Ultrahigh bandwidth applications of microcombs: Neural networks and optical data transmission

M.Tan,¹ X.Xu,² B.Corcoran,² J.Wu,¹ A. Boes,³ T.Nguyen,³ S.Chu,⁴ B.Little,⁵ R.Morandotti,^{6,7} A.Mitchell,³ and D.J. Moss¹

¹Optical Sciences Centre, Swinburne University of Technology, Hawthorn, VIC 3122, Australia

²Department of Electrical and Computer Systems Engineering, Monash University, Clayton, 3800 VIC, Australia

³School of Engineering, RMIT University, Melbourne, VIC 3001, Australia

⁴Department of Physics, City University of Hong Kong, Tat Chee Avenue, Hong Kong, China.

⁵Xi'an Institute of Optics and Precision Mechanics Precision Mechanics of CAS, Xi'an, China.

⁶INRS-Énergie, Matériaux et Télécommunications, 1650 Boulevard Lionel-Boulet, Varennes, Québec, J3X 1S2, Canada.

⁷Inst. of Fundamental and Frontier Sciences, University of Electronic Science and Technology of China, Chengdu 610054, PR China

ABSTRACT

We report ultrahigh bandwidth applications of Kerr microcombs to optical neural networks and to optical data transmission, at data rates from 44 Terabits/s (Tb/s) to approaching 100 Tb/s. Convolutional neural networks (CNNs) are a powerful category of artificial neural networks that can extract the hierarchical features of raw data to greatly reduce the network complexity and enhance the accuracy for machine learning tasks such as computer vision, speech recognition, playing board games and medical diagnosis [1-7]. Optical neural networks can dramatically accelerate the computing speed to overcome the inherent bandwidth bottleneck of electronics. We use a new and powerful class of micro-comb called soliton crystals that exhibit robust operation and stable generation as well as a high intrinsic efficiency with an extremely low spacing of 48.9 GHz. We demonstrate a universal optical vector convolutional accelerator operating at 11 Tera-OPS/s (TOPS) on 250,000 pixel images for 10 kernels simultaneously — enough for facial image recognition. We use the same hardware to sequentially form a deep optical CNN with ten output neurons, achieving successful recognition of full 10 digits with 900 pixel handwritten digit images. We also report world record high data transmission over standard optical fiber from a single optical source, at 44.2 Terabits/s over the C-band, with a spectral efficiency of 10.4 bits/s/Hz, with a coherent data modulation format of 64 QAM. We achieve error free transmission across 75 km of standard optical fiber in the lab and over a field trial with a metropolitan optical fiber network. Our work demonstrates the ability of optical soliton crystal micro-combs to exceed other approaches in performance for the most demanding practical optical communications applications.

Keywords: Optical neural networks, neuromorphic processor, microcomb, convolutional accelerator, data transmission

1. INTRODUCTION

Artificial neural networks, collections of nodes with weighted connections can, with proper feedback to adjust the network parameters, “learn” and perform complex operations for face recognition, speech translation, playing board

Figure 1 | Operation principle of the Tera-FLOPS photonic convolution accelerator. EOM: electro-optical Mach-Zehnder modulator. SMF: single mode telecommunications fibre.

games and medical diagnosis [1-4]. While classic fully connected feedforward networks face challenges in processing extremely high-dimensional data, convolutional neural networks (CNNs), inspired by the biological behavior of the visual cortex system, can abstract the representations of input data in their raw form, and then predict their properties with both unprecedented accuracy and greatly reduced parametric complexity [5]. CNNs have been widely applied to computer vision, natural language processing and other areas [6, 7].

The capability of neural networks is dictated by the computing power of the underlying neuromorphic hardware. Optical neural networks (ONNs) [8-18] are promising candidates for next-generation neuromorphic computation, since they have the potential to overcome the bandwidth bottleneck of their electrical counterparts [6, 19-22] and achieve ultra-high computing speeds, enabled

by the >10 THz wide optical telecom band [14]. ONNs are attracting a great deal of attention with recent breakthroughs and reviews [13-18]. Operating in analog frameworks, they avoid the limitations imposed by the energy and time consumed during reading and storing data back and forth, known as the von Neumann bottleneck [19]. Significant progress has been made in highly parallel, high-speed and trainable ONNs [8-18, 23-27], including approaches that have the potential for full integration on a single photonic chip [8,12,14,15], in turn offering an ultra-high computational density. However, there are still opportunities for significant improvements in ONNs. Processing large-scale data, for practical real-life computer vision tasks, remains challenging because they are primarily fully connected structures where their input scale is determined solely by hardware parallelism. This leads to tradeoffs between the network scale and footprint. Our work represents the first report of the extremely high computing speeds that ONNs are capable of.

Here, we report ultra-high bandwidth applications of optical Kerr soliton crystal microcombs [28 – 87] including the demonstration of an optical convolution accelerator to process and compress large-scale data [13, 14] at 11 TeraOP/s rates, as well as world record high bandwidth optical data transmission at 44Tb/s from a single chip source [34]. Both achievements are performed through the use of soliton crystal Kerr microcombs. The ONN is achieved by interleaving wavelength, temporal, and spatial dimensions with the microcomb. We achieve a vector computing speed as high as 11.322 TOPS and use it to process 250,000 pixel images with 10 convolution kernels at 3.8 TOPs. The convolution accelerator is fully and dynamically reconfigurable, and scalable, and can serve as both a convolutional accelerator front-end with multiple and simultaneous parallel kernels, as well as forming an optically deep CNN with fully connected neurons, with the same hardware. We demonstrate a CNN and successfully apply it to the recognition of full ten digit (0-9) handwritten images, achieving an accuracy of 88%. Our optical neural network represents a major step towards realizing monolithically integrated ONNs and is enabled by our use of an integrated microcomb chip. Moreover, our accelerator scheme is stand alone and universal — fully compatible with either electrical or optical interfaces. Hence, it can serve as a universal ultrahigh bandwidth data compressing front end for any neuromorphic hardware — either optical or electronic — making massive-data machine learning for real-time, ultrahigh bandwidth data possible.

Our record data transmission speed of 44.2 Terabits/s (Tb/s) across standard optical fibre with a single optical source, also achieves a very high spectral efficiency of 10.4 bits/s/Hz. Spectral efficiency is critically important since it directly governs how much total bandwidth can be realized in a system. As for the CNN demonstration, we use soliton crystals for the optical source. Soliton crystals are very stable and robust and can be easily generated. Further, they have a very high intrinsic conversion efficiency that, together with the low soliton micro-comb FSR of 48.9 GHz, enabled us to use a high coherent modulation data format of 64 quadrature amplitude modulation (QAM). We demonstrate error free data transmission across 75 km of standard optical fibre in our lab and in a field trial in an installed metropolitan area optical fibre testbed network in the Melbourne region. Our work directly highlights the capability of Kerr microcombs to outperform other sources for optical communications systems.

2. PRINCIPLE OF OPERATION

Figure 1 shows the operation principle of the photonic convolutional accelerator (CA), featuring high-speed electrical signal input and output data ports, while Figure 2 shows a detailed experimental configuration. The data vector input X is serially encoded with the intensity of temporal symbols in an electrical waveform at a symbol rate $1/\tau$ (baud), where τ is the symbol period. The convolution kernel is likewise represented by a weight vector W of length R that is used to encode the optical power of the microcomb lines by spectral shaping with a Waveshaper. The temporal waveform X is then multi-cast onto the kernel wavelength channels via electro-optical modulation, generating the replicas weighted by W . Next the optical waveform is transmitted through a dispersive delay with a delay step between adjacent wavelength channels equal to the symbol duration of X , thus achieving time and wavelength interleaving. Finally, the delayed and weighted replicas are summed via high speed photodetection so that each time slot yields a convolution between X and W for a given convolution window, or receptive field. Thus, the convolution window effectively slides at the modulation speed matching the baud rate of X . Each output symbol is the result of R multiply-and-accumulate operations, with the computing speed given by $2R/\tau$ OPS. Since the speed of this process scales with both the baud rate and number of wavelengths, it can be dramatically boosted into the TOP/s regime by using the massively parallel wavelength channels of a microcomb. Further, the input data X length is unlimited - the convolution accelerator can process arbitrarily large-scale data, limited only by the electronics. Likewise, the kernels number and length are arbitrary, limited only by the number of wavelengths. We achieve simultaneous convolution of multiple kernels by adding additional sub-bands of R wavelengths for each kernel. Following multicasting and dispersive delay, the sub-bands (kernels) are demultiplexed and detected separately with high speed photodetectors, generating a separate electronic waveform for each kernel.

While the convolutional accelerator typically processes vectors, it can operate on matrices for image processing by flattening the matrix into a vector. The precise way that this is done determines both the sliding convolution window's stride and the equivalent matrix computing speed. Our flattening method sets the receptive field (convolution slot) to slide with a horizontal stride of unity (ie., every matrix input element has a corresponding convolution output) and a vertical stride that scales with the size of the convolutional kernel. The larger vertical stride effectively results in sub-sampling across the vertical direction of the raw input matrix, equivalent to a partial pooling function [88] in addition to the convolution. This results in an effective reduction (or overhead) in matrix computing speed that scales inversely with the size of the kernel, so that a 3 x 3 kernel results in a speed reduction overhead by 1/3. While this can be eliminated by a variety of means to produce convolutions with a symmetric stride and hence no speed overhead, this is actually not necessary for most applications. Finally, this approach is highly flexible and reconfigurable without any change in hardware - we use same system for the convolutional accelerator for image processing as well as to form an optical deep learning CNN which we use to perform a separate series of experiments. The convolutional accelerator hardware forms both the input processing stage as well as the fully connected neuron layer of the CNN. The system can achieve matrix multiplication by simply sampling a single time slot of the output waveform, since the vector dot product is equivalent to the special convolution case where the two input vectors $\mathbf{X}$ and $\mathbf{W}$ have the same length.

Figure 3 shows a detailed example of the photonic convolution accelerator operating in two different modes. The left panel shows the system performing convolution operations, that are used for the large stand-alone convolution image processing and the convolutional layer of the CNN. The right panel shows the system performing matrix operations which are used as the fully connected layer of the optical CNN. Considering that the experimentally demonstrated configurations are too complex to be presented clearly, in Figure 3 we show a simplified configuration of input data and weights to illustrate the operation principle of our system. The length of $\mathbf{W}$ and $\mathbf{X}$ shown in this figure are $R = 4$ and $L = 13$ for the case of convolution operations, and $R = L = 4$ for the fully connected layer for matrix operations, respectively.

The schematic of the TOPS photonic convolution accelerator is illustrated in the left panel of Figure 3. The input data vector (length L) and weight vector (length R) is first multiplexed in the time and wavelength domains, respectively. The input data vector is represented by the intensities of the temporal symbols in a stepwise electrical waveform $\mathbf{X}[n]$ (n denotes discrete temporal locations of the symbols, $n \in [1, L+R-1]$), where $\mathbf{X}[n]$ is the electrical input of the accelerator. The weight vector of the kernel is imprinted onto the optical power of the shaped comb lines as $\mathbf{W}[R-i+1]$, at the i^{th} wavelength channel ($i \in [1, R]$, where i increases with wavelength). The input electrical waveform $\mathbf{X}[n]$ is first broadcast onto the shaped comb lines via electro-optical modulation. Thus the weighted replica at the i^{th} wavelength channel is $\mathbf{W}[R-i+1] \cdot \mathbf{X}[n]$. Next, the optical signals across all wavelengths are progressively shifted in the time domain via an optical time-of-flight buffer, which provides a wavelength-sensitive (dispersive) delay with a delay step τ (the difference in delay between adjacent wavelengths) equal to the symbol duration (inverse of the Baud rate) of $\mathbf{X}[n]$. Hence, the shifted replica becomes $\mathbf{W}[R-i+1] \cdot \mathbf{X}[n-i]$. Finally, the replicas of all wavelengths are summed via photo-detection as

$$\mathbf{Y}[n] = \sum_{i=1}^R \mathbf{W}[R-i+1] \cdot \mathbf{X}[n-i] = (\mathbf{W} * \mathbf{X})[n] \quad (1)$$

where each calculated symbol $\mathbf{Y}[n]$ within the range of $[R+1, L+1]$ denotes the dot product between $\mathbf{W}$ and a certain region of $\mathbf{X}$ (this region is defined by the sliding receptive field as $[n-R : n-1]$ or $[n-R, n-R+1, n-R+2, \dots, n-1]$). By simply reading different time slots of the output signal, a convolution is achieved between the weight vector and the input data, thus generating extracted feature maps (matrix convolution outputs) of the input image. While higher order dispersion in the dispersive delay can, in principle, degrade performance, in our experiments this was not a factor.

In addition, the convolution accelerator can also perform matrix multiplication operations, as illustrated in the right panel of Figure 3. The matrix multiplication operations can be treated as a special case of convolution operations when the two input vectors (the pooled and flattened feature maps, and the flattened synaptic weights for the fully connected layer) are the same length ($R=L$). Figure 3 shows an example with $R=L=4$. Here, we assume the input data vector $\mathbf{X}_{FC}[n]$ and the weight vector $\mathbf{W}_{FC}[R-i+1]$ both have a length R ($i \in [1, R]$, $n \in [1, R]$). Thus, according to Eq. 1, the output waveform after photodetection is

$$\mathbf{Y}_{FC}[n] = \sum_{i=1}^R \mathbf{W}_{FC}[R-i+1] \cdot \mathbf{X}_{FC}[n-i] \quad (2)$$

By sampling at the time slot denoted by $n=R+1$, the matrix multiplication result of the two input vectors is therefore

$$\mathbf{Y}_{FC}[R+1] = \sum_{i=1}^R \mathbf{W}_{FC}[R-i+1] \cdot \mathbf{X}_{FC}[R+1-i] = \sum_{i=1}^R \mathbf{W}_{FC}[i] \cdot \mathbf{X}_{FC}[i] \quad (3)$$

3. EXPERIMENT

3.1 Optical soliton crystal micro-combs

Optical frequency combs, composed of discrete and equally spaced frequency lines, are extremely powerful for optical frequency metrology [28]. Micro-combs offer the full power of optical frequency combs, but in an integrated form with much smaller footprint [28-34]. They have enabled many breakthroughs in high-resolution optical frequency synthesis [32], ultrahigh-capacity communications [33, 34], complex quantum state generation [35 - 43], advanced microwave signal processing [67 - 87], and more. Figure 4 shows a schematic of our optical microcomb chip as well as typical

spectra and pumping curves. We use a class of microcomb called soliton crystals that have a crystal-like profile in the angular domain of tightly packed self-localized pulses within micro-ring resonators [34, 47, 48]. They form naturally in micro-cavities with appropriate mode crossings, without complex dynamic pumping or stabilization schemes (described by the Lugiato-Lefever equation [28, 46]). They are characterized by distinctive optical spectra (Fig. 4f) which arise from spectral interference between the tightly packaged solitons circulating along the ring cavity. Soliton crystals exhibit

deterministic generation arising from interference between the mode crossing-induced background wave and the high intra-cavity power (Fig. 4c). In turn this enables simple and reliable initiation via adiabatic pump wavelength sweeping [34] that can be achieved with manual detuning (the intracavity power during pump sweeping is shown in Fig. 4d). The key to the ability to adiabatically sweep the pump is that the intra-cavity power is over 30x higher than single-soliton states (DKS), and very close to that of spatiotemporal chaotic states [28, 34]. Thus, the soliton crystal has much less thermal detuning or instability arising from the ‘soliton step’ that makes resonant pumping of DKS states more challenging. It is this combination of ease of generation and conversion efficiency that makes soliton crystals highly attractive. The coherent soliton crystal microcomb (Figure 4) was generated by optical parametric oscillation in a single integrated MRR (Fig. 4a, 4b) fabricated in CMOS-compatible Hydex [22, 23, 34] glass, similar to silicon oxynitride, featuring a $Q > 1.5$ million, radius $592\ \mu\text{m}$, and a low FSR of ~ 48.9 GHz. The pump laser (Yenista Tunics – 100S-HP) was boosted by an optical amplifier (Pritel PMFA-37) to initiate the parametric oscillation. The soliton crystal microcomb yielded over 90 channels over the C-band (1540-1570 nm), offering adiabatically generated low-noise frequency comb lines with a small footprint $< 1\ \text{mm}^2$ and low power consumption ($> 100\ \text{mW}$ [34]).

Figure 5 | Image processing results. a. Kernel weights and shaped microcomb optical spectrum. b. The input electrical waveform of the image (grey lines are theory and blue experimental). c. The convolved results of the 4th kernel that performs top Sobel image processing. d. The weight matrices of the kernels and corresponding recovered images.

Figure 6. Experimental and simulated large scale facial image processing achieved by convolutional accelerator with 10 kernels.

3.2 Matrix Convolution Accelerator

Figure 2 shows the experimental setup for the full matrix convolutional accelerator that we use to process a classic 500×500 face image. The system performs 10 simultaneous convolutions with ten 3×3 kernels to achieve distinctive image processing functions. The weight matrices for all kernels were flattened into a composite kernel vector W containing all 90 weights (10 kernels with $3 \times 3 = 9$ weights each), which were then encoded onto the optical power of 90 microcomb lines by an optical spectral shaper (Waveshaper), each kernel occupying its own frequency band of 9 wavelengths. The wavelength channels were supplied by the soliton crystal microcomb described in the previous section, producing 90 wavelengths in the C-band (1540-1570 nm) at a spacing ~ 48.9 GHz over a ~ 36 nm bandwidth [34].

Figure 5 shows the experimental results of the image processing. Figure 5a depicts the kernel weights and shaped microcomb optical spectrum while the input electrical waveform of the image (grey lines are theoretical and blue

experimental waveforms) are in Figure 5b. Figure 5c displays the convolved results of the 4th kernel that performs a top Sobel image processing function (grey lines are theory and red experimental). Finally, Figure 5d shows the weight matrices of the kernels and corresponding recovered images.

The raw 500×500 input face image was flattened electronically into a vector X and encoded as the intensities of 250,000 temporal symbols with a resolution of 8 bits/symbol (limited by the electronic arbitrary waveform generator (AWG)), to form the electrical input waveform via a high-speed electrical digital-to-analog converter, at a data rate of 62.9 Giga Baud (time-slot $\tau=15.9$ ps) (Fig. 5b). The waveform duration was 3.975 μ s for each image corresponding to a processing rate for all ten kernels of over 1/3.975 μ s, equivalent to 0.25 million of these ultra-large-scale images per second.

The input waveform X was then multi-cast onto the 90 shaped comb lines via electro-optical modulation, yielding replicas weighted by the kernel vector W. Following this, the waveform was transmitted through ~2.2 km of standard single mode fibre with a dispersion ~17ps/nm/km. The fibre length was carefully chosen to induce a relative temporal shift in the weighted replicas with a progressive delay step of 15.9 ps between adjacent wavelengths, exactly matching the duration of each input data symbol τ , resulting in time and wavelength interleaving for all ten kernels.

The 90 wavelengths were then de-multiplexed into 10 sub-bands of 9 wavelengths, each sub-band corresponding to a kernel, and separately detected by 10 high speed photodetectors. The detection process effectively summed the aligned symbols of the replicas (the electrical output waveform of one of the kernels (kernel 4) is shown in Fig. 5c). The 10 electrical waveforms were converted into digital signals via ADCs and resampled so that each time slot of each of the waveforms corresponded to the dot product between one of the convolutional kernel matrices and the input image within a sliding window (i.e., receptive field). This effectively achieved convolutions between the 10 kernels and the raw input image. The resulting waveforms thus yielded the 10 feature maps (convolutional matrix outputs) containing the extracted hierarchical features of the input image (Figure 5d).

The convolutional vector accelerator made full use of time, wavelength, and spatial multiplexing, where the convolution window effectively slides across the input vector X at a speed equal to the modulation baud-rate — 62.9 Giga Symbols/s. Each output symbol is the result of 9 (the length of each kernel) multiply-and-accumulate operations, thus the core vector computing speed (i.e., throughput) of each kernel is $2 \times 9 \times 62.9 = 1.13$ TOPS. For ten kernels computed in parallel the overall computing speed of the vector CA is therefore $1.13 \times 10 = 11.3$ TOPS, or $11.321 \times 8 = 90.568$ Tb/s (reduced slightly by the optical signal to noise ratio (OSNR)). This speed is over 500 x the fastest ONNs reported to date.

For the image processing matrix application demonstrated here, the convolution window had a vertical sliding stride of 3 (resulting from the 3×3 kernels), and so the effective matrix computing speed was $11.3/3 = 3.8$ TOPs. Homogeneous strides operating at the full vector speed can be readily achieved by duplicating the system with parallel weight-and-delay paths (see below), although we found that this was unnecessary. While the length of the input data processed here was 250,000 pixels, the convolution accelerator can process data with an arbitrarily large scale, the only practical limitation being the capability of the external electronics.

To achieve the designed kernel weights, the generated microcomb was shaped in power using liquid crystal on silicon based spectral shapers (Finisar WaveShaper 4000S). We used two WaveShapers in the experiments - the first was used to flatten the microcomb spectrum while the precise comb power shaping required to imprint the kernel weights was performed by the second, located just before the photo-detection. A feedback loop was employed to improve the accuracy of comb shaping, where the error signal was generated by first measuring the impulse response of the system with a Gaussian pulse input and comparing it with the ideal channel weights. Figure 6 shows the experimental and theoretical large scale facial image processing results achieved by the matrix convolutional accelerator with ten convolutional kernels. The electrical input data was temporally encoded by an arbitrary waveform generator (Keysight M8195A) and then multicast onto the wavelength channels via a 40 GHz intensity modulator (iXblue). For the 500×500 image processing, we used sample points at 62.9 Giga samples/s to form the input symbols. We then employed a 2.2 km length of dispersive fibre that provided a progressive delay of 15.9 ps/channel, precisely matched to the input baud rate.

Considering that the convolutional accelerator fundamentally operates on vectors, for applications to image processing, the input data is in the form of matrices and so it needs to be flattened into vectors. (see video presentation). We follow a common approach where the raw input matrix is first sliced horizontally into multiple sub-matrices, each with a height equal to that of the convolutional kernel. The sub-matrices were then flattened into vectors and connected head-to-tail to form the desired vector. This flattening method equivalently makes the receptive field slide with a horizontal stride of 1 and a vertical stride equal to the height of the convolutional kernel. We note that a small stride (such as a horizontal stride of 1) ensures that all features of the raw data are extracted, while a large stride (3 or 5) reduces the overlap

Figure 7. Deep level CNN. Feature maps are the convolutional matrix outputs while the fully connected layers embody the neural network component.

Figure 8. The architecture of the optical CNN, including a convolutional layer, a pooling layer, and a fully connected layer.

between the sliding convolution windows and effectively subsamples the convolved feature maps, thus partially serving as a pooling function. A stride of 4 was used for the AlexNet [88]. We note that although the homogeneous strides are generally used more often in digitally implemented CNNs, inhomogeneous convolution strides (unequal horizontal and vertical strides) such as those used here are often used, and in most cases, including ours, do not limit the performance. Our performance was verified by the high recognition success rate of the CNN for full 10 digit prediction. Further, homogeneous convolutions can be achieved by duplicating the weight-and-delay paths (each including a modulator, a spool of dispersive fibre, a de-multiplexer and multiple photo-detectors) of the accelerator.

3.3 Deep Learning Optical Convolutional Neural Network

The convolutional accelerator architecture presented here is fully and dynamically reconfigurable and scalable with the same hardware system. We were thus able to use the accelerator to sequentially form both a front-end convolution processor as well as a fully connected layer, together yielding an optical deep CNN. We applied the CNN to the recognition of full 10 (0-9) handwritten digit images. Figure 7 shows the overall architecture of the deep (multiple) level CNN structure. The feature maps are the convolutional matrix outputs while the fully connected layers embody the neural network. Figure 8 shows the architecture of the optical CNN, including a convolutional layer, a pooling layer, and a fully connected layer. Figure 9 shows the detailed experimental schematic of the optical CNN. The left side is the input front end convolutional accelerator while the right is the fully connected layer - both the deep learning optical CNN. The microcomb supplies the wavelengths for both the convolution accelerator as well as the fully connected layer. The electronic digital signal processing (DSP) module used for sampling and pooling is external.

The convolutional layer (Fig. 9, left) performs the heaviest computing duty of the entire network, generally taking 55% to 90% of the total computing power. The digit images – 30×30 matrices of grey-scale values with 8-bit resolution – were flattened into vectors and multiplexed in the time-domain at 11.9 Giga Baud (time-slot $\tau = 84$ ps). Three 5×5 kernels were used, requiring 75 microcomb lines, resulting in a vertical stride of 5. The dispersive delay was achieved with ~ 13 km of SMF to match the data baud-rate. The wavelengths were de-multiplexed into the three kernels which were detected by high speed photodetectors and then sampled and nonlinearly scaled with digital electronics to recover the extracted hierarchical feature maps of the input images. The feature maps were then pooled electronically and flattened into a vector (Eq. 2,3) $\mathbf{X}_{FC}$ ($72 \times 1 = 6 \times 4 \times 3$) per image that formed the input data to the fully connected layer.

The fully connected layer had 10 neurons, each corresponding to one of the 10 categories of handwritten digits from 0 to 9, with the synaptic weights represented by a 72×10 weight matrix $\mathbf{W}_{FC}^{(l)}$ (ie., ten 72×1 column vectors) for the l^{th} neuron ($l \in [1, 10]$) – with the number of comb lines (72) matching the length of the flattened feature map vector $\mathbf{X}_{FC}$. The shaped optical spectrum at the l^{th} port had an optical power distribution proportional to the weight vector $\mathbf{W}_{FC}^{(l)}$, thus serving as the equivalent optical input of the l^{th} neuron. After being multicast onto the 72 wavelengths and progressively delayed, the optical signal was weighted and demultiplexed with a single Waveshaper into 10 spatial output ports — each corresponding to a neuron. Since this part of the network involved linear processing, the kernel wavelength weighting could be implemented either before the EO modulation or at a later stage just before photodetection. The advantage of the latter is that both the demultiplexing and weighting can then be achieved with a single Waveshaper. Finally, the different node/neuron outputs were obtained by sampling the 73rd symbol of the convolved results. The final output of the optical CNN was represented by the intensities of the output neurons (Figure 10), where the highest

Figure 9. Schematic of optical CNN. Left: input front end convolutional accelerator; Right: fully connected layer. Microcomb source supplies wavelengths for both. The electronic digital signal processing (DSP) used for sampling and pooling is external.

intensity for each tested image corresponded to the predicted category. The peripheral systems, including signal sampling, nonlinear function and pooling, were implemented electronically with digital signal processing hardware, although some of these functions (e.g., pooling) can be performed in the optical domain with the VCA. Supervised network training was performed offline electronically (see below).

We experimentally tested 50 x 8-bit resolution images each 30×30 of the handwritten digit dataset with the deep optical CNN. The confusion matrix (Figure 11) shows an accuracy of 88% for the generated predictions, in contrast to 90% for the numerical results calculated on an electrical digital computer. The computing speed of the CA component of the deep optical CNN was $2 \times 75 \times 11.9 = 1.785$ TOPS, or 14.3 Tb/s. To process image matrices with 5×5 kernels, the convolutional layer had a matrix flattening overhead of 5, yielding an image computing speed of $1.785/5 = 357$ Giga OPS. The computing speed of the fully connected layer was 119.8 Giga-OPS (see below). The waveform duration was $30 \times 30 \times 84 \text{ ps} = 75.6 \text{ ns}$ for each image, and so the convolutional layer processed images at the rate of $1/75.6 \text{ ns} = 13.2$ million handwritten digit images per second.

We note that handwritten digit recognition, although widely employed as a benchmark test in digital hardware, is still (for full 10 digit (0 - 9) recognition) beyond the capability of existing analog reconfigurable ONNs. Digit recognition requires a large number of physical parallel paths for fully-connected networks (e.g., a hidden layer with 10 neurons requires 9000 physical paths), which poses a huge challenge for current nanofabrication techniques. Our CNN represents the first reconfigurable and integrable ONN capable not only of performing high level complex tasks such as full handwritten digit recognition, but at TOP speeds.

For the convolutional layer of the CNN, we used 5 sample points at 59.421642 Giga Samples/s to form each single symbol of the input waveform, which also matched with the progressive time delay (84 ps) of the 13km dispersive fibre. The generated electronic waveforms for 50 images served as the electrical input signal for the convolutional and fully connected layers, respectively.

For the convolutional accelerator in both the CA and CNN experiments - the 500×500 image processing experiment and the convolutional layer of the CNN - the second Waveshaper simultaneously shaped and de-multiplexed the wavelength channels into separate spatial ports according to the configuration of the convolutional kernels. As for the fully connected layer, the second Waveshaper simultaneously performed the shaping and power splitting (instead of de-multiplexing) for the ten output neurons. The de-multiplexed or power-split spatial ports were sequentially detected and

measured. However, these two functions could readily be achieved in parallel with a commercially available 20-port optical spectral shaper (WaveShaper 16000S, Finisar) and multiple photodetectors.

The negative channel weights were achieved using two methods. For the 500×500 image processing experiment and the convolutional layer of the CNN, the wavelength channels of each kernel were separated into two spatial outputs by the WaveShaper according to the signs of the kernel weights, and then detected by a balanced photodetector (Finisar XPDV2020). Conversely, for the fully connected layer the weights were encoded in the symbols of the input electrical waveform during the electrical digital processing stage. Both of these methods to impart negative weights were successful. Finally, the electrical output waveform was sampled and digitized by a high-speed oscilloscope (Keysight DSOZ504A, 80 Giga Symbols/s) to extract the final convolved output. For the CNN, the extracted outputs of the convolution accelerator were further processed digitally, including rescaling to exclude the loss of the photonic link via a reference bit, and then mapped onto a certain range using a nonlinear tanh function.

The pooling layer’s functions were also implemented digitally, following the algorithm introduced in the network model. The residual discrepancy between experiment and calculations, for both the recognition and convolving functions, was due to the deterioration of the input waveform caused by performance limitations of the electrical arbitrary waveform generator. Addressing this would lead to greater accuracy and closer agreement with numerical calculations.

3.4 Network training and digital processing

For the deep learning (multiple level) optical CNN, we employed datasets from the MNIST (Modified National Institute of Standards and Technology) handwritten digit database [89]. The dataset contained 60000 images as the training set and 10000 images as the test set. The structure of the CNN in this work (Figure 7) was determined empirically using trial-and-error, which is a standard approach for neural networks. In our case this was greatly aided by the fact that the network structure (number of synapses and neurons) can be reconfigured dynamically without any change in hardware. The 28×28 input data was first padded with zeros into a 30×30 image and then sliced into a 5×180 matrix and convolved with the 5×5 kernels. This slicing operation equivalently made the receptive field slide horizontally with a stride = 1 across the rows and a vertical stride = 5 across the columns of the 30×30 input data (corresponding to the 900 input nodes). Then the $6 \times 26 \times 3$ feature map was pooled (using average pooling) to a smaller dimension of $6 \times 4 \times 3$. Finally, the matrix was further flattened into a 72×1 vector that served as input nodes for the fully connected layer, which in turn generated the predictions using the 10 output neurons. The nonlinear function we used after the convolutional layer, the pooling function and the fully connected layer was the tanh function. Although other nonlinear functions such as ReLU are widely used, we used this tanh function since it can be realized with a saturating electrical amplifier.

The training necessary to acquire pre-trained weights and biases was performed offline with a digital computer. The Back Propagation algorithm [90] was employed to adjust the weights. To validate the hyper-parameters of the CNN, we performed a 10-fold cross validation using the 60000 samples of the training dataset, where the training set was separated into 10 subsets and each was then used to test the trained network (6000 samples) with the rest of the 9 subsets

(54000 samples). The test sets were assessed by both the optical CNN (50 images) and an electronic computer (10000 images) for comparison.

Figure 6 shows the experimental and simulated large scale facial image processing results achieved by the convolutional accelerator with ten convolutional kernels. It shows the experimental results of large 500×500 face image processing, including the recorded waveforms and the recovered images. Figure 12 shows the fully connected layer architecture and experimental results. The left panel depicts the experimental setup, similar to the convolutional layer. The right panel shows top: the experimental results for one output neuron, including the shaped comb spectrum; middle: the pooled feature maps of the digit “3” and the corresponding input electrical waveform (the grey lines are theory and red experimental; bottom: output waveform of the neuron and sampled intensities. The supplementary materials (SM) of Ref [14] shows the full experimental results of the CNN. Including the shaped impulse response of the convolutional layer that has 3 kernels and 75 wavelengths, or weights, in total. Also shown in the SM is a figure depicting the shaped impulse responses for the ten neurons, each with 72 synapses, at the fully connected layer. The fifty handwritten digits tested during our experiments are shown in the SM, with their corresponding encoded electrical waveform, which served as the electrical input of the convolutional layer. The electronic waveform generated from the extracted feature maps served as the input of the fully connected layer. Ref [14] SM shows the full experimental results of the CNN including the recorded waveforms and the recovered/sampled

Figure 12. Fully connected layers. Architecture and experimental results. The left panel depicts the experimental setup, similar to the convolutional layer. The right panel shows top: the experimental results for one output neuron, including the shaped comb spectrum; middle: the pooled feature maps of the digit “3” and the corresponding input electrical waveform (the grey lines are theory and red experimental; bottom: output waveform of the neuron and sampled intensities.

outputs including the experimental results of the ten output neurons in the fully connected layer.

Since there are no common standards in the literature for classifying and quantifying the computing speed and processing power of ONNs, we explicitly outline the performance definitions that we use in characterizing our performance. We follow the approach that is widely used to evaluate electronic micro-processors. The computing power of the convolution accelerator—closely related to the operation bandwidth—is denoted as the throughput, which is the number of operations performed within a certain period. Considering that in our system the input data and weight vectors originate from different paths and are interleaved in different dimensions (time, wavelength, and space), we use the temporal sequence at the electrical output port to define the throughput in a more straightforward manner.

At the electrical output port, the output waveform has $L+R-1$ symbols in total (L and R are the lengths of the input data vector and the kernel weight vector, respectively), among which $L-R+1$ symbols are the convolution results. Further, each output symbol is the calculated outcome of R multiply-and-accumulate operations or $2R$ OPS, with a symbol duration τ given by that of the input waveform symbols. Thus, considering that L is generally much larger than R in practical convolutional neural networks, the term $(L-R+1)/(L+R-1)$ would not affect the vector computing speed, or throughput, which (in OPS) is given by

$$\frac{2R}{\tau} \cdot \frac{L-R+1}{L+R-1} \approx \frac{2R}{\tau} \quad (4)$$

As such, the computing speed of the vector convolutional accelerator demonstrated here is $2 \times 9 \times 62.9 \times 10 = 11.321$ Tera-OPS for ten parallel convolutional kernels).

We note that when processing data in the form of vectors, such as audio speech, the effective computing speed of the accelerator would be the same as the vector computing speed $2R/\tau$. Yet when processing data in the form of matrices, such as for images, we must account for the overhead on the effective computing speed brought about by the matrix-to-vector flattening process. The overhead is directly related to the width of the convolutional kernels, for example, with 3-by-3 kernels, the effective computing speed would be $\sim 1/3 * 2R/\tau$, which still is in the TOP regime due to the high parallelism brought about by the time-wavelength interleaving technique.

For the convolutional accelerator, the output waveform of each kernel (with a length of $L-R+1=250,000-9+1=249,992$) contains $166 \times 498 = 82,668$ useful symbols that are sampled out to form the feature map, while the rest of the symbols are discarded. As such, the effective matrix convolution speed for the experimentally performed task is slower than the

vector computing speed of the convolution accelerator by the overhead factor of 3, and so the net speed then becomes

$$11.321 \times 82,668 / 249,992 = 11.321 \times 33.07\% = 3.7437 \text{ TOPS.}$$

For the deep CNN the convolutional accelerator front end layer has a vector computing speed of $2 \times 25 \times 11.9 \times 3 = 1.785$ TOPS while the matrix convolution speed for 5×5 kernels is $1.785 \times 6 \times 26 / (900 - 25 + 1) = 317.9$ Giga-OPS. For the fully connected layer of the deep CNN, according to Eq. (4), the output waveform of each neuron would have a length of $2R-1$, while the useful (relevant output) symbol would be the one locating at $R+1$, which is also the result of $2R$ operations. As such, the computing speed of the fully connected layer would be $2R / (\tau * (2R-1))$ per neuron. With $R = 72$ during the experiment and ten neurons simultaneous

Table I

Performance Comparison for Optical Neural Networks

Parameter Approach	Input data dimension	Computing speed		Scalability & reconfigurability*	Integrated components
		FLOPS	Bit resolution [†]		
Lens [11]	784	<1	8-bit	Level 1	None
Coupler array [8]	4	<1	N.A.	Level 2	Weight and sum circuits
Reservoir [9]	371 (up to 1113)	~ 17.6 G	N.A.	Level 2	None
Phase change material [12]	15	<1	1-bit	Level 2	Weight and sum circuits
Broadcast and weight [21]	8	~ 1 G	N.A.	Level 3	None
Time and wavelength interleaving (this work)	250,000 (theoretically unlimited)	11.321 T	8-bit	Level 4	90-wavelength light source

Bit resolution[†]: this parameter, adopted from digital protocols, can roughly reflects the intensity resolution of analog ONNs. For example, analog ONNs with bit resolution = 8-bit can process input data with $2^8=256$ discrete levels, such as the benchmark MNIST handwritten digits. This parameter here for our work and other existing ONN is determined by the bit-resolution of the input data, and for some work is not available (N. A.) since the relevant details of the tested dataset is not provided.

Scalability & reconfigurability*:

Level 1—can hardly reconfigure the synaptic weights; Level 2—the synaptic weights can be reconfigured, but the network structure, including the number of layers and neurons in each layer, can hardly be reconfigured; Level 3—the weights and network structure can be reconfigured, but the computing function of each layer cannot be reconfigured. Level 4—the weights and network structure can be dynamically reconfigured, the function of each layer can also be reconfigured (such as from convolutional layers to fully connected layers).

operating, the effective computing speed of the matrix multiplication would be $2R / (\tau^*(2R-1)) \times 10 = 2 \times 72 / (84 \text{ps} \times (2 \times 72 - 1)) = 119.83 \text{ Giga-OPS}$.

In addition, the intensity resolution (bit-resolution for digital systems) for analog ONNs is mainly limited by the signal-to-noise ratio (SNR). To achieve 8-bit resolution, the SNR of the system needs to be $> 20 \cdot \log_{10}(28) = 48 \text{ dB}$. This was achieved by our accelerator and so our speed in Tb/s is close to the speed in OPs $\times 8$ – not reduced by our OSNR.

3.5 Performance comparison

Here, we review recent progress of optical neuromorphic hardware (Table 1). This section is not comprehensive but focuses on the leading results that address the most crucial technical issues for optical computing hardware. The input data dimension directly determines the complexity of the processing task. In real-life scenarios, the input data dimension is generally very large, for example, a human face image would require over 60,000 pixels. Thus, to make optical computing hardware eventually useful, the input data dimension would need to be at least over 20,000. In this work we demonstrate processing of images containing 250,000 pixels, which is 224 x higher than previous reports.

The computing speed is perhaps the most important parameter for computing hardware and is the main strength of optical approaches. Although there has not been a widely accepted definition of optical hardware computing speed, the key issue is the number of data sets that are processed within a certain time period - i.e., how many images can be processed per second. As such, although in some approaches [8, 11, 12], the latency is low due to the short physical path lengths, the computing speed remains very low due to the absence of high-speed data interfaces (i.e., input and output nodes are not updated at a high rate). Although other approaches [9, 27] offer high-speed data interfaces, their computing parallelism is not high and so their speed is similar to the input data rate. In our work, through the use of high-speed data interfaces (62.9 Giga Baud) and time-wavelength interleaving, we achieved a record computing speed of 11.321 Tera-OPS, $> 500 \times$ higher than previous reports.

Finally, the scalability and reconfigurability determines the versatility of the optical computing hardware. Approaches that cannot dynamically reconfigure the synapses [11] (marked as “Level 1” in the table) are barely trainable. Approaches at Level 2 [9, 12, 27] support online training, however, they can only process a specific task since the network structure is fixed once the device is fabricated. For approaches [27] at Level 3, different tasks can be processed although the function of each layer is fixed, which limits the hardware from implementing more complex operations other than matrix multiplication. Our work represents the first approach that operates at Level 4 with full dynamic reconfigurability in all respects. Here, the synaptic weights can be reconfigured by programming the WaveShaper. Further, the number of synapses per neuron can be reconfigured by reallocating the wavelength channels with the demultiplexer. The number of layers can be reconfigured by changing the number of stacked devices. Finally, the computing function can be switched between convolution and matrix multiplication by changing the sampling method. The degree of integration directly determines the potential computing density (processing capability per unit footprint). For approaches not well suited to integration [8, 11, 27], the potential computing density is low. While other approaches achieve limited integration of the weight and sum circuits [8, 12] - probably the most challenging issue — advanced integrated light sources have not been demonstrated. The performance of the light source directly determines the performance of the overall hardware in both input data scale [8] and number of synaptic connections per neuron [12]. The mm^2 sized microcomb offers a large number of precisely-spaced wavelengths, which enhances the overall parallelism and computing density, representing a major step towards the full integration of optical computing hardware.

4. DISCUSSION

This approach can be readily scaled in performance in terms of input data size, as well as network size and speed. The data size is limited in practice only by the memory of the electrical digital-to-analog converters, and so in principle it is possible to process 4K-resolution (4096×2160) images. By integrating 100 photonic convolution accelerators layers (still much less than the 65536 processors integrated in the Google TPU [21]), the optical CNN would be capable of solving much more difficult image recognition tasks at a vector computing speed of $100 \times 11.3 = 1.130 \text{ Peta-OPS}$. Further, the optical CNN presented here supports online training, since the optical spectral shaper used to establish the synapses can be dynamically reconfigured as fast as 500 ms or faster with integrated optical spectral shapers [91].

Although we had a non-trivial optical latency of $0.11 \mu\text{s}$ introduced by dispersive fibre spool, this did not affect the operational speed. Moreover, the latency of the delay function can be virtually eliminated (to $< 200 \text{ ps}$) by using integrated highly dispersive devices such as photonic crystals or customized chirped Bragg gratings [92] or even tunable dispersion compensators [93, 94]. Finally, current nanofabrication techniques can enable significantly higher levels of

integration of the convolutional accelerator. The micro-comb source itself is based on a CMOS compatible platform that is intrinsically designed for large-scale integration. Other components such as the optical spectral shaper, modulator, dispersive media, de-multiplexer and photodetector have all been realized in integrated form [91, 92, 95].

5. OPTICAL DATA TRANSMISSION AT 44 TERABITS/S

In this work we also use soliton crystal microcombs [47] to achieve a world record speed for data transmission across standard optical fibre with a single optical source [34]. We achieve a line rate of 44.2 Terabits/s (Tb/s) utilizing only the C-band, with a very high spectral efficiency of 10.4 bits/s/Hz. Spectral efficiency is critically important since it directly governs how much total bandwidth can be realized in a system. Soliton crystals display very stable and robust operation and generation as well as a very high intrinsic conversion efficiency that, all taken together with the low soliton micro-comb FSR of 48.9 GHz, enabled us to use a high coherent modulation data format of 64 quadrature amplitude modulation (QAM). We demonstrate error free data transmission across 75 km of standard optical fibre in our lab and in a field trial in an installed metropolitan area optical fibre testbed network in the Melbourne region. Our results were underpinned by the ability of soliton crystals to operate without stabilization or feedback control, but only open loop systems, significantly reducing the amount of instrumentation required.

Currently, 100's of terabits per second are transmitted every moment across the world's fibre optic networks and the global bandwidth is growing at a rate of 25% /yr [96]. Ultrahigh capacity data links that use parallel massive wavelength division multiplexing (WDM) systems combined with coherent advanced modulation formats [97], are critical to meet this demand. Space-division multiplexing (SDM) is another emerging approach where multiple signals are transmitted either over multiple core or multiple mode fibre, or both [98]. In parallel with all of this, there is a growing movement towards very short links but still with very high capacity, particularly for data centres. Even just ten years ago, long haul networks such as undersea links spanning thousands of kilometres, used to dominate the global infrastructure, but nowadays the demand has dramatically shifted towards smaller scale applications including data centres as well as metropolitan area networks (10s to 100s of kilometres in size). These trends demand highly compact, energy efficient and low-cost devices. Photonic integrated circuits are the only approach that can address these needs, where the optical source is absolutely key to each link, and therefore has the greatest need to meet these requirements. The ability to generate all wavelengths on a single chip that is both integrated and compact, replacing many lasers, will yield the highest benefits [99-101].

Kerr optical microcombs have attracted a great deal of interest and one of their main applications has been in this area. They have successfully been used as optical sources for ultra-high bandwidth optical fiber data transmission [33 - 34]. A key factor has been achieving the capacity to modelock all of the microcomb lines, and this has been characterized by the

discovery of new states of temporal optical soliton oscillation that include feedback-stabilized Kerr combs [102], dark solitons [101] and dissipative Kerr solitons (DKS) [33]. The last one (DKS) has achieved the greatest success, being the basis of extremely high data transmission rates across the full C and L telecom bands, at a rate of 30 Tb/s using only a single source, and 55 Tb/s using two microcombs [33]. Despite this, micro-combs need to be even more stable and simpler and robust in both operation and generation, in order to meet the demands of real-world installed fibreoptic systems [34, 97-101]. They particularly must work without the need for complicated stabilization feedback, preferably in uncomplicated open-loop fashion and without the need for complicated pumping

Figure 13. Field trial network in greater Metropolitan Melbourne.

schemes that DKS states need in order to be generated.

Furthermore, the conversion efficiency from pump to comb lines must be much higher and their threshold pump power much lower. Systems that use microcombs also must achieve a much higher spectral efficiency (SE) since to date they have only achieved about $\frac{1}{4}$ of the theoretical maximum. Spectral efficiency is an absolutely key and fundamental parameter that limits the total data capacity of systems [97, 98].

This paper reports a world-record high bandwidth for optical fibre data transmission using standard single mode fibre together with a single optical source. Our use of soliton crystals [47, 34], based on CMOS - compatible chips [28, 103, 34 - 87], enabled us to reach a transmission data rate of 44.2 Tb/s using only a single chip – an increase of almost 50% [33, 102]. More importantly, we report an improvement, by a factor of 3.7, in the SE, achieving 10.4 bits / s / Hz - a record high value for microcombs. We do this through the use of a high coherent modulation format of 64 QAM, together with a microcomb that has a very low spacing, or FSR, of 48.9 GHz. We only use the telecom C-band, leaving room for significant expansion in our capacity. We report experiments in the lab with 75 km of fibre as well as over an installed metropolitan optical fibre network (Figure 13). These results were made possible because of the highly and stable and robust generation and operation of the soliton crystals, together with their very high natural efficiency.

Soliton crystal oscillation states in micro-resonators that have a crystalline type of profile along the resonator path, forming in the angular domain of tightly packed self-localized pulses within micro-ring resonators [47]. They can occur in integrated ring resonators that have a higher order mode crossing. Further, they do not need the dynamic and very complicated pumping schemes or elaborate stabilization that self-localised DKS states need [104]. The basis of their stable behaviour originates from the fact that their intra-cavity power is dramatically higher than DKS states. In fact, it is very similar to the power levels of the chaotic temporal states [47, 105]. As a result, there is a very small difference in power levels in the cavity when the soliton crystal states are created out of chaos, and so there is no change in the resonant frequency. It is this self-induced frequency detuning arising from thermal instability due to the soliton step that renders pumping of DKS states, for example, challenging [106]. The combined effect of natural stability and robust and simple manual generation and the overall efficiency of soliton crystals that makes them extremely attractive for very high bandwidth data transmission.

a) CMOS compatible micro-comb source

The micro-ring resonator (MRR) for the soliton crystal comb generation was the same as described above for the optical neural network experiments. The soliton crystal device and the soliton crystal comb spectra are shown in Figure 4. As before, the microcomb had a 48.9 GHz FSR, producing a soliton crystal output with a spectrum spanning across > 80 nm while pumping at 1.8 watts of CW power at a wavelength of 1550nm. The soliton crystal micro-comb was preceded first by the primary comb and displayed a variation in comb line powers at ± 0.9 dB, for ten different incidents of initiation, and was achieved by sweeping the wavelength manually from 1550.300 - 1550.527 nm. This clearly proves the micro-comb turn-key generation repeatability for our devices. Out of the total number of generated comb lines, eighty were chosen from the 3.95 THz, 32 nm wide C-band window at 1536 – 1567 nm.

The MRR chip was mounted on a Peltier cooler, monitored by a standard NTC temperature sensor. The temperature was maintained with a thermo-electric cooler (TCM-M207) at 25°C, within 0.1°C of accuracy. The laser was set to standard running mode, with no extra steps made to stabilise the output frequency. Soliton crystal generation was

Figure 14. Alternate comb generation states and comb stability measurements. a) Measured spectrum of an alternate soliton crystal generation state achieved through manual wavelength tuning of the pump laser. b) Captured spectrum of a modulation instability ('chaotic') state, also achieved via manual wavelength tuning. c) Final complete soliton crystal spectrum used in the transmission experiments.

achieved by automated wavelength tuning, in turn reducing the system complexity compared to other micro-comb generation schemes [33]. We measured the internal conversion efficiency of our soliton crystals to be 42% for the whole spectrum, and 38% when selecting the 80 lines over the C-band, highlighting that over 90% of our available comb power is compatible with standard C-band equipment. The high performance of SCs including both robustness and efficiency [47, 105] stems from the fact that in SCs the resonator is virtually completely filled with solitons making the intracavity

energy very close to the chaotic state, whereas the energy of DKS states largely resides in the CW background rather than the single soliton pulse. Because soliton crystals are tightly packed systems of self-localized pulses, they have more than ten times higher intra-cavity power than the DKS regime — in fact, close to the power of the spatiotemporal chaotic states [47, 34]. Therefore, when sweeping the pump wavelength, switching between chaotic and soliton crystal states does not introduce significant changes in intra-cavity power. Hence, soliton crystals do not suffer from thermal detuning due to the characteristic ‘soliton step’ observed for single soliton states. This has the important consequence that they can be initiated through adiabatic pump wavelength sweeping. Hence, soliton crystals are highly robust and can provide stable micro-combs without the need for complex feedback systems [34, 105]. This intrinsic robustness is central for enabling the use of micro-combs outside of the laboratory. While they have been successfully exploited for microwave photonics [68], their potential for coherent optical communications has only been reported recently [34].

The evolution of the soliton crystal micro-comb while tuning the pump laser manually is shown in Figure 4. The open loop generation of the soliton crystal comb along with the low (~ 50 GHz) line spacing and wide bandwidth of the high conversion efficiency first ‘lobe’ (4 THz) makes these combs attractive for application to compact, high rate transceivers for optical communications. For comparison, the actual soliton spectrum used in the experiments is shown again in (Figure 14c). The automatic wavelength sweeping over a pre-determined wavelength range (1550.300 – 1550.527 nm) to generate the wanted soliton crystal state was achieved by wavelength tuning the laser in 1pm steps, with half a second between each step. This tuning was achieved by simply setting the wavelength to different values on a Yenista Tunics T100HP, remotely controlling the unit via GPIB through Python. The tuning rate was found through trial

Figure 15. Soliton crystal micro-comb communications experiment. A CW laser, amplified to 1.8W, pumped a 48.9 GHz FSR micro-ring resonator, producing a micro-comb from a soliton crystal oscillation state. The comb was flattened and optically demultiplexed to allow for modulation, with the data multiplexed before the transmission through fibres with EDFA amplification. At the receiver, each channel was demultiplexed before reception. ECL, edge-coupled laser, WSS wavelength-selective switch.

and error to allow for reproducible generation of the desired soliton crystal state. While the generation of this state seems robust and repeatable, we make no firm claims to deterministic generation of soliton crystal states, as we have not modelled the generation of the desired state with thermal terms introduced into the LLE, which would allow for the simulation of the effects of thermo-optic chaos [28]. We note that there have been successful demonstrations of deterministic micro-comb generation, although these involve auxiliary systems which greatly increase overall system complexity [108, 109]. In order to estimate the internal conversion efficiency from the pump to comb line within the micro-ring resonator, we analyse the relative powers of the pump line and comb lines emitted from the drop port of the device. Measurement of light from the drop port of our 4-port, dual-bus device accurately reflects the light within the resonator. We define internal conversion efficiency as the power ratio between the pump line at 1551.05 nm and the comb lines, consistent with [28, 103, 110, 111]. Analysing the spectrum gathered on a standard OSA running at 0.06 nm resolution, we measure an internal conversion efficiency of 42.1%. This compares favourably with dark [112] (20%) and bright [103, 111] solitons ($< 0.6\%$), as expected. Clearly, when taking other factors into account, such as coupling loss, proportion of comb lines used, etc., this will naturally drop, as it will for all devices. We do note that our chips are fibre pigtailed with on-chip mode converters, resulting in extremely low fibre-chip coupling loss < 0.5 dB. Further, after selecting the 80 comb lines over the C-band our internal efficiency was 38.1%. This highlights that our soliton crystal state provides most of its power ($> 90\%$) in a useful bandwidth for standard C-band optical communications equipment. Once established, the robustness of soliton crystal states to relatively large laser frequency drifts was verified by observing that the same stable oscillation state was maintained, and with a maximum power variation of only +2.4 dB to -1.7 dB for the 80 lines used for the transmission experiments even while tuning the pump laser by more than 12 pm (1.5 GHz) — much larger than the actual variation of the pump laser wavelength in the transmission experiments. We

measured the power stability of the comb lines (Figure 4) over 66 hours, with spectra captured every 15 minutes with only open loop control (standard thermo-electric controllers) and without manual intervention to stabilise the micro-ring resonator chip. Figure 4a shows the measured spectrum for the 80 C-band comb lines along with the standard deviation (SD in dBm - given by the error bars), showing a relative SD (in dB, Figure 4b) of about -14dB over the 66 hour period. Conceptually, these measurements demonstrate that the fiber-coupled MRR device can operate as a separate, independent, plug-in element that multiplies the number of coherent carriers produced by an independent laser by 80 times. We emphasize, however, that in fact it is the net superchannel transmission that is the ultimate test of the comb's performance – not the power stability measurements – and so OSNR is far more important. As long as the required OSNR is maintained, high fidelity data transmission will be supported, which was the case in our experiments. Further, due to the self-stabilizing nature of soliton states in micro-resonators [28, 103], there are fundamental reasons to expect even greater open-loop long-term stability than we have shown in our current measurements. Finally, any improvements to the system such as the addition of feedback control loops, the use of a pump laser with greater frequency stability, better design of the micro-ring resonator thermal stabilization, or the avoidance of comb flattening, would enhance our superchannel transmission performance even further.

In our experiments we used an external laser source and amplifier to pump the micro-ring used to generate the optical frequency comb. We note that recent demonstrations of hybrid integration of pump lasers to generate single DKS microcombs [44, 68, 113-115], as well as advanced techniques such as injection locking [44], can equally be applied to soliton crystals – in fact probably even more easily given their much simpler generation process. Hybrid integrated pump sources also yield much greater energy efficiency - the states in [113] were produced with 4 dBm (2.5 mW) pump power. Moreover, the external cavity structure shown in [113] is intrinsically compatible with the much less complex and slower tuning required by soliton crystals.

Integrated laser arrays can now produce high quality optical carriers. Although a full analysis comparing the different approaches is beyond the scope of this work, we make a number of comments. To support tightly spaced superchannel multiplexing with 80 separately integrated lasers with the same performance as micro-combs, it would require precise wavelength locking of all 80 lasers to compensate for variations due to fabrication error. Moreover, the footprint of 80 discrete high-quality lasers may be prohibitively large. Micro-combs, on the other hand, have been realized with hybrid integrated pump sources and have achieved ultralow thresholds <10 mW [44]. Microcombs are intrinsically locked to the micro-resonator free-spectral range, without any feedback requirements. They also have a small footprint, being based on a single high quality resonator and single pump laser source.

a) Systems experiments

We performed 2 experiments: the first across 75 km of single mode optical fiber in the lab and the second in a field trial using a metropolitan network in the greater Melbourne area, also based on standard SMF, which linked Monash University's Clayton campus to the RMIT campus in the Melbourne CBD. A map of the metropolitan network used for the system field trial is given in Figure 13, while the soliton crystal device and the comb spectra are shown in Figure 4. The experimental setup for the demonstration of high capacity optical data transmission is shown in Figure 15 (simplified overview) and in Figure 16 in more detail.

The transmission link was comprised of two fibre cables connecting labs at RMIT University (Swanston St., Melbourne CBD) and Monash University (Wellington Rd, Clayton). These cables were routed from the labs access panels, to an interconnection point with the AARNet's fibre network. The fibre links were a mix of OS1 and OS2 standard cables and include both subterranean and aerial paths. There was no active equipment on these lines, providing a direct dark fibre connection between the two labs. The total loss for these cables was 13.5 dB for the RMIT-Monash link and 14.8 dB for the Monash-RMIT paths. The cable lengths as measured by OTDR were both 38.3 km (totalling 76.6 km in loop-back configuration). At Monash, an EDFA was remotely monitored and controlled using a 1310 nm fibre-ethernet connection running alongside the C-band test channels. The comb was amplified to 19 dBm before launch, at Monash, and on return to RMIT. The installed network fiber for the field trial presented a different testing platform to the spooled fibres used in lab. Splices and connections along the link between the two labs provide a source of uncontrolled back-reflections and limit the amount of power that can safely be sent over the network given the risk of connector burns and even fibre fuses from reflective interfaces. Coupled with the higher losses of installed (legacy) fiber links, this provides a challenging platform for high spectral efficiency optical communications where maximising signal to noise ratio is key to enabling high capacities. Moreover, the operation over legacy fibre links covering typical suburban distances demonstrates that it is possible to leverage installed fiber infrastructure for next generation metropolitan/regional area systems, which have been experiencing a higher growth in required capacity than long-haul networks, actually surpassing them in 2017 [34, 116]. This is particularly important due to the cost of laying new fiber in

installed ducting being on the order of \$30k / mile [34, 117]. It also demonstrates the feasibility of system upgrades using micro-comb-based transceivers to extend the useful lifetime of installed fiber systems.

The microcomb featured a 48.9 GHz FSR, producing a soliton crystal output with a spectrum spanning across > 80 nm while pumping at 1.8 watts of CW power at a wavelength of 1550nm. The soliton crystal micro-comb was preceded first by the primary comb and displayed a variation in comb line powers of $< \pm 0.9$ dB, for ten different incidents of initiation, and was achieved by sweeping the wavelength manually from 1550.300 - 1550.527 nm. This demonstrates the micro-comb turn-key generation repeatability. Since the nonuniform spectrum of soliton crystal combs has been viewed as a weakness, we chose to flatten the optical frequency comb first so that all lines were of equal power, even though this is not necessary and actually introduces unnecessary impairments – both in our experiments and other micro-comb demonstrations (e.g. [33, 34]). All comb lines were wavelength demultiplexed into separate waveguides and sent to separate modulators. It is then straightforward to adjust the comb line power by variable attenuators, amplifiers or even by varying the RF drive amplitude to the modulators. We implemented comb flattening for several reasons: i) to prove our system operation under the most demanding conditions, ii) to pre-empt the criticism that the nonuniform spectrum of SCs is a limitation, and iii) to facilitate easy comparisons with prior art. Since avoiding flattening would reduce impairments and improve our performance (by increasing the OSNR of the higher power comb lines and their ability to carry higher spectral efficiency modulation formats, and by eliminating the loss of the extra WaveShaper), it does not represent a limitation to SC based transmission.

The soliton crystal micro-comb was flattened in two stages by two independent programmable optical filters (Finisar WaveShaper 4000S). The WaveShapers had an insertion loss of 5dB each, in addition to any variable attenuation. The first had a static filter shape set to equalize each comb line to within about 1 dB of each other, to coarsely match the generic shape of the soliton crystal state we chose to use. The second programmable filter was set each time that a new soliton crystal state was initiated, to equalize the comb line powers to within < 1 dB of each other, although we note that it was often unnecessary to change the filter profile when generating a new soliton crystal. Spectral shaping in a WDM transceiver using a comb source involved minimal extra complexity since only the addition of attenuators after the WDM demultiplexer were required to route each comb line to a separate modulator. The comb was then amplified by a further polarization maintaining EDFA (Pritel PMFA-20-IO), before being divided for modulation. Prior to modulation, the optical signal-to-noise ratio (OSNR) of the individual comb lines was > 28 dB. Comb flattening is not necessary either in our experiments or other micro-comb demonstrations [33, 34], since all comb lines are typically wavelength demultiplexed into separate waveguides and sent to separate modulators. It is then straightforward to adjust the comb line power by variable attenuators, amplifiers or even by varying the RF drive amplitude to the modulators. In fact, we expect better performance without comb flattening since the higher power comb lines would need less attenuation and/or amplification before modulation, resulting in a higher OSNR, while the lower power comb lines would have essentially the same performance as reported here. Furthermore, using the raw spectrum would avoid the loss of the extra WaveShaper. Therefore, avoiding flattening (working with the raw spectrum) would in fact yield even higher system performance.

Out of the total number of generated comb lines, eighty were chosen from the 3.95 THz, 32 nm wide C-band window at 1536 – 1567 nm. The spectrum was then flattened using a WaveShaper. Following this the number of wavelengths was doubled to 160, corresponding to a 24.5 GHz spacing, to increase the spectral efficiency. This was accomplished with a single sideband modulation technique that generated both even and odd channels that were not correlated. We then grouped six wavelengths, with the rest of the bands supporting data loaded channels based on the same even-odd structure. We were able to use a record high order 64 QAM coherent modulation format that modulated the whole comb at a baud rate of 23 Giga-baud, that achieved 94% utilization of the available spectrum. After the spectrum was flattened with the WaveShaper, the number of wavelengths doubled to 160, corresponding to a 24.5 GHz spacing, to increase the spectral efficiency. This was accomplished with a single sideband modulation technique that generated both even and odd channels that were not correlated. We then grouped 6 wavelengths, with the rest of the bands supporting data loaded channels based on the same even-odd structure. We used 64 QAM coherent modulation at 23 Giga-baud, with 94% utilization of the available spectrum.

Figure 16. Experimental set-up. ECL: External cavity laser, PM-EDFA: Polarization maintaining EDFA, Pol. Cont.: Polarization controller, uRR: micro-ring resonator, TEC: Thermo-electric controller, WSS: Wavelength selective switch, PM 3dB: Polarization maintaining 50%/50% optical power coupler/splitter, PBS: Polarization beam splitter, PMF: Polarization maintaining fibre, CMZM: Complex Mach-Zehnder modulator (dual nested type), AWG: Arbitrary waveform generator, RF Amps.: microwave amplifiers, WDM: Wavelength division multiplexer (1310/1550 nm), fibre/ether: Gigabit Ethernet to fibre channel media converter, Coh. Rx: integrated optical hybrid and balanced photodiode coherent receiver, DSO: Digital storage oscilloscope (real-time). For the EDFAs, LW2020 denotes Lightwaves2020 MOABF17 optical amplifiers, PT is a Photonic Technologies 'Powerflat' amplifier. f) depicts the concept of a receiver for every sub-band.

6. RESULTS AND DISCUSSION

Figures 17 and 18 show the experimental results. We conducted two transmission experiments, sending data over 75 km of single mode fibre in the laboratory as well as a field trial over an installed metropolitan-area fibre network connecting the Melbourne City campus of RMIT and the Clayton campus of Monash University, spanning greater metropolitan Melbourne. From the micro-comb, 80 lines were selected over the C-band (32 nm wide, 3.95 THz window from 1536 – 1567 nm) which were then flattened with a WaveShaper and the number of wavelengths doubled to 160 (equivalent to a spacing of 24.5 GHz) to optimize the spectral efficiency with a single-sideband modulation scheme to generate odd/even decorrelated test channels. We then combined a test band of six channels, with the remaining bands providing loading channels having the same odd-and-even channel structure. We used a record high order format of 64 QAM to modulate the entire comb at 23 Giga-baud, resulting in the utilization of 94% of the available spectrum.

Spectra of the comb at key points are given in Figure 17 a-c. Figure 17d shows constellation diagrams for the signal at 194.34 THz. In back-to-back configuration (i.e. with transmitter directly connected to receiver) we measured signal quality (Q2, from error vector magnitude) at almost 18.5 dB, dropping to near 17.5 dB when transmitting the fully modulated comb through the test links. Figure 18a shows the transmission performance using the bit error ratio (BER) for each channel as a metric, showing the 20% threshold for soft-decision forward error correction (SD-FEC), a common benchmark for performance, using a proven code, at a BER of 4×10^{-2} [118]. Three scenarios were investigated: i) a direct connection between the transmitter stage to the receiver (back-to-back, B2B), after transmission through ii) in-lab fibre and iii) the field trial. Transmission globally degraded the performance of all channels, as expected.

All results were below the given FEC limit, but since using SD-FEC thresholds based on BER is less accurate for higher order modulation formats and for high BERs [119], we additionally used generalized mutual information (GMI) to calculate the system performance. Figure 18b plots the GMI for each channel and its associated SE, with lines given to indicate projected overheads. We achieved a raw bitrate (line-rate) of 44.2 Tb/s, which translates to an achievable coded rate of 40.1 Tb/s (in B2B), dropping to 39.2 Tb/s and 39.0 Tb/s for the lab and field trial transmission experiments, respectively. These yielded spectral efficiencies of 10.4, 10.2 and 10.1 b/s/Hz.

This data rate represents an increase of nearly 50% (see below) over the highest reported result from a single integrated device [33]. Even more importantly the SE is enhanced even more, being a factor of 3.7. This is notable considering that we performed our experiments under the most demanding conditions, including not using any closed-loop feedback, stabilization or elaborate initiation schemes, as well as with the use of full comb flattening (equalization). Even though it is not necessary, we implemented flattening mainly to address possible concerns that the nonuniform spectra of soliton crystals might pose a limitation. Given that we achieved our results with flattening, and avoiding it would eliminate the impairments it introduced, thus improving our performance, it therefore does not represent a limitation. The same argument holds true for closed-loop feedback control of the micro-comb.

The record high spectral efficiency and absolute bandwidth that we achieve were enabled by the very high conversion efficiency we achieved between the pump and the soliton crystal comb lines [47, 105]. Again, as mentioned this results from the very small power step in the cavity that occurs when the soliton crystals are generated from the chaos states.

We only used the telecom C-band, and yet the bandwidth of the microcomb was larger than 80 nm. Therefore, wavelengths in both the L (1565-1605 nm) and even S (1500-1535 nm) bands could easily be used. In fact, even broader bandwidths can be achieved by increasing the power, by varying the pump wavelength, by engineering the dispersion or further methods. This would yield a 3x increase in bandwidth resulting in >120 Tb/s with only a single source.

Achieving even lower spacings, or FSRs, with micro-combs would yield yet higher SEs since the quality of the signal increases for smaller baud rates. This may result in a smaller overall comb bandwidth however. For our experiments, the use of single sideband modulation allowed multiplexing two channels using one single wavelength, which cut the comb spacing by a factor of 2 while enhancing the back-to-back performance that was limited by transceiver noise. This was made possible by the stability of the soliton crystals. Conversely, electro-optic modulation has also been used to subdivide the micro-comb repetition rate, and this would also create broader comb bandwidths. This, however, would require locking the comb FSR spacing to an external RF source, although this is feasible since sub megahertz stabilization of microcombs has been achieved [120-122]. Furthermore, increasing the comb conversion efficiency by using recently reported laser cavity-soliton micro-combs [45] will offer a powerful way to increase the system capacity as well as the quality of the signal. For recently installed networks, our approach can easily be complemented by spatial division multiplexing based on multiple core fibre [123], yielding bandwidths of more than a petabit/s with a single source. Our results build on our work on soliton crystal combs for RF signal processing [67-87]. This paper represents the most advanced demonstration of micro-combs for ease of generation, coherence, stability, noise, efficiency, and others, and is a direct result of the superior soliton crystal microcomb qualities.

IV. PERFORMANCE COMPARISON

Table II summarizes key results from the literature comparing the system performance metrics for both integrated sources and rack-mounted equipment in both standard fibre or calculated on a per-mode basis for multicore fibre. See footnote [124] for details on how we calculate our performance metrics. Prior to this work, the best result (per core) was from [33], where, a single microcomb supported 30.1 Tb/s over the C and L bands, when using a standard tuneable laser coherent receiver. This is not only the best published result using a single micro-comb, but closely resembles our experiment (single micro-comb at transmitter, single tuneable laser at the receiver as a local oscillator). Note that our system uses only the C-band [34], while we improve the data rate because of the higher spectral efficiency. Reference [33] provides three different system set-ups. The first is closest to our own demonstration, and we use this in the main text for a direct comparison. In this case a single micro-comb is used at the transmitter, and a single laser employed as a local oscillator at the receiver. In the second demonstration in [33], two frequency interleaved combs are used in the transmitter, which, although resulting in a higher overall data rate, resulted in a lower per-comb spectral efficiency, and consequently a lower overall per-comb data rate than the single-comb demonstration. In the third demonstration in [33], a microcomb source is used as a local oscillator in the receiver, with a single comb employed at the transmitter side. In this case, overall system performance was improved to achieve a higher spectral efficiency and hence a higher overall rate, although lower than what we have demonstrated here. High modulation formats have also been achieved with dark solitons [101], yet at a lower overall data rate, primarily due to the large comb line spacing that limits the spectral efficiency. In ref. [101], the same modulation cardinality (i.e. 64 QAM) was used; however, the high comb line spacing and relatively low symbol rate translated to a low spectral occupancy, resulting in a low spectral efficiency and aggregate data rate. The dark soliton state used in this demonstration also seems to require feedback to stabilize the state, increasing system complexity.

The work of [98] used a comb generator based on a benchtop pulsed seed laser source combined with waveguide spectral broadening. To provide a fully integrated system, this source would need to be on-chip. The focus in that experiment was using novel, proprietary multi-core fibre to achieve a 30-fold increase in bandwidth over the fibre in this spatially multiplexed system, to reach 0.66 Petabits/s. On a per-mode basis, Ref. [98] yields 25.6 Tb/s/mode, a lower per-mode capacity than this work and [33]. We note that both our approach and that of [33] are able to take advantage of SDM techniques to scale the overall bandwidth by using multi-core fibre. Although spectral utilization in [98] was high, the per-mode spectral efficiency is lower than our demonstration, and used 16-QAM modulation.

In contrast, soliton-based microcomb generation from an integrated hybrid chip has been demonstrated [44, 114]. For completeness, we also compare against record per-mode results for non-chip based sources - both benchtop-scale and rack mounted comb sources, and with traditional multiple discrete laser based WDM systems. The highest achieved per-mode data rate from a single comb used a non-integrated rack-mounted, or benchtop fibre-based comb [123] – i.e., orders of magnitude larger (and more expensive) than chip-based microcombs, and which has not been realized in integrated form. There, the per-mode capacity across the C & L bands was 97.75 Tb/s, by using high cardinality modulation (64-QAM), high spectral utilization (24.5 Gbd on a 25 GHz grid) and efficient forward error correcting codes. Our work shows that the expansion to the L-band with comb spacings similar to our own can enable even higher aggregate rates.

The highest aggregate rate achieved over single mode fibre [125] used multiple discrete light sources. Again, that demonstration leveraged both C & L bands, although the achieved spectral efficiency is not significantly greater than in our demonstration, even when using 256-QAM relatively high code rate overheads were required (40-50%). This indicates that a single micro-comb source with quality similar to our own can provide performance comparable to multiple discrete laser sources. We note that the results presented in [125] were achieved using highly specialized fibres and with specialized amplification, neither of which are available in typical networks. The high aggregate rate they present was in part due to the full use of C- and L-bands, compared with the single C-band system we demonstrate here.

Our high spectral efficiency is partly enabled by the baud rate of the channels we modulate. The optimum baud rate in systems where OSNR is not the dominant performance limiter is the subject of on-going research. In systems with high OSNR, noise added by the transmitter and receiver limit performance, either from thermal noise from electrical components, quantization noise from the A/D and D/A converters, as well as noise and distortion from modulation and photodetection. This has proved important for superchannel reception [126, 127]. Also, we note that ultra-high spectral efficiency (>15 b/s/Hz) modulation has been based on low baud rates (between 6-10 Gbd) [128, 129]. It is on this basis that we suggest that combs with a lower FSR may enable a higher spectral efficiency, which should also improve the single-device data rate. The trade-off is that a reduction of line spacing results in a lower power per-line, impacting the individual comb line OSNR (for the same bandwidth), which may ultimately pose a limit to the spectral efficiency [120]. Finally, soliton crystals should be achievable in mid-IR frequencies for a very wide range of applications [131-136].

Figure 17. Spectra of soliton crystal frequency comb after flattening (a), modulation and transmission through 75 km spooled in-lab fibre (b) field-trial link (c). Spectrum (a) measured with 12.5 GHz resolution to resolve individual comb lines, (b) and (c) are plotted at 50 GHz resolution to illustrate average channel powers. Flattening equalised the comb line power to within 1 dB. After modulation and amplification, channels were shaped by the EDFA gain spectrum. The inset in (b) depicts the test channel spectra captured with a 150 MHz resolution optical spectrum analyser, showing the odd and even sub-bands modulated onto each comb line in the test band. d) Constellation diagrams for a comb line at 193.4 THz (1550.1 nm) for X- and Y-polarization channels. ‘Back to back’ is the transmitter directly connected to the receiver, ‘75 km in-lab fibre’ indicates transmission over 75 km of fibre in the lab, ‘76.6 km field fibre’ is after transmission in field-trial link. BER and Q^2 related to the constellations are noted.

Figure 18. a) BER for each comb line. Blue circles points indicate performance of channels in B2B configuration, red squares dots are for performance after transmission through 75 km of in-lab spooled fibre, while green triangles are after transmission through the 76.6 km installed metropolitan-area fibre link. An indicative FEC threshold is given at 4×10^{-2} , corresponding to a pre-FEC error rate for a 20% soft-decision FEC based on spatially-coupled LDPC codes [130] (dashed line). After transmission, all channels were considered to be error-free, b) GMI and spectral efficiency measured for each comb line. GMI was calculated after normalization to scale measured constellations in order to account for received signal-to-noise ratio (SNR). Lines are for 20% and 10% overheads. Spectral efficiency was derived from GMI, and the ratio of symbol rate to comb spacing. GMI indicates a higher overall capacity than BER with the indicated SD-FEC threshold, as GMI assumes the adoption of an ideal code for the system. For B2B, GMI (SE) varied between 11.3 b/symb. (10.6 b/s/Hz) and 10.9 b/symb. (10.3 b/s/Hz). After in-lab fibre transmission, the achievable per-channel GMI (SE) varied between 11.0 b/symb. (10.4 b/s/Hz) and 10.7 b/symb. (10.1 b/s/Hz), with the same range observed for the installed field-trial fibres. We estimate the capacity from the sum of the GMIs, multiplied by the symbol rate.

Table II

Line Rate	Net Rate	SE	Transmission	Reference
30.1 Tb/s	28.0 Tb/s	2.8 b/s/Hz	75 km SMF in-lab	[33], 1 st
27.5 Tb/s	25.1 Tb/s	2.6 b/s/Hz	75 km SMF in-lab	[33], 2 nd
37.2 Tb/s	34.6 Tb/s	3.5 b/s/Hz*	75 km SMF in-lab	[33], 3 rd
44.2 Tb/s	40.1 Tb/s	10.4 b/s/Hz	B2B (0 km)	[34]
44.2 Tb/s	39.2 Tb/s	10.2 b/s/Hz	75 km SMF in-lab	[34]
44.2 Tb/s	39.0 Tb/s	10.1 b/s/Hz	76.6 km SMF installed	[34]
4.8 Tb/s*	4.4 Tb/s	0.98 b/s/Hz*	80 SMF km in-lab	[101]
25.6 Tb/s*	22.0 Tb/s*	3.2 b/s/Hz*	9.6 km, 30-core	[98]
117.3 Tb/s*	97.75 Tb/s	9.8 b/s/Hz*	31 km, 22-core	[123]**
174 Tb/s*	120 Tb/s	10.8 b/s/Hz*	70 km 'Z'-type SMF	[125]***

Note – results highlighted in yellow are not from chip-based sources but from rack mounted or bench top based sources, included for completeness.

Table II. Key systems performance metrics, per comb source used in the transmitter, and on a per mode basis. '*' indicates that this figure was not directly provided in the reference, and so is inferred from data provided. '**' indicates a demonstration using a commercial benchtop comb source, '***' indicates a traditional WDM result using multiple laser sources.

7. CONCLUSION

We demonstrate a universal optical vector convolutional accelerator operating at 11 TOPS on images of 250,000 pixels with 8-bit resolution for 10 kernels simultaneously — enough for facial image recognition. We use the same hardware to form a deep optical CNN with 10 output neurons, achieving recognition of full 10 digits with 900 pixel handwritten images at 88% accuracy. Our approach is scalable and trainable to more complex networks for demanding applications to unmanned vehicles and real-time video recognition. We also demonstrate a world record ultra-high bandwidth optical transmission with a single source over standard optical fiber. We use soliton crystal micro-combs with a low FSR of 48.9GHz. Our results arise from this low comb spacing and the efficient, broad bandwidth, and stable nature of soliton crystals. They are low noise, coherent and can be initialised and operated with simple open-loop control.

References

1. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
2. Schalkoff, R. J. Pattern recognition. Wiley Encyclopedia of Computer Science and Engineering (2007).
3. Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
4. Silver, D. et al. Mastering the game of Go without human knowledge. *Nature* **550**, 354 (2017).
5. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet Classification with Deep Convolutional Neural Networks. *Commun. ACM* **60**, 84–90 (2017).
6. Yao, P. et al. Fully hardware-implemented memristor convolutional neural network. *Nature* **577**, 641–646 (2020).
7. Lawrence, S. et al., "Face recognition: A convolutional neural-network approach", *IEEE transactions on neural networks* **8**, 98–113 (1997).
8. Shen, Y. et al. Deep learning with coherent nanophotonic circuits. *Nature Photonics* **11**, 441 (2018).
9. Larger, L. et al. High-speed photonic reservoir computing using a time-delay-based architecture: Million words per second classification. *Phys. Rev. X* **7**, 011015 (2017).
10. Peng, H., Nahmias, M. A., Lima, T. F. d., Tait, A. N. & Shastri, B. J. Neuromorphic Photonic Integrated Circuits. *IEEE Journal of Selected Topics in Quantum Electronics* **24**, 1–15, doi:10.1109/JSTQE.2018.2840448 (2018).
11. Lin, X. et al. All-optical machine learning using diffractive deep neural networks. *Science* **361**, 1004–1008 (2018).
12. Feldmann, J. et al., "All-optical spiking neurosynaptic networks with self-learning capabilities", *Nature* **569**, 208–214 (2019).
13. X. Xu et al., "Photonic perceptron based on a Kerr microcomb for scalable high speed optical neural networks", *Laser and Photonics Reviews*, vol. 14, no. 8, 2000070 (2020). DOI: 10.1002/lpor.202000070.
14. X. Xu, et al., "11 TOPs photonic convolutional accelerator for optical neural networks", *Nature* **589**, 44–51 (2021).
15. Feldmann, J. et al., "Parallel convolutional processing using an integrated photonic tensor core", *Nature* **589**, 52–58 (2021).
16. B. J. Shastri et al., "Photonics for artificial intelligence and neuromorphic computing", *Nature Photonics* **15**, (2) 102–114 (2021).
17. H. Wu, and Q. Dai, "Artificial intelligence accelerated by light", *Nature* **589**, 25–26 (2021).
18. G. Wetzstein et al., "Inference in artificial intelligence with deep optics and photonics", *Nature* **588** (7836), 39–47 (2020).
19. Ambrogio, S. et al., "Equivalent-accuracy accelerated neural-network training using analogue memory", *Nature* **558**, 60 (2018).
20. Esser, S. K. et al., "Convolutional networks for fast, energy-efficient neuromorphic computing", *Proc. of the National Academy of Sciences* **113**, 11441 (2016).
21. Graves, A. et al. Hybrid computing using a neural network with dynamic external memory. *Nature* **538**, 471–476 (2016).
22. Miller, D. A. B. Attojoule Optoelectronics for Low-Energy Information Processing and Communications. *Journal of Lightwave Technology* **35**, 346–396 (2017).
23. Appeltant, L. et al. Information processing using a single dynamical node as complex system. *Nature Communications* **2**, 468 (2011).

24. Chang, J., Sitzmann, V., Dun, X., Heidrich, W. & Wetzstein, G. Hybrid optical-electronic convolutional neural networks with optimized diffractive optics for image classification. *Scientific Reports* **8** (2018).
25. Vandormael, K. et al., "Experimental demonstration of reservoir computing on a silicon photonics chip", *Nature Communications* **5**, 3541 (2014).
26. Brunner, D. et al., "Parallel photonic information processing at gigabyte/s data rates using transient states", *Nature Comm.* **4**, 1364 (2013).
27. Tait, A. N. et al., "Demonstration of WDM weighted addition for principal component analysis", *Optics Express* **23**, 12758-12765 (2015).
28. Pasquazi, A. et al. Micro-combs: a novel generation of optical sources. *Physics Reports* **729**, 1-81 (2018).
29. Moss, D. J. et al., "New CMOS-compatible platforms based on silicon nitride and Hydex for nonlinear optics", *Nature photonics* **7**, 597 (2013).
30. Kippenberg, T. J., Gaeta, A. L., Lipson, M. & Gorodetsky, M. L. Dissipative Kerr solitons in optical microresonators. *Science* **361**, 567 (2018).
31. Savchenkov, A. et al. Tunable optical frequency comb with a crystalline whispering gallery mode resonator. *Phys. Rev. Lett.* **101**, 093902 (2008).
32. Spencer, D. T. et al. An optical-frequency synthesizer using integrated photonics. *Nature* **557**, 81-85 (2018).
33. Marin-Palomo, P. et al. Microresonator-based solitons for massively parallel coherent optical communications. *Nature* **546**, 274 (2017).
34. B. Corcoran, et al., "Ultra-dense optical data transmission over standard fiber with a single chip source", *Nature Communications*, vol. 11, Article:2568 (2020).
35. Kues, M. et al. Quantum optical microcombs. *Nature Photonics* **13**, (3) 170-179 (2019). doi:10.1038/s41566-019-0363-0
36. P. Roztocki et al., "Complex quantum state generation and coherent control based on integrated frequency combs", *Journal of Lightwave Technology* **37** (2) 338-347 (2019).
37. S. Sciara et al., "Generation and Processing of Complex Photon States with Quantum Frequency Combs", *IEEE Photonics Technology Letters* **31** (23) 1862-1865 (2019). DOI: 10.1109/LPT.2019.2944564.
38. L. Caspani, C. Reimer, M. Kues, et al., "Multifrequency sources of quantum correlated photon pairs on-chip: a path toward integrated Quantum Frequency Combs," *Nanophotonics*, vol. 5, no. 2, pp. 351-362, 2016.
39. C. Reimer et al., "Generation of multiphoton entangled quantum states by means of integrated frequency combs," *Science*, vol. 351, no. 6278, pp. 1176-1180, 2016.
40. M. Kues, et al., "On-chip generation of high-dimensional entangled quantum states and their coherent control", *Nature*, vol. 546, no. 7660, pp. 622-626, 2017.
41. P. Roztocki et al., "Practical system for the generation of pulsed quantum frequency combs," *Optics Express*, vol. 25, no.16, 18940-18949, 2017.
42. Y. Zhang, et al., "Induced photon correlations through superposition of two four-wave mixing processes in integrated cavities", *Laser and Photonics Reviews*, vol. 14, no. 7, pp. 2000128, 2020. DOI: 10.1002/lpor.202000128
43. C. Reimer, et al., "High-dimensional one-way quantum processing implemented on d-level cluster states", *Nature Physics*, vol. 15 (2) 148 (2019).
44. Stern, B., Ji, X., Okawachi, Y., Gaeta, A. L. & Lipson, M. Battery-operated integrated frequency comb generator. *Nature* **562**, 401 (2018).
45. H. Bao, et al., Laser cavity-soliton microcombs, *Nature Photonics*, vol. 13, no. 6, pp. 384-389, Jun. 2019.
46. Lugiato, L. A., Prati, F. & Brambilla, M. *Nonlinear Optical Systems*, (Cambridge University Press, 2015).
47. Cole, D. C., Lamb, E. S., Del'Haye, P., Diddams, S. A. & Papp, S. B. Soliton crystals in Kerr resonators. *Nature Photonics* **11**, 671 (2017).
48. Wang, W., et al., Robust soliton crystals in a thermally controlled microresonator, *Opt. Lett.*, **43**, 2002 (2018).
49. Bao, C., et al., Direct soliton generation in microresonators, *Opt. Lett.*, **42**, 2519 (2017).
50. M. Ferrera et al., "CMOS compatible integrated all-optical RF spectrum analyzer", *Optics Express*, vol. 22, (18) 21488 (2014).
51. A. Pasquazi, et al., "Sub-picosecond phase-sensitive optical pulse characterization on a chip", *Nature Photonics*, vol. 5, no. 10, pp. 618-623 (2011).
52. M. Kues, et al., "Passively modelocked laser with an ultra-narrow spectral width", *Nature Photonics*, vol. 11, no. 3, pp. 159, 2017.
53. L. Razzari, et al., "CMOS-compatible integrated optical hyper-parametric oscillator," *Nature Photonics*, vol. 4, no. 1, 41-45, 2010.
54. M. Ferrera, et al., "Low-power continuous-wave nonlinear optics in doped silica glass integrated waveguide structures," *Nature Photonics*, vol. 2, no. 12, pp. 737-740, 2008.
55. M. Ferrera et al., "On-Chip ultra-fast 1st and 2nd order CMOS compatible all-optical integration", *Opt. Express*, vol. 19, (23)pp. 23153-23161 (2011).
56. D. Duchesne, M. Peccianti, M. R. E. Lamont, et al., "Supercontinuum generation in a high index doped silica glass spiral waveguide," *Optics Express*, vol. 18, no. 2, pp. 923-930, 2010.
57. H Bao, L Olivieri, M Rowley, ST Chu, BE Little, R Morandotti, DJ Moss, ... "Turing patterns in a fiber laser with a nested microresonator: Robust and controllable microcomb generation", *Physical Review Research* **2** (2), 023395 (2020).
58. M. Ferrera, et al., "On-chip CMOS-compatible all-optical integrator", *Nature Communications*, vol. 1, Article 29, 2010.
59. A. Pasquazi, et al., "All-optical wavelength conversion in an integrated ring resonator," *Optics Express*, vol. 18 (4) 3858 (2010).
60. A. Pasquazi, Y. Park, J. Azana, et al., "Efficient wavelength conversion and net parametric gain via Four Wave Mixing in a high index doped silica waveguide," *Optics Express*, vol. 18, no. 8, pp. 7634-7641, 2010.
61. M. Peccianti, M. Ferrera, L. Razzari, et al., "Subpicosecond optical pulse compression via an integrated nonlinear chirper," *Optics Express*, vol. 18, no. 8, pp. 7625-7633, 2010.
62. Little, B. E. et al., "Very high-order microring resonator filters for WDM applications", *IEEE Phot. Technol. Lett.* **16**, 2263(2004).
63. M. Ferrera et al., "Low Power CW Parametric Mixing in a Low Dispersion High Index Doped Silica Glass Micro-Ring Resonator with Q-factor > 1 Million", *Optics Express*, vol.17, no. 16, pp. 14098-14103 (2009).
64. M. Peccianti, et al., "Demonstration of an ultrafast nonlinear microcavity modelocked laser", *Nature Comm.*, vol. 3, 765, 2012.
65. A. Pasquazi, et al., "Self-locked optical parametric oscillation in a CMOS compatible microring resonator: a route to robust optical frequency comb generation on a chip," *Optics Express*, vol. 21, no. 11, pp. 13333-13341, 2013.

66. A. Pasquazi, et al., "Stable, dual mode, high repetition rate mode-locked laser based on a microring resonator," *Optics Express*, vol. 20, no. 24, pp. 27355-27362, 2012.
67. Wu, J. et al., "RF Photonics: An Optical Microcombs' Perspective", *IEEE Journal of Selected Topics in Quantum Electronics* **24**, 1-20 (2018).
68. Xu, X., et al., Photonic microwave true time delays for phased array antennas using a 49 GHz FSR integrated micro-comb source, *Photonics Research*, **6**, B30-B36 (2018).
69. T. G. Nguyen et al., "Integrated frequency comb source-based Hilbert transformer for wideband microwave photonic phase analysis," *Opt. Express*, vol. 23, no. 17, pp. 22087-22097, Aug. 2015.
70. X. Xu, J. Wu, M. Shoeiby, T. G. Nguyen, S. T. Chu, B. E. Little, R. Morandotti, A. Mitchell, and D. J. Moss, "Reconfigurable broadband microwave photonic intensity differentiator based on an integrated optical frequency comb source," *APL Photonics*, vol. 2, no. 9, 096104, Sep. 2017.
71. X. Xu, M. Tan, J. Wu, R. Morandotti, A. Mitchell, and D. J. Moss, "Microcomb-based photonic RF signal processing", *IEEE Photonics Technology Letters*, vol. 31 no. 23 1854-1857, 2019.
72. X. Xu, et al., "Broadband RF channelizer based on an integrated optical frequency Kerr comb source," *Journal of Lightwave Technology*, vol. 36, no. 19, pp. 4519-4526, 2018.
73. X. Xu, et al., "Continuously tunable orthogonally polarized RF optical single sideband generator based on micro-ring resonators," *Journal of Optics*, vol. 20, no. 11, 115701. 2018.
74. X. Xu, et al., "Orthogonally polarized RF optical single sideband generation and dual-channel equalization based on an integrated microring resonator," *Journal of Lightwave Technology*, vol. 36, no. 20, pp. 4808-4818. 2018.
75. M. Tan et al, "Orthogonally polarized Photonic Radio Frequency single sideband generation with integrated micro-ring resonators", *IOP Journal of Semiconductors*, Vol. 42 (2021). DOI:10.13140/RG.2.2.17603.63523.
www.iopscience.iop.org/article/shaid/4c314db6c6dcfe19c275efde2b99a22329ba658eda52233b8ab4b8a08b55a7bfX
76. Xu, et al., "Advanced adaptive photonic RF filters with 80 taps based on an integrated optical micro-comb source," *Journal of Lightwave Technology*, vol. 37, no. 4, pp. 1288-1295 (2019).
77. X. Xu, et al., "Broadband microwave frequency conversion based on an integrated optical micro-comb source", *Journal of Lightwave Technology*, vol. 38 no. 2, pp. 332-338, 2020.
78. M. Tan, et al., "Photonic RF and microwave filters based on 49GHz and 200GHz Kerr microcombs", *Optics Comm.* vol. 465,125563, Feb. 22. 2020.
79. X. Xu, et al., "Broadband photonic RF channelizer with 90 channels based on a soliton crystal microcomb", *Journal of Lightwave Technology*, Vol. 38, no. 18, pp. 5116 - 5121, 2020. doi: 10.1109/JLT.2020.2997699.
80. X. Xu, et al., "Photonic RF and microwave integrator with soliton crystal microcombs", *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 67, no. 12, pp. 3582-3586, 2020. DOI:10.1109/TCSIL.2020.2995682.
81. X. Xu, et al., "Photonic RF phase-encoded signal generation with a microcomb source", *J. Lightwave Technology*, vol. 38, no. 7, 1722-1727, 2020.
82. X. Xu, et al., "High performance RF filters via bandwidth scaling with Kerr micro-combs," *APL Photonics*, vol. 4 (2) 026102. 2019.
83. M. Tan, et al., "Microwave and RF photonic fractional Hilbert transformer based on a 50 GHz Kerr micro-comb", *Journal of Lightwave Technology*, vol. 37, no. 24, pp. 6097 – 6104, 2019.
84. M. Tan, et al., "RF and microwave fractional differentiator based on photonics", *IEEE Transactions on Circuits and Systems: Express Briefs*, vol. 67, no.11, pp. 2767-2771, 2020. DOI:10.1109/TCSIL.2020.2965158.
85. M. Tan, et al., "Photonic RF arbitrary waveform generator based on a soliton crystal micro-comb source", *Journal of Lightwave Technology*, vol. 38, no. 22, pp. 6221-6226 (2020). DOI: 10.1109/JLT.2020.3009655.
86. M. Tan, X. Xu, J. Wu, R. Morandotti, A. Mitchell, and D. J. Moss, "RF and microwave high bandwidth signal processing based on Kerr Micro-combs", *Advances in Physics X*, VOL. 6, NO. 1, 1838946 (2021). DOI:10.1080/23746149.2020.1838946.
87. X. Xu, et al., "Advanced RF and microwave functions based on an integrated optical frequency comb source," *Opt. Express*, vol. 26 (3) 2569 (2018).
88. Krizhevsky A, Sutskever I, Hinton G E. Image net classification with deep convolutional neural networks. *Advances in neural information processing systems* **25**, 1097-1105 (2012).
89. LeCun, Y., et al., Gradient-based learning applied to document recognition. *Proc. of the IEEE* **86**, 2278-2324 (1998).
90. Bishop, C. M. *Neural networks for pattern recognition*. (Oxford university press, 1995).
91. Metcalf, A. J. et al. Integrated line-by-line optical pulse shaper for high-fidelity and rapidly reconfigurable RF-filtering. *Optics Express* **24**, 23925-23940 (2016).
92. Sahin, E et al., "Large, scalable dispersion engineering using cladding-modulated Bragg gratings on a Si chip", *App. Phys. Lett.* **110**, 161113 (2017).
93. D. J. Moss et al., "Tunable dispersion and dispersion slope compensators for 10Gb/s using all-pass multicavity etalons", *IEEE Phot. Technology Letters*, vol. 15, no. 5, 730-732 (2003).
94. L.M. Lunardi et al., "Tunable dispersion compensators based on multi-cavity all-pass etalons for 40Gb/s systems", *J. Lightwave Technology*, vol. 20, (12) 2136 (2002).
95. Wang, C. et al. Integrated lithium niobate electro-optic modulators operating at CMOS-compatible voltages. *Nature* **562**, 101 (2018).
96. Cisco Visual Networking Index: Forecast and Methodology, 2016–2021 (Cisco, September 2017);
www.cisco.com/c/en/us/solutions/collateral/serviceprovider/visual-networking-index-vni/complete-white-paper-c11-481360.html

97. Winzer, P.J. & Neilson, D.T., From scaling disparities to integrated parallelism: A decathlon for a decade, *J. Lightwave Technol.*, **35**, 1099–1115 (2017).
98. Hu, H. et al., Single-source chip-based frequency comb enabling extreme parallel data transmission, *Nat. Photon.*, **12**, 469 (2018).
99. Ataie, V. et al. Ultrahigh count coherent WDM channels transmission using optical parametric comb-based frequency synthesizer, *J. Lightwave Technol.* **33**, 694–699 (2015).
100. Hillerkuss, D. et al. 26 Tbit s⁻¹ line-rate super-channel transmission utilizing all-optical fast Fourier transform processing, *Nature Photonics* **5**, 364–371 (2011).
101. Fülöp, A., et al., V., High-order coherent communications using modelocked dark-pulse Kerr combs from microresonators, *Nat. Comm.*, **9**, 1598 (2018)
102. Pfeifle, J. et al., Optimally Coherent Kerr Combs Generated with Crystalline Whispering Gallery Mode Resonators for Ultrahigh Capacity Fiber Communications, *Phys. Rev. Lett.* **114**, 093902 (2015)
103. Gaeta, A.L., Lipson, M., Kippenberg, T.J., Photonic-chip-based frequency combs, *Nature Photonics* **13**, 158–169 (2019).
104. Lugiato, L. A., Prati, F. & Brambilla, M. *Nonlinear Optical Systems*, (Cambridge University Press, 2015).
105. Wang, W., et al., Robust soliton crystals in a thermally controlled microresonator, *Opt. Lett.*, **43**, 2002 (2018).
106. Bao, C., et al., Direct soliton generation in microresonators, *Opt. Lett.*, **42**, 2519 (2017).
107. A. Aadhi et al., “Highly reconfigurable hybrid laser based on an integrated nonlinear waveguide”, *Optics Express* **27** (18) 25252 (2019). DOI: 10.1364/OE.27.025251.
108. Zhou, H., et al., Soliton bursts and deterministic dissipative Kerr soliton generation in auxiliary-assisted microcavities, *Light: Science and Applications*, **8**, 50 (2019).
109. Kang, Z., et al., Deterministic generation of single soliton Kerr frequency comb in microresonators by a single shot pulsed trigger, *Opt. Express*, **26**, 326548 (2018).
110. Del'Haye, P., et al., “Optical frequency comb generation from a monolithic microresonator”, *Nature* **450**, 1214–1217 (2007).
111. Kippenberg, T. J., Holzwarth, R. & Diddams, S. A. Microresonator-based optical frequency combs. *Science* **332**, 555–559
112. Xue, X. et al., Mode-locked dark pulse Kerr combs in normal-dispersion microresonators. *Nat. Photonics* **9**, 594–600 (2015).
113. Kang, Z., et al., Deterministic generation of single soliton Kerr frequency comb in microresonators by a single shot pulsed trigger, *Opt. Express*, **26**, 326548 (2018)
114. Raja, A.S., et al., Electrically pumped photonic integrated soliton microcomb, *Nat. Comms.*, **10**, 680 (2019)
115. Pavlov, N.G., Narrow-linewidth lasing and soliton Kerr microcombs with ordinary laser diodes, *Nat. Photon.*, **12**, 694 (2018)
116. Cisco Corp., Cisco Visual Networking Index: Forecast and Trends, 2017–2022 White Paper, (accessed 6 Aug. 2019) www.cisco.com/c/en/us/solutions/service-provider/visual-networking-index-vni/index.html#~complete-forecast
117. S. Grubb, et al., Field trials: Is it make or break for innovative technologies?, *Proc. OFC 2018*, S2A (2018)
118. K. Schuh, et al., Single Carrier 1.2 Tbit/s Transmission over 300 km with PM-64 QAM at 100 GBaud, *Proc. Optical Fiber Communications (OFC), Th5B.5*, San Diego, CA (2017).
119. Alvarado, A., et al., Replacing the Soft-Decision FEC Limit Paradigm in the Design of Optical Communication Systems, *J. Lightwave Technol.* **34**, 707 (2016).
120. Torres-Company, V., et al., Laser Frequency Combs for Coherent Optical Communications, *J. Lightwave Technol.* doi: 10.1109/JLT.2019.2894170 (2019).
121. Obrzud, E., Lecomte, S. & Herr, T., Temporal solitons in microresonators driven by optical pulses, *Nat. Photon.*, **11**, 600 (2017).
122. Papp, S.B., et al., Microresonator frequency comb optical clock, *Optica* **1**, 10 (2014).
123. Puttnam, B., et al., 2.15 Pb/s Transmission Using a 22 Core Homogeneous Single-Mode Multi-Core Fiber and Wideband Optical Comb, *Proc. European Conference on Optical Communications (ECOC), PDP 3.1*, Valencia (2015).
124. Calculation of system performance metrics. After signal reconstruction using DSP, we measured system performance based on 3 metrics: bit-error rate (BER), error-vector magnitude (EVM) and generalized mutual information (GMI). BER is measured by decoding a 1.1-Mbit-long random bit sequence, grey-coded onto the 64-QAM constellation, so a BER of 9×10^{-5} provides 100 errors. Error-vector magnitude gives an alternative metric directly related to the effective signal-to-noise ratio (SNR) at the receiver in the presence of uniform Gaussian noise. We use EVM to calculate signal quality factor (Q2 [dB]) in Figs 8,9 as $20 \log_{10}(1/\text{EVM}^2)$, with $\text{EVM} = 1/\text{SNR}^{0.5}$. For high-order modulation formats, generalized mutual information (GMI) gives more accurate performance measure than BER (assuming a soft-decision forward error correction threshold). We use GMI as the key performance figure (i.e. net data rate and spectral efficiency) where the achievable capacity (b/s) is the sum of individual channel GMIs (b/symbol), multiplying by the symbol rate (symbols/s). Spectral efficiency (b/s/Hz) can also be calculated from GMI, by taking the mean of the channel GMIs (b/symbol), multiplying by the symbol rate (symbols/s) and dividing by channel spacing (Hz)
125. Ionescu, M., et al., 91 nm C+L hybrid distributed Raman–erbium-doped fibre amplifier for high capacity subsea transmission, *Proc. ECOC 2018*, doi:10.1109/ECOC.2018.8535151 (2018).
126. R. Maher, Signal processing for high symbol rate transmission: Challenges and Opportunities, *Proc. Adv. Photonics 2018*, SpW3G.2 (2018).
127. L. Galdino et al., The trade-off between transceiver capacity and symbol rate, *Proc. OFC 2018*, W1B.4 (2018)
128. Olsson, S.L.I. et al., Record-high 17.3-bit/s/Hz spectral efficiency transmission over 50 km using probabilistically shaped PDM 4096-QAM, *Proc. OFC 2018*, Th4C.5 (2018).
129. Chen, X., et al., 16384-QAM transmission at 10 Gbd over 25-km SMF using polarization-multiplexed probabilistic constellation shaping, *Proc ECOC 2019*, PD.3.3 (2019).

130. Winzer, P. J., Neilson, D. T. & Chraplyvy, A. R., “Fiber-optic transmission and networking: the previous 20 and the next 20 years”, *Opt. Express* **26**, 24190–24239 (2018).
131. A. Della Torre et al., “Mid-infrared supercontinuum generation in a low-loss germanium-on-silicon waveguide”, *APL Photonics* Vol. 6, 016102 (2021); doi: 10.1063/5.0033070.
132. M. Sinobad, et al., “Mid-infrared supercontinuum generation in silicon-germanium all-normal dispersion waveguides”, *Optics Letters*, Vol. 45 (18), 5008-5011 (2020). DOI: 10.1364/OL.402159.
133. M. Sinobad et al., “High coherence at f and 2f of a mid-infrared supercontinuum in a silicon germanium waveguide”, *IEEE Journal of Selected Topics in Quantum Electronics* Vol. 26 (2) 8201008 (2020). DOI:10.1109/JSTQE.2019.2943358.
134. M. Sinobad et al., “Dispersion trimming for mid-infrared supercontinuum generation in a hybrid chalcogenide Si-Ge waveguide”, *Journal of the Optical Society of America B*, Vol. 36 (2) A98-A104 (2019). DOI: 10.1364/JOSAB.36.000A98.
135. M. Sinobad et al., “High brightness mid-infrared octave spanning supercontinuum generation to 8.5 μ m in chip-based Si-Ge waveguides”, *Optica*, Vol. 5 (4) 360-366 (2018). DOI:10.1364/OPTICA.5.000360.
136. L. Jin et al., *Applied Physics Letters Photonics*, Vol, 5 Article 056106, (2020). DOI:10.1063/5.0002941
137. L. Carletti et al., “Nonlinear optical properties of Si-Ge waveguides in the mid-infrared”, *Optics Express* Vol. 23 (7) 8261–8271 (2015).