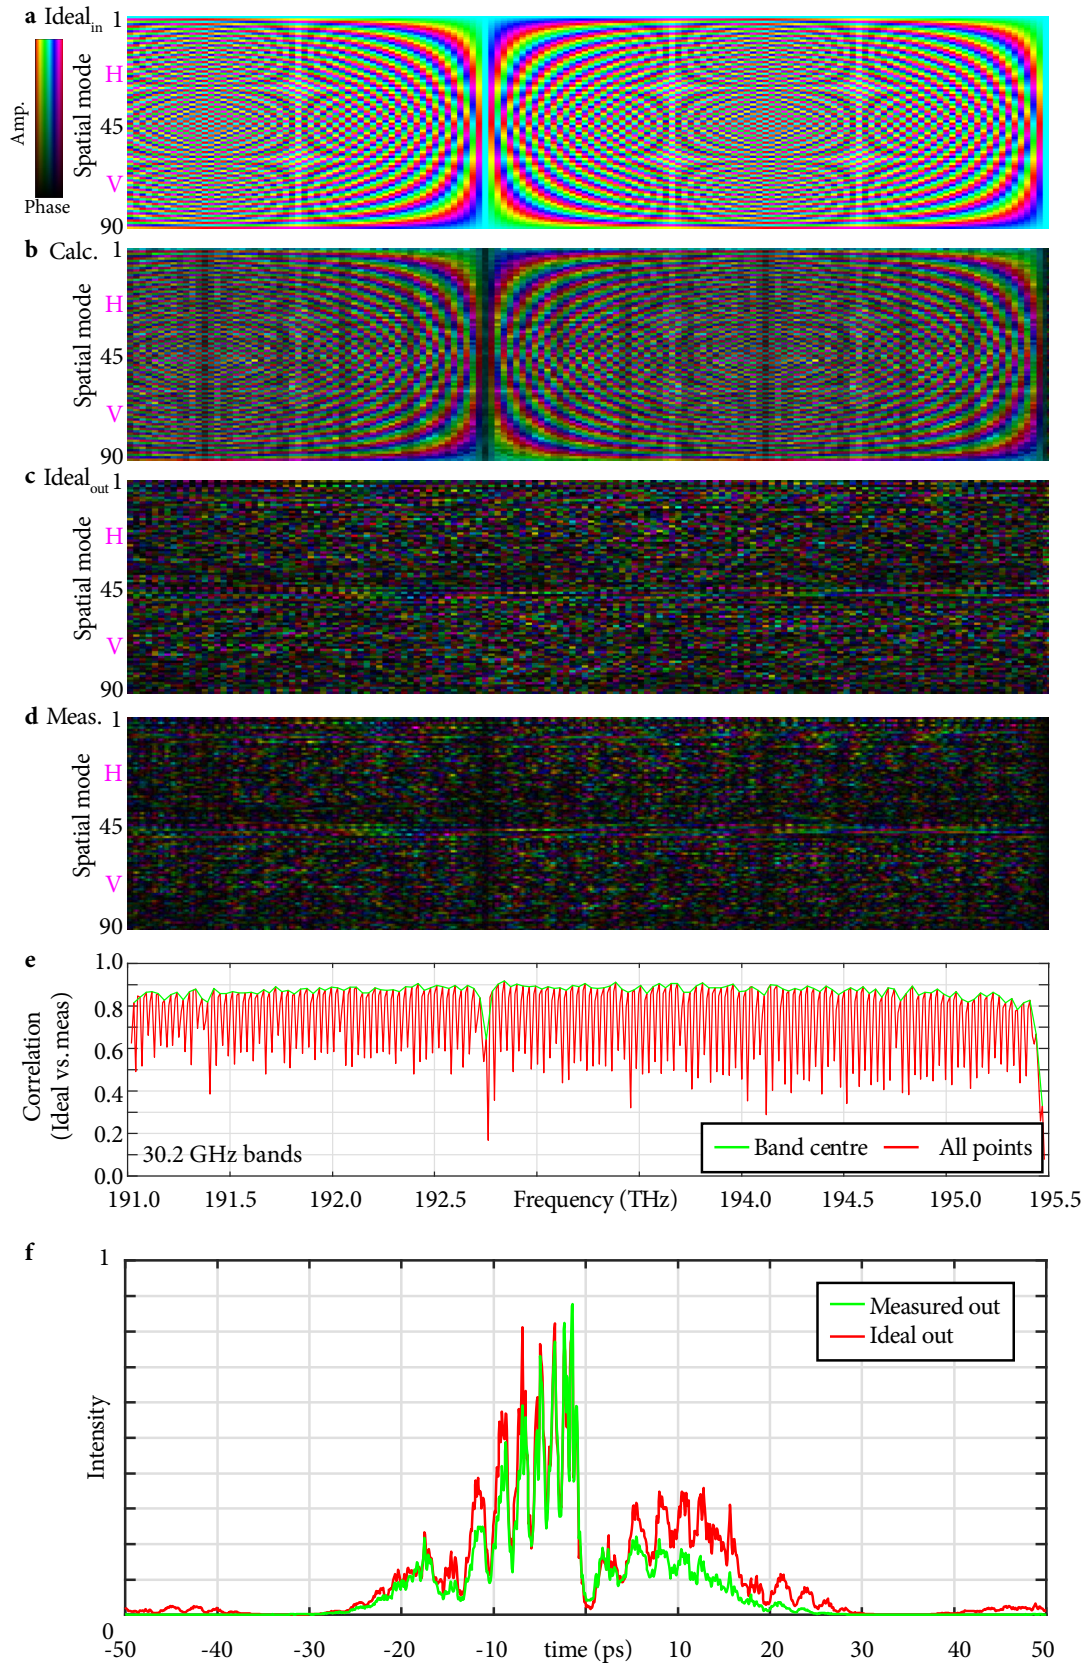
Supplementary Figure 13 | Characterisation of discrete Fourier transform spatiospectral states generated by the SLM for 30 GHz spectral bands.
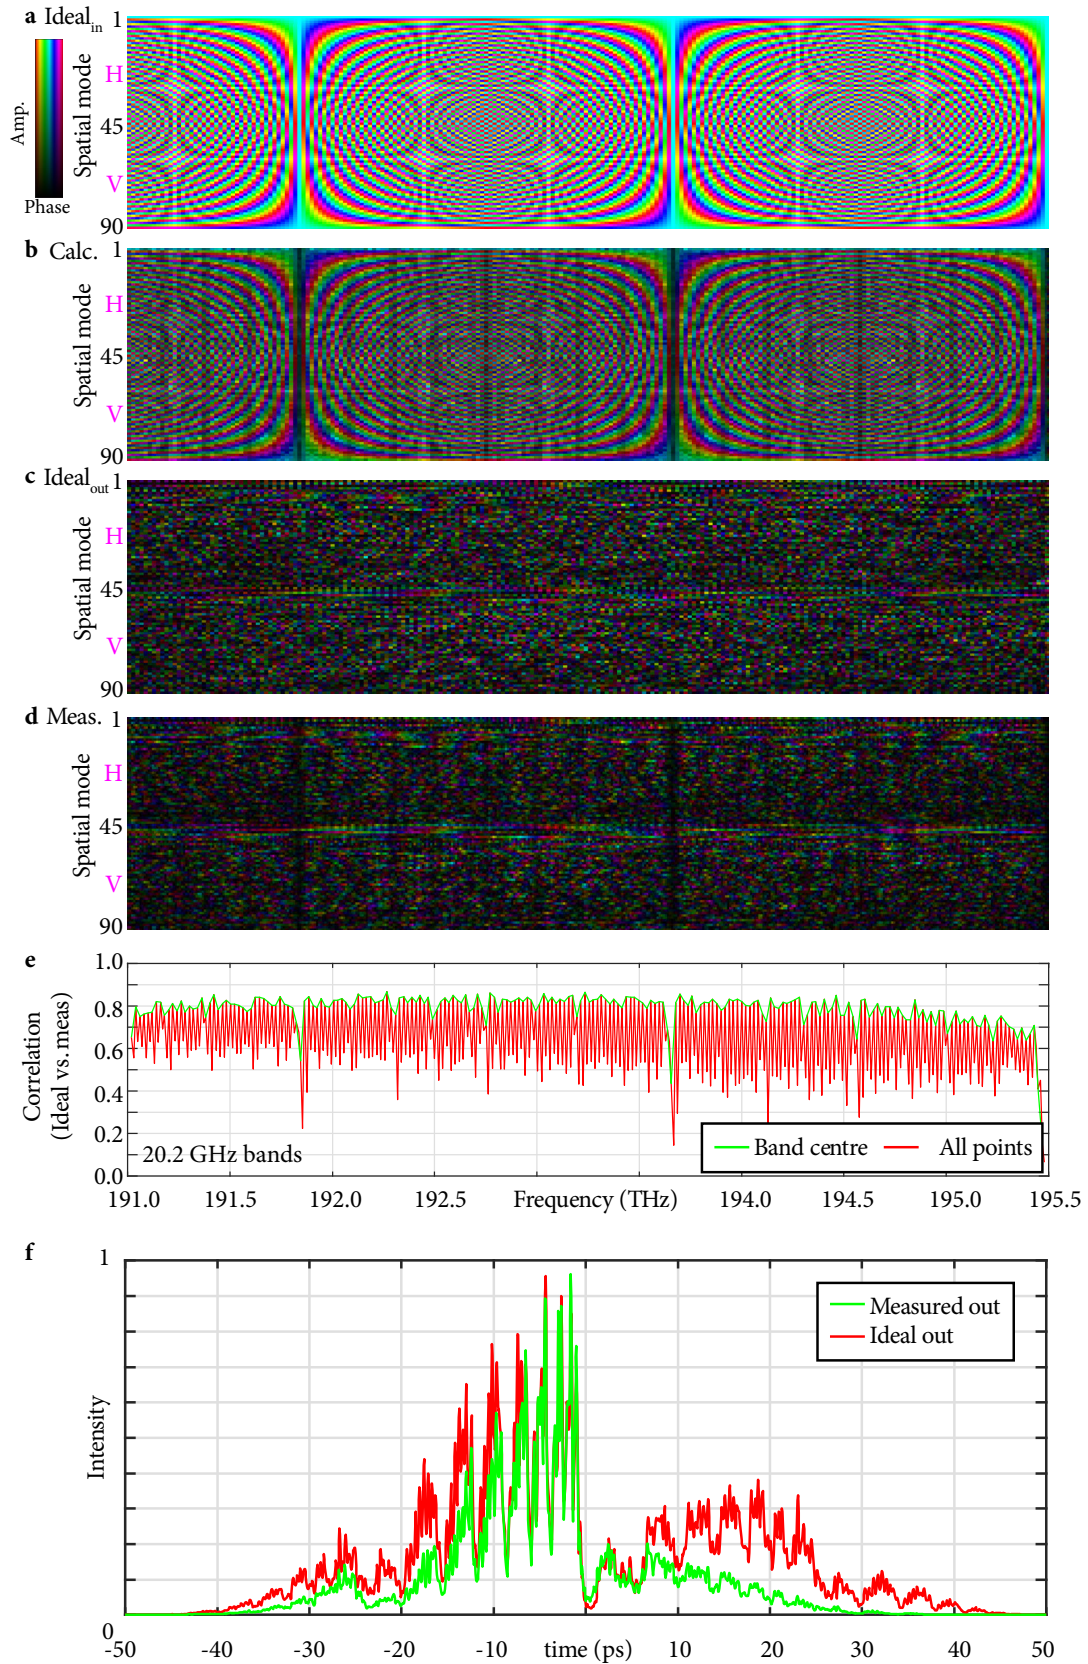
Supplementary Figure 14 | Characterisation of discrete Fourier transform spatiospectral states generated by the SLM for 20 GHz spectral bands.
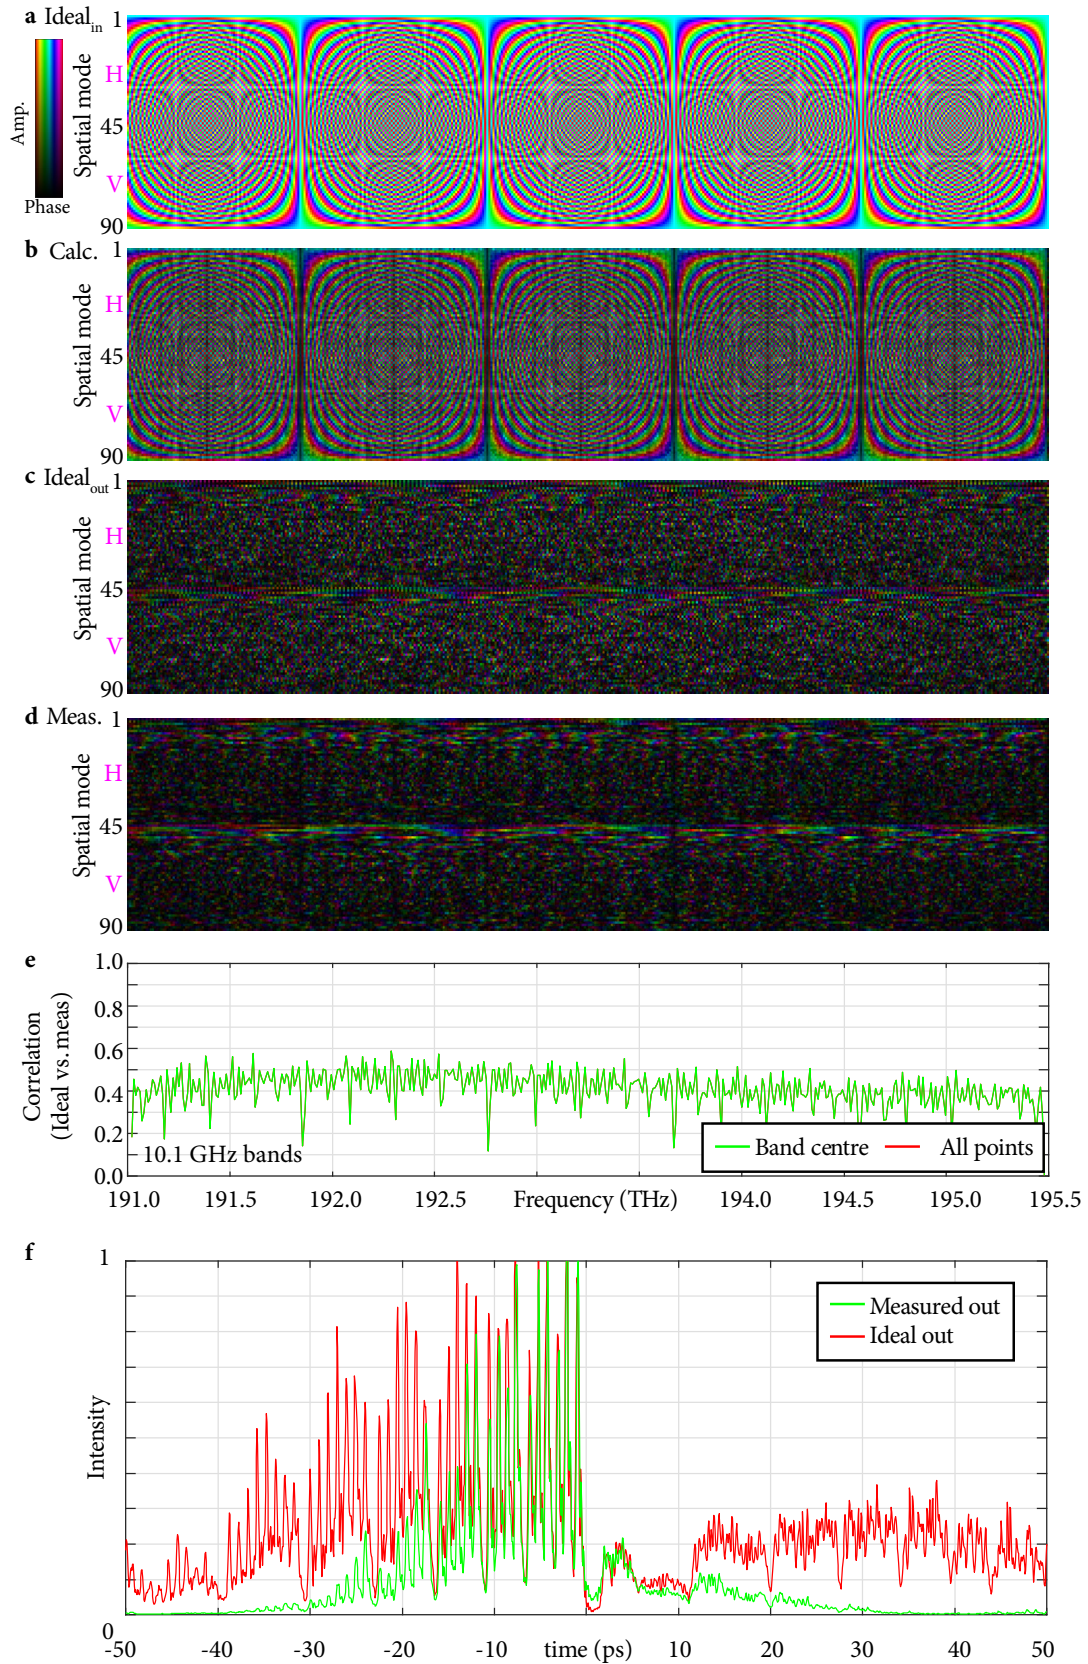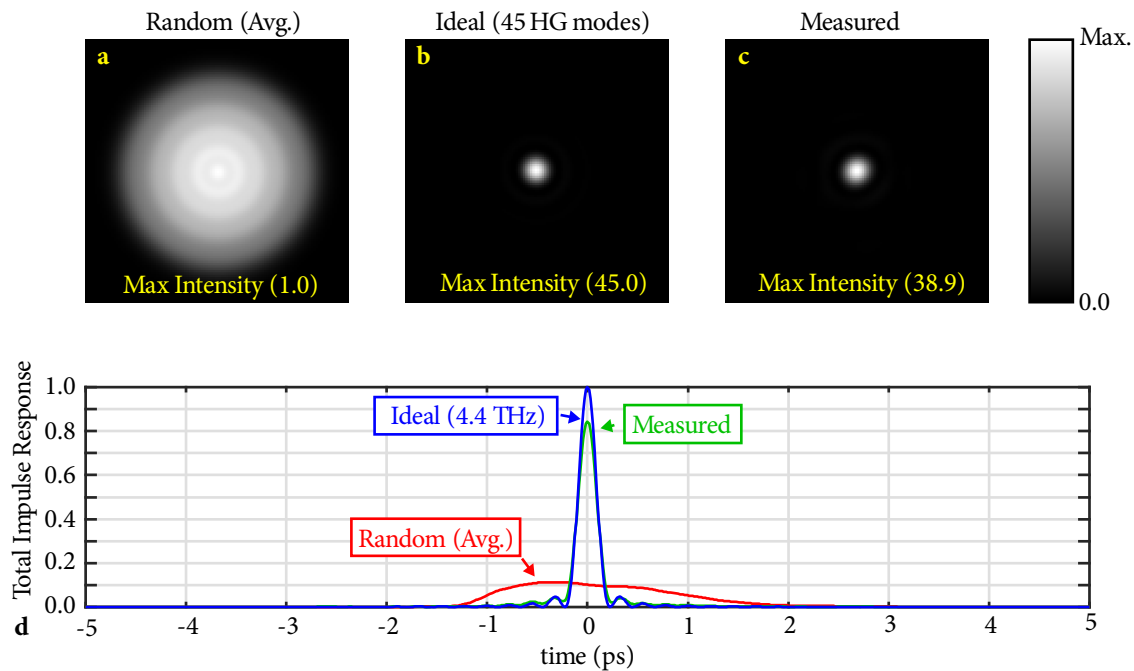
Supplementary Figure 15 | Characterisation of discrete Fourier transform spatiospectral states generated by the SLM for 10 GHz spectral bands.

# Supplementary Note 6 : Spatial and temporal intensity enhancement example

Supplementary Figure 16 | Example of spatial and temporal intensity enhancement. a, Average spatial intensity over all random output spatial states. **b,** Output spatial state with theoretical peak intensity possible using 45 HG modes. **c,** Measured output spatial state generated at end of MMF. **d,** Impulse response for average random output spatial state (red), theoretical peak intensity for bandwidth limited pulse (blue) and measured impulse response (green).

In this demonstration, a 45 degree linearly polarised, diffraction-limited, and bandwidth-limited spatiotemporal focus is generated at the end of the MMF. The corresponding intensity enhancement is then compared with a random output state, as well as the theoretical maximum possible intensity given the number of spatial modes and spectral bandwidth supported by the system. Supplementary Figure 16a is the reference average random output spatial state. Averaged over all possible superpositions of HG modes, which is equal to the sum of the intensities of each of the 45 HG basis spatial modes. Supplementary Figure 16 is a diffraction-limited spot as approximated by a superposition of the 45 HG basis modes. This is theoretically the most intense focus which can be created using the basis modes. Supplementary Figure 16c is the corresponding spatial state measured at 0 ps at the output of the MMF. When the spatial states are normalised to unity total power, the peak intensity of the spot is 38.9 times the average random state, approximately 86% of the theoretical maximum of 45. In this comparison polarisation has been averaged, meaning it does not contribute to the measure of intensity enhancement. If the metric was instead defined as the intensity enhancement at the output of the MMF in a given polarisation, an additional factor of two would be gained as the average random state power would be divided amongst two polarisations.

Supplementary Figure 16d is a similar illustration to Supplementary Figure 16a-c in the temporal rather than spatial domain. The average random impulse response, illustrated in red, is over all spatial and polarisation modes with equal power at the output of the MMF. This is the temporal response of the spatial state of Supplementary Figure 16a. An impulse response similar to Supplementary Figure 9b, but normalised for equal output power rather than input power. The ideal response shown in blue is a sinc pulse, corresponding with the full bandwidth of the system (4.46 THz), and has a peak intensity at 0 ps that is 9.85 times higher than the average random state. The measured response shown in green has a peak intensity of 8.31 times the average state or 84% of the theoretical maximum for a bandwidth limited pulse.

Supplementary Note 7 : Spatial/polarisation state fidelity
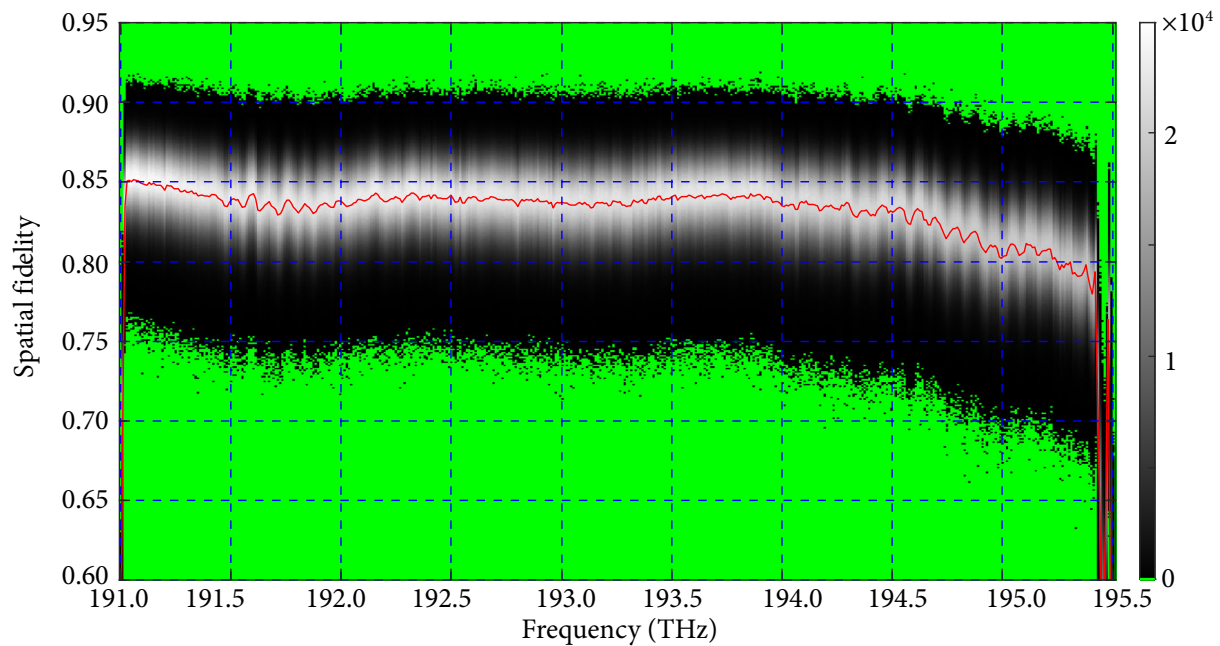
Supplementary Figure 17 | Distribution of fidelity for random spatial/polarisation target states. Given the measured transfer matrices for the device, one million random spatial/polarisation states per ~10 GHz spectral band are propagated back and forth through the system to yield a measure the achievable fidelity of spatial/polarisation states when maximising the power delivered to the target state, given the presence of mode dependent losses. Green indicates exactly zero states over all trials in that specific 0.001 wide histogram bin. Red line indicates average fidelity.

For the results presented in this work, a given target spatial/polarisation state is specified for each frequency component at the distal end of the multimode fibre. These states are numerically back-propagated using the conjugate transpose of the measured transfer matrices, to find the corresponding input state. For an ideal system, these transfer matrices would be unitary to within some overall loss common to all spatial and polarisation modes. All spatial/polarisation modes would have the same loss, and the conjugate transpose of the transfer matrix would also be its inverse. For such a system, ignoring any additional imperfections associated with generating superposition states as discussed previously in the section 'Spatiospectral stress test', a desired spatial/polarisation target state could be generated with maximum power and perfect fidelity. Theoretically, for an imperfect system with mode dependent loss (Supplementary Figure 8), "back-propagation" using the inverse of the transfer matrix will yield an input state capable of generating the output state with perfect fidelity, at the expense of some additional loss. Using the conjugate transpose of the transfer matrix for back-propagation as we do in this work, has the opposite trade-off. It yields the input state which would deliver the maximum amount of power to the target state, at the expense of overall fidelity.

Supplementary Figure 17 is a histogram which illustrates the degradation of spatial fidelity due to mode dependent loss when using numerical back-propagation by the conjugate transpose of the measured transfer matrices. For each ~10GHz band, one million random spatial/polarisation states are generated numerically. For a random state, each of the 90 modes is assigned a random amplitude between 0 and 1, and a random phase, and then normalised to unity total power. These states are then back-propagated through the multimode fibre to the input using the conjugate transpose of the transfer matrices, and then forward propagated using the transfer matrices to return to the distal end of the multimode fibre. This state is then renormalised to unity total power, and overlapped with the target state to yield the spatial fidelity. That is, the proportion of the output power which is in the target state. The results of Supplementary Figure 17 are linked to the results presented in Supplementary Figure 6 and 8, which also relate the mode dependent losses. The histogram of Supplementary Figure 17 indicates the number of random states per million trials per 10 GHz frequency band, which were found in each 0.001 wide fidelity bin. The mean fidelity within the 4.4 THz operating band is 0.83, and is indicated for each frequency with the red line. The distribution of fidelities is not

Gaussian, skewing towards low fidelities, but approximating as a normal distribution yields a standard deviation of 0.04. Highlighted in green, are histogram bins which recorded exactly zero states over all trials.

Supplementary Note 8 : Frequently Asked Questions

Isn't time reversal just phase conjugation?

Sometimes the terms are used interchangeably, however using the terminology we prefer[15], time reversal represents broadband control, whereas traditional phase conjugation is CW or narrowband. Traditional phase conjugation has no temporal features and is spatial-only. Consisting of a single spatial wavefront for all delays. Time reversal described in the frequency domain is broadband phase conjugation[15]. Each spectral component has its own wavefront, which in turn creates spatiotemporal features. Time reversal is spatiotemporal, whereas CW phase conjugation is only spatial.

How is this device different from previous wavefront shaping experiments?

This is the first optical wavefront shaping device which can arbitrarily and independently control all light's degrees of freedom; 2D space, time/frequency and polarisation. There is no coupling between these degrees of freedom and control of each maps to a unique position on the SLM. For example, assigning a certain amplitude and phase to any spatial/polarisation mode, does not define the properties of any other spatial or spectral component.

There have been many demonstrations over the years with some partial ability to control the spatial, polarisation and temporal aspects of an output beam. However due to the way those devices work, some types of beams are not possible, as components of the beam that would need to be independently controlled are not independently controllable. Imagine tracing the different spatial/polarisation and spectral components of the target beam back through the optical system from the target to the wavefront controlling device(s), typically an SLM. For previous demonstrations, for an arbitrary target, there would be at least two components of the beam, for example spatial modes and spectral modes, that must be controlled independently in the plane of the SLM in order to generate the target, yet both end up mapped to the same position on the SLM. Attempting to control one, also controls the other, and the target field is not accessible.

For example, some demonstrations make use of a highly scattering media which strongly couples the spatial and temporal properties[16,17]. Hence, by controlling only the spatial properties at the input, both spatial and temporal properties at the output can be manipulated. However again, this is only partial control. For an arbitrary spatiotemporal output field, the corresponding time reversed input field will be spatiotemporal; a non-separable function of both space and time. Yet only the subset of target fields that are spatiotemporal at the output, but purely spatial at the input can be accessed by these methods. Our system does not rely on spatiotemporal coupling in the complex media itself to address the temporal degree of freedom. Hence the device can operate in uncoupled scenario like free-space, all the way to a completely coupled scattering media. The limitations of our system are in terms of the maximum amount of spatiotemporal detail (addressable number of spatial and temporal/spectral modes), but not in terms of the types of coupling scenarios (none, weak, strong) it can handle.

Another common trade-off is a system which uses some dispersive element like a grating, to map a spatial dimension to the spectral degree of freedom which is in turn controlled by the 2D surface of an SLM[18–20]. In these cases, one spatial degree of freedom is not independently controllable from the spectral degree of freedom. To take a specific example, if we look at Fig. 1d of Sun et al.[19], we can see that one spatial dimension at the sample plane, maps directly to one spectral dimension. Each spectral component incident on the sample has its own incident angle, and the two components, spatial and spectral, cannot be controlled completely independently. Thinking of the system in reverse, trace an arbitrary wavelength back from the sample at an arbitrary angle. Trace the same wavelength forward from the laser source. The forward and backward paths will not overlap on the SLM for an arbitrary incident angle at the sample. Our system is similar to these approaches, however it use the 1D-to-2D spatial transformation of the MPLC[3] to get around this limitation. The MPLC transformation allows the SLM to address the spatial degree of freedom as if it is a one-dimensional degree of freedom. The MPLC acts as a kind of 1D 'lookup table' from which 2D output beams can be selected from a 1D list. Our device generates beams which are spatiotemporal in both the near and far-field.

Similar to McCabe et al.[18] where the spectral components are recombined with a grating, but different from Sun et al[19] and Hernandez et al.[20] in which the spectral components are recombined using a lens and hence are only temporally short at the focus.

Why not physically back-propagate your target field?

Our approach is similar to techniques such as DORT[21,22] in time reversed acoustics, or transfer matrix based phase conjugation approaches for CW light[23]. Numerical back-propagation using a transfer matrix has many advantages over physical back-propagation, especially for sophisticated beam types. A numerically back-propagated beam is more experimentally accurate as it eliminates the source of error associated with physically generating complicated fields. For relatively simple beams types, which have the same spatial profile at all delays (space-time separable), such as spatiotemporal focus, physical generation is relatively straightforward and could be done using a single-mode fibre, perhaps coupled to a traditional spectral pulse shaper. However part of the point of this paper is the generation of arbitrary vector spatiotemporal beams. Not only the relatively simple separable beam types, but rather beams which can have any value of amplitude and phase in any polarisation state, at any position at any point in time. These are the class of beams required to create time reversed waves in the general case, and a class of beams not previously demonstrated. Generating these beam types is considerably more difficult than space-time separable beams like spatiotemporal focus, and our system presented here is the first system capable of doing so. Hence to physically back-propagate these beam types, two copies of our optical systems would need to be built; one to physically generate back-propagated waves, and one to generate forward-propagated waves. Even so, if there was an advantage to taking this approach, building two systems could be justified. However as the numerical back-propagation approach using a transfer matrix is more experimentally accurate, allows any field to be delivered to the target from a relatively small number of measurements, and is also applicable to non-reciprocal system such as fibre amplifiers containing isolators, there is not a strong use-case for physical back-propagation. Numerical back-propagation is more accurate, has less optical components and is applicable in more experimental scenarios.

Why do you generate your spatiotemporal fields in the frequency domain rather than time domain?

Ultrafast pulse manipulation in optics is typically performed in the frequency domain, as the bandwidths required are beyond the capabilities of electronics. Many applications for which we'd expect a device such as this to be useful, would have an optical input from a source such as a pulsed laser.

An alternative approach would be a time domain implementation, perhaps using large numbers of Mach Zehnder modulators driven electrically. This would have similarities to a mode division multiplexing type system in optical telecommunications[8], but where the phase between spatial modes is carefully controlled at the transmitter and processing is performed optically at the transmitter side, rather than digitally at the receiver side. An advantage of this approach would be the ability to generate very long impulse responses, limited mostly by computer memory. The disadvantages would be comparatively low bandwidth (GHz rather than THz), and a potentially very complicated optical system. A time domain implementation in optics would be similar to a mode division multiplexing transmission system in optical fibre communications[8], but with special attention to the phase relationship between spatial channels.

Why do you characterise in the frequency domain not the time domain?

As mentioned above, our system implements its functionality using a frequency domain approach. As such, it makes sense to characterise the device in that same domain. Our optical characterisation apparatus based on swept-wavelength digital holography, is similar to a multi-port vector network analyser (VNA) in microwave systems, with all the usual advantages when it comes to impulse response measurements. Specifically, it has superior dynamic range, more accurate calibration and large flat bandwidth. In the frequency domain, we measure 444 CCD images over a span of 4.46 THz. One image for each spectral slice. For a spatiotemporal focus type demonstration, each of these spectral slices has the same power and uses the full dynamic range of the camera. The equivalent time domain characterisation would consist of a sequence of 443 blank images, with a single frame containing the information of interest. In terms of calibration, similar to a VNA, corrections are frequency dependent rather than delay dependent

which increases experimental accuracy when operating in the frequency domain. For example, removing the dispersion difference between the reference arm of the interferometer and the device-under-test is trivial in the frequency domain. An additional advantage is that each measurement in the frequency domain also corresponds with a particular position on the SLM in the device itself. This is particularly important for building and debugging devices such as these.

Having said all that, characterising in the frequency domain is not always an option, and time domain can become necessary. This is either because the system being characterised is nonlinear, or because of other experimental considerations. For example, if operating at 800 nm, a high quality swept-frequency laser source may not be available, yet a high-quality pulsed source is. Leaving no choice but to characterise in the time domain.

This system uses a multimode optical fibre, is it applicable to free-space/other scattering/complex media?

Yes, in much the same way that a traditional single-mode fibre attached spectral pulse shaper can be used in either context. Designing a fibre-coupled device naturally results in free-space compatibility, however a free-space device is not automatically compatible with fibre. Stated another way, free-space supports all the spatial modes a fibre does, but a fibre does not support all the modes free-space does. A convenient aspect of our device is it addresses the spatial dimension of the beam in the Hermite-Gaussian basis. Which makes it well suited to both fibre and free-space applications. HG modes are not only eigenmodes of parabolic index fibres in the weak-guidance approximation, but also eigenfunctions of the Fourier transform, solutions to the paraxial wave equation and eigenmodes of laser resonators. As opposed to a spatial basis such as grids of spots in real or k-space, the HG basis allows light to be focused anywhere in the near or far-field, without dead-space. The HG basis also losslessly couples to fibres, whereas grids in real or k-space do not. It is also a more efficient basis for addressing a finite aperture[24] compared with simple grids.

Traditional spectral pulse shapers often have single-mode fibre outputs, which acts as a convenient conduit through which the light can be routed to the part of the optical system, where it is ultimately required. It is a fibre-coupled device, applicable to both free-space and fibre applications. Our MMF attached spectral pulse shaper can be thought of in a similar fashion. Just as some additional dispersion compensation may be applied to a traditional spectral pulse shaper to compensate for the attached SMF, additional spatiotemporal dispersion compensation can be applied to our spectral pulse shaper to compensate for the attached MMF if required. As with a traditional single-mode spectral pulse shaper, this assumes the power levels involved are compatible with fibre and are not sufficient to cause unwanted nonlinear effects. However the thresholds would be much higher in this case as due to the use of MMF rather than SMF.

The compatibility with optical fibre could also be important for future applications of spatiotemporal beam shaping for nonlinear optics.

Supplementary References

1. Weiner, A. M. Femtosecond pulse shaping using spatial light modulators. *Rev. Sci. Instrum.* **71**, 1929–1960 (2000).

2. Baxter, G. *et al.* Highly programmable wavelength selective switch based on liquid crystal on silicon switching elements. *Optical Fiber Communication Conference, 2006 and the 2006 National Fiber Optic Engineers Conference. OFC 2006* 3 pp. (2006). doi:10.1109/OFC.2006.215365

3. Fontaine, N. K. *et al.* Packaged 45-Mode Multiplexers for a 50μm Graded Index Fiber. in *2018 European Conference on Optical Communication (ECOC)* 1–3 (2018). doi:10.1109/ECOC.2018.8535302

4. Sillard, P. *et al.* 50 μm Multimode Fibers for Mode Division Multiplexing. *J. Light. Technol.* **34**, 1672–1677 (2016).

5. Fontaine, N. K. *et al.* Laguerre-Gaussian mode sorter. *Nat. Commun.* **10**, 1865 (2019).

6. Carpenter, J. Laguerre-Gaussian mode sorter (Technical Version). (2019). Available at: https://www.youtube.com/watch?v=TqAr4UFtoqw.

7. Carpenter, J., Eggleton, B. J. & Schröder, J. Complete spatiotemporal characterization and optical transfer matrix inversion of a 420 mode fiber. *Opt. Lett.* **41**, 5580–5583 (2016).

8. Ryf, R. *et al.* High-Spectral-Efficiency Mode-Multiplexed Transmission Over Graded-Index Multimode Fiber. in *2018 European Conference on Optical Communication (ECOC)* 1–3 (2018). doi:10.1109/ECOC.2018.8535536

9. Saleh, B. E. A. & Teich, M. C. *Fundamentals of Photonics.* (John Wiley & Sons, Inc., 1991). doi:10.1002/0471213748

10. Prada, C., Manneville, S., Spoliansky, D. & Fink, M. Decomposition of the time reversal operator: Detection and selective focusing on two scatterers. *J. Acoust. Soc. Am.* **99**, 2067–2076 (1996).

11. Yang, G., Dong, B., Gu, B., Zhuang, J. & Ersoy, O. K. Gerchberg–Saxton and Yang–Gu algorithms for phase retrieval in a nonunitary transform system: a comparison. *Appl. Opt.* **33**, 209–218 (1994).

12. Plöschner, M. & Čižmár, T. Compact multimode fiber beam-shaping system based on GPU accelerated digital holography. *Opt. Lett.* **40**, 197–200 (2015).

13. Fontaine, N. K. *et al.* Laguerre-Gaussian mode sorter supplementary information. (2019). doi:10.14264/uql.2019.81

14. Lazarev, G., Chen, P.-J., Strauss, J., Fontaine, N. & Forbes, A. Beyond the display: phase-only liquid crystal on Silicon devices and their applications in photonics [Invited]. *Opt. Express* **27**, 16206–16249 (2019).

15. Derode, A., Tourin, A. & Fink, M. Time reversal versus phase conjugation in a multiple scattering environment. *Ultrasonics* **40**, 275–280 (2002).

16. Aulbach, J., Gjonaj, B., Johnson, P. M., Mosk, A. P. & Lagendijk, A. Control of light transmission through opaque scattering media in space and time. *Phys. Rev. Lett.* **106**, 103901 (2011).

17. Rotter, S. & Gigan, S. Light fields in complex media: Mesoscopic scattering meets wave control. *Rev. Mod. Phys.* **89**, 15005 (2017).

18. McCabe, D. J. Spatio-temporal focusing of an ultrafast pulse through a multiply scattering medium. *Nat. Commun* **2**, 447 (2011).

19. Sun, B. *et al.* Four-dimensional light shaping: manipulating ultrafast spatiotemporal foci in space and time. *Light Sci. Appl.* **7**, 17117 (2018).

20. Hernandez, O. *et al.* Three-dimensional spatiotemporal focusing of holographic patterns. *Nat. Commun.* **7**, 11928 (2016).

21. Prada, C. & Fink, M. Eigenmodes of the time reversal operator: A solution to selective focusing in multiple-target media. *Wave Motion* **20**, 151–163 (1994).

22. Lehman, S. K. & Devaney, A. J. Transmission mode time-reversal super-resolution imaging. *J. Acoust. Soc. Am.* **113**, 2742–2753 (2003).

23. Popoff, S. M. *et al.* Measuring the Transmission Matrix in Optics: An Approach to the Study and Control of Light Propagation in Disordered Media. *Phys. Rev. Lett.* **104**, 100601 (2010).

24. Zhao, N., Li, X., Li, G. & Kahn, J. M. Capacity limits of spatially multiplexed free-space communication. *Nat. Photonics* **9**, 822 (2015).