






# Ultrafast optical circuit switching for data centers using integrated soliton microcombs

Arslan Sajid Raja<sup>1,3</sup>, Sophie Lange<sup>2,3</sup>, Maxim Karpov <sup>1,3</sup>, Kai Shi<sup>2,3</sup>, Xin Fu<sup>1</sup>, Raphael Behrendt<sup>2</sup>, Daniel Cletheroe<sup>2</sup>, Anton Lukashchuk <sup>1</sup>, Istvan Haller<sup>2</sup>, Fotini Karinou<sup>2</sup>, Benn Thomsen<sup>2</sup>, Krzysztof Jozwik<sup>2</sup>, Junqiu Liu <sup>1</sup>, Paolo Costa<sup>2</sup>, Tobias Jan Kippenberg <sup>1</sup>✉ & Hitesh Ballani <sup>2</sup>✉

Due to the slowdown of Moore's law, it will become increasingly challenging to efficiently scale the network in current data centers utilizing electrical packet switches as data rates grow. Optical circuit switches (OCS) represent an appealing option to overcome this issue by eliminating the need for expensive and power-hungry transceivers and electrical switches in the core of the network. In particular, optical switches based on tunable lasers and arrayed waveguide grating routers are quite promising due to the use of a passive core, which increases fault tolerance and reduces management overhead. Such an OCS-network can offer high bandwidth, low network latency and an energy-efficient and scalable data center network. To support dynamic data center workloads efficiently, however, it is critical to switch between wavelengths at nanosecond (ns) timescales. Here we demonstrate ultrafast OCS based on a microcomb and semiconductor optical amplifiers (SOAs). Using a photonic integrated  $\text{Si}_3\text{N}_4$  microcomb, sub-ns (<520 ps) switching along with the 25-Gbps non-return-to-zero (NRZ) and 50-Gbps four-level pulse amplitude modulation (PAM-4) burst mode data transmission is achieved. Further, we use a photonic integrated circuit comprising an Indium phosphide based SOA array and an arrayed waveguide grating to show sub-ns switching (<900 ps) along with 25-Gbps NRZ burst mode transmission providing a path towards a more scalable and energy-efficient wavelength-switched network for data centers in the post Moore's Law era.

<sup>1</sup>Swiss Federal Institute of Technology Lausanne (EPFL), CH-1015 Lausanne, Switzerland. <sup>2</sup>Microsoft Research, 21 Station Road, Cambridge CB1 2FB, UK. <sup>3</sup>These authors contributed equally: Arslan Sajid Raja, Sophie Lange, Maxim Karpov, Kai Shi. ✉email: [tobias.kippenberg@epfl.ch](mailto:tobias.kippenberg@epfl.ch); [hitesh.ballani@microsoft.com](mailto:hitesh.ballani@microsoft.com)

The recent growth in data center traffic due to emerging applications such as machine learning, resource disaggregation, and large-scale data analytics, poses tight requirements on the network in terms of low latency, high bandwidth, and high scalability. Current data center networks comprise a multi-tier fabric of electrical switches and optical transceivers that consume a significant amount of power and this is expected to exacerbate due to the slowdown of Moore's law<sup>1,2</sup>.

Optical circuit switching (OCS) has been proposed as an alternative technology to overcome these challenges; it can provide high bandwidth and low network latency (due to the lack of buffers in the network core), and by allowing for a flat data center topology, the network can be more energy-efficient as fewer electrical switches and transceivers are required<sup>3</sup>. In OCS the servers are optically interconnected directly without needing the opto-electronic-opto (O-E-O) conversion. Micro-electro-mechanical (MEMS) based OCS systems are commercially available with the possibility of high port counts (~1000 ports)<sup>4,5</sup>. However, MEMS-based OCS architectures suffer from slow switching time (~ms), which can result in severe bandwidth under-utilization. Some emerging workloads are characterized by very small packets, e.g., in a recent key-value store application, more than 97% of the packets have a size of 576 bytes or less<sup>6</sup>. Assuming today's 50-Gbps speeds, it takes <93 ns to transmit a 576-byte packet. Therefore, achieving a reconfiguration overhead lower than 10% would require a nanosecond switching time.

Many photonic chip-based techniques have been used to demonstrate the fast optical switching (ns) such as semiconductor optical amplifiers (SOAs)<sup>7</sup>, electro-absorption modulators (EAMs)<sup>8</sup>, and Mach-Zehnder interferometers (MZIs)<sup>9,10</sup>. The arrayed waveguide grating routers (AWGRs)<sup>11,12</sup>, in conjunction with a tunable laser (TL), is a promising OCS technology as the core of the network is passive. This improves fault tolerance and it is also beneficial for future scaling because, unlike with today's electrical switches, the core may not need to be upgraded when data rates increase. However, state-of-the-art tunable lasers with custom drivers can have a tuning latency of 10 ns or higher<sup>13,14</sup> while custom lasers and drivers have been combined to achieve a latency of around 5 ns<sup>15,16</sup>, which make them unsuitable for our targets. A recent proposal addresses this issue by showing that using a disaggregated laser architecture in which the wavelength generation is decoupled from the wavelength selection, it is possible to achieve sub-nanosecond laser tuning time<sup>17,18</sup>. One possible implementation of this architecture relies on a discrete bank of lasers as a multi-wavelength source and a wavelength selector containing photonic integrated arrayed waveguide grating (AWG) and SOAs. The SOAs are preferred as switching elements due to fast and lossless switching, chip-scale integration, and high extinction ratio (ER)<sup>7</sup>.

In this paper we propose and experimentally demonstrate a disaggregated tunable transceiver that uses photonic chip-based soliton microcombs as a multi-wavelength source, which are generated in high-quality (*Q*) microresonators exhibiting third-order nonlinearity ( $\chi^3$ ) and anomalous dispersion<sup>19,20</sup>. The soliton microcombs have been used in many system-level applications, for example, distance ranging (LiDAR)<sup>21–23</sup>, microwave photonics<sup>24–27</sup>, optical coherence tomography (OCT)<sup>28,29</sup>, and coherent communication<sup>30–32</sup>. The remarkable improvements in the losses of high-confinement Si<sub>3</sub>N<sub>4</sub> integrated waveguides, a widely used platform for integrated photonics<sup>33,34</sup>, allows the generation of solitons using the on-chip laser<sup>35–37</sup>. The broadband-bandwidth operation across C- and L-bands, precise spacing control to match ITU frequency grid (50/100 GHz), low power consumption, and wafer-scale scalability make the soliton microcomb an attractive solution for an ultra-fast wavelength switched system in comparison with a bank of lasers<sup>30</sup>. Further, no extra guard band is required for frequency stabilization or to align each comb channel to the frequency grid. A key advantage

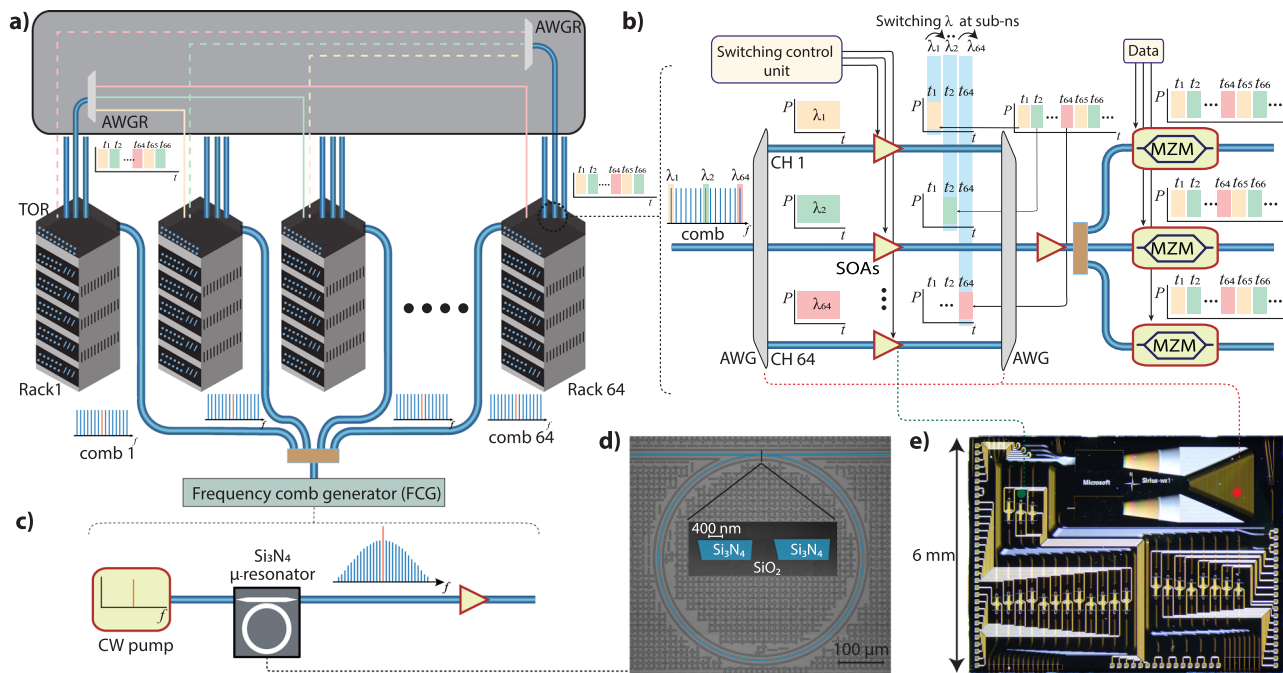
of the disaggregated transceiver design is that the multi-wavelength comb source can be shared across many (e.g. 64) nodes (racks or servers) in parallel instead of using 64 separate comb sources, using a split-and-amplify architecture (Fig. 1a). The source can thus be treated as a shared infrastructure element such as the power source in today's data centers. This allows for an appealing division of functionality since the power consumed by the comb source is amortized as the power efficiency of the end-to-end system converges to the efficiency of the amplifiers while allowing for a high-quality and stable light source that can be rapidly wavelength-tuned (cf. discussion on power analysis and Supplementary Information (SI)).

Here, we use a Si<sub>3</sub>N<sub>4</sub> based soliton microcomb as a multi-wavelength source to show ultrafast optical wavelength switching. In a proof-of-concept experiment, we demonstrate switching of >20 individual comb channels at sub-ns time scale (<520 ps) using discrete SOAs. Further, 25-GBd (baud-rate) non-return-to-zero on-off keying (NRZ) and four-level pulse amplitude modulation (PAM-4) burst mode transmission systems along with ultrafast switching are shown. Then, a more compact switching system consisting of PIC-based AWG and SOAs is implemented to show sub-ns switching and 25-GBd NRZ burst mode transmission, indicating the potential utilization of such a miniaturized system to mitigate power and scaling issues.

## Results

**OCS architecture based on soliton microcomb.** Figure 1 shows the OCS architecture containing the soliton microcomb, SOAs, AWGs in the switch and AWGRs to route the data across many racks (cf. Methods). This architecture allows further parallelization of different resources by sharing the soliton across many racks and modulators for power efficiency and parallel data transmission, respectively (Fig. 1a, b). The multi-wavelength source is generated by pumping a packaged Si<sub>3</sub>N<sub>4</sub> microresonator fabricated using the photonic Damascene reflow<sup>38,39</sup> process enabling a mean intrinsic *Q*-factor (*Q*<sub>0</sub>) of >15 million. Initially, a multi-soliton is initiated by performing a scan over resonance (forward tuning)<sup>20</sup> and then a single soliton is generated via backwards switching<sup>40</sup> (Fig. 2b). The soliton is amplified using a low-noise and compact EDFA resulting in a comb with a maximum power of up to -4 dBm and an optical signal-to-noise ratio (OSNR) of >34 dB. The post-amplified soliton as shown in Fig. 2c is de-multiplexed using a 100G spaced 48-channel AWG (~1525–1564 nm) providing 30 dB isolation. The individual comb channels are initially switched using discrete SOAs with a small-signal gain of ~11–13 dB at 1550 nm. Figure 2d shows 10–90% rise and fall times of 493 and 395 ps, respectively, for a single microcomb carrier centered at 1555 nm (CH 37 of AWG) when applying a current of 120 mA to operate the SOA. Similarly, more than 20 comb channels (1540–1564 nm) in C-band are tested individually to show sub-ns switching (cf. SI). Even though not tested due to the unavailability of an L-band AWG, the current results indicate that more than 40 comb channels in the L-band could be used as well.

By scaling to thousands of nodes, we can take advantage of the fact that today's servers or top-of-the-rack (ToR) switches comprise multiple channels. For example, the recently announced NVIDIA Ampere A100 GPU<sup>41</sup>, supports 2.4 Tbps of bandwidth by combining 48 channels, each operating at 50 Gbps. Therefore, assuming 40 wavelengths and SOAs per transmitter, we can connect up to 48 × 40 = 1920 other nodes. If, instead, we consider a rack-based deployment, the latest ToR switches<sup>42</sup> have 512 SERDESes (i.e., 256 uplinks) that could allow to interconnect up to 256 × 40 = 25,600 racks, which is an order of magnitude higher than the number of racks in even large data centers today.



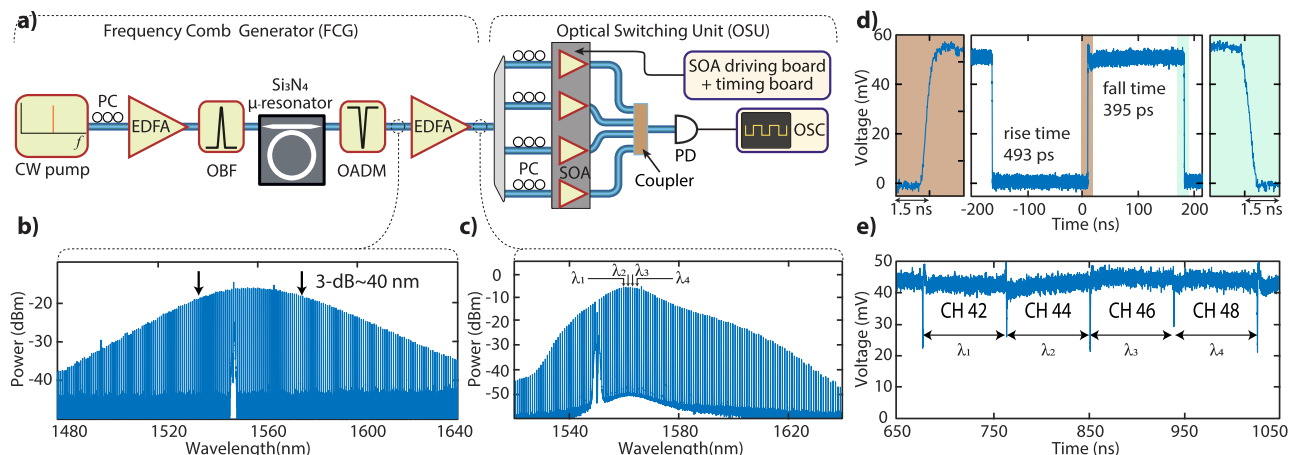
**Fig. 1 Concept of optical circuit switching (OCS) using a photonic chip-based  $\text{Si}_3\text{N}_4$  soliton microcomb and semiconductor optical amplifiers (SOAs).** **a** Interconnection of 64 racks via an arrayed waveguide grating router (AWGR) for implementing fast OCS. In this model, distinct wavelengths are assigned to each rack at each time slot. At each receiver, a 10–20 nanosecond (ns) time slot is assigned for each rack on a round-robin basis for data transmission. The switching module is placed on the top of rack (TOR) switch. A single comb source that is post amplified via cascaded amplifiers to attain high optical power per line can be distributed among many racks. The 64 individual combs (comb1, ..., comb 64) split from a central frequency comb generator (FCG) can be distributed across 64 different racks as a multiwavelength laser, making this architecture more power-efficient and flexible. The multiple data-carrying optical carriers are routed using a passive AWGR to the assigned racks. **b** Each comb channel is transmitted to SOAs after de-multiplexing, where a control signal (turning on/off current) is applied to switch between the comb channels at sub-ns. The comb channels (10-ns slots) are encoded with data using Mach-Zehnder modulators (MZMs) and transmitted to the relevant racks. The multiple MZMs shown here indicate that this architecture can be scaled further to establish links between more racks. **c** The multi-wavelength source based on the chip-scale soliton microcomb is generated by pumping with a single laser. Microscope images of a  $\text{Si}_3\text{N}_4$  microresonator (**d**) and a photonic chip (**e**) containing an AWG and SOAs. The inset in (**d**) shows a false color SEM image of a  $\text{Si}_3\text{N}_4$  microresonator’s coupling section.

However, this also means that a node is directly connected to any other node in the data center through only one of its uplink channels. To ensure that any pair of nodes can still communicate with their full bandwidth, the network routes traffic between any pair of nodes through all other nodes in the network. Such detour routing imposes a throughput overhead although the throughput can be at most  $2\times$  worse than an ideal switch and thus, can be compensated by doubling the per-node network bandwidth. Furthermore, detour routing offers several advantages including the fact that it obviates the need for explicitly scheduling network traffic, which has been a key bottleneck for practical and deployable optical switching in data center. We provide more details of our network architecture and the trade-offs imposed by our topology and routing choice in a separate paper<sup>18</sup>.

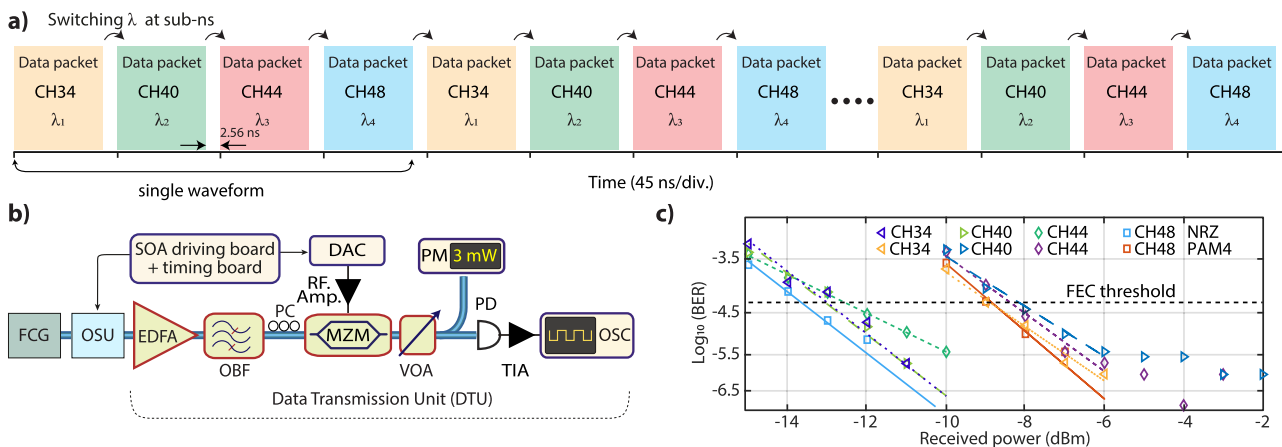
**Ultra-fast wavelength switching and data transmission via discrete SOAs.** For a proof-of-concept system-level demonstration, fast switching within four different comb channels is performed. Figure 2e shows sub-ns switching between four different comb channels with a wavelength separation of  $\sim 5.6$  nm. A guard zone of 2.56 ns is used to allow a smooth switching between adjacent comb channels while an external reference clock (timing board) aligns the different switching signals. Distinct currents are applied to the SOAs to achieve a constant output power and to compensate for the non-uniform comb power per line or SOA’s gain (Fig. 2e). While the sub-ns switching of four channels with a maximum 20 nm separation has been demonstrated, the

maximum channel separation is mainly limited by the optical band-pass filter (OBF) (cf. SI).

In the following experiment, we show 25 Gbps (NRZ) and 50 Gbps (PAM-4) burst mode data transmission while switching between four comb channels using the setup shown in Fig. 3b. The four optical carriers after switching are further amplified to overcome the insertion loss ( $\sim 7$  dB) of the 20-GHz Mach-Zehnder modulator (MZM) which is operating at the quadrature point. In addition to eliminating comb channels in the next order FSR of the AWG, the OBF is utilized to suppress the out-of-band SOA and EDFA amplified spontaneous emission (ASE) noise. The burst mode sequence at 25 GBaud symbol rate, generated by the arbitrary waveform generator is applied to the MZM with a random sequence of  $2^{15}$ - and  $2^{16}$ -bits for NRZ and PAM-4 respectively. The electrical waveform is amplified using a trans-impedance amplifier (TIA) after detecting it on a fast photodiode having 50-GHz bandwidth. The amplified waveform is acquired using a real-time oscilloscope with 160 GSamples/s sampling rate. The digital signal processing (DSP), explained in detail in ref. 17, is performed offline to obtain the bit error ratio (BER). The received optical power (ROP) vs. (BER) of the system, as shown in Fig. 3c, is characterized by changing the optical power of incoming waveform via a variable optical attenuator (VOA). A BER of below  $5 \times 10^{-5}$ , which is the threshold for forward error correction (FEC) in data center transmission systems corresponding to the KP4 FEC<sup>43</sup>, is achieved for both NRZ and PAM-4 at a ROP of  $\sim -12$  dBm and  $\sim -8$  dBm respectively. The BER error



**Fig. 2 Experimental demonstration of sub-ns optical circuit switching (OCS) of single and four different microcomb channels.** **a** A packaged  $\text{Si}_3\text{N}_4$  microresonator is pumped using a continuous wave (CW) pump amplified via an erbium-doped fiber amplifier (EDFA) to generate a soliton microcomb. A soliton state is initiated via forward tuning. After filtering the pump using an optical add-drop multiplexer (OADM), a low-noise amplifier is used to further amplify the soliton comb. The post-amplified comb is de-multiplexed using a 48-channels 100G AWG, and the individual comb channels are routed to semiconductor amplifiers (SOAs) which are controlled via a custom-designed switching board. The output signal is detected using a photodiode (PD) and recorded on an oscilloscope (OSC). PC polarization controller, OBF (narrow) optical bandpass filter. **b** A single soliton spectrum with 3 dB bandwidth around 40 nm and no post amplification. **c** Post-amplified soliton spectrum having maximum power up to  $-4$  dBm in comb lines close to pump line (1550 nm). **d** The 10%–90% rise and fall time of 493 and 395 ps respectively for single comb channel (CH 37: 1554.9 nm). The left and right insets show the zoomed-in view of rising and falling signals. **e** The four different microcomb channels simultaneously switching at sub-ns time scale with wavelength separation  $\sim 4.9$  nm. A guard zone of  $\sim 2.56$  ns is used to allow smooth switching between two adjacent channels. The ER is lower in comparison with **(d)** as signal intensity does not drop to zero while switching between adjacent channels (see “Methods”, Switching control unit). CH-42: 1558.988 nm, CH-44: 1560.613 nm, CH-46: 1562.238 nm, and CH-48: 1563.863 nm.



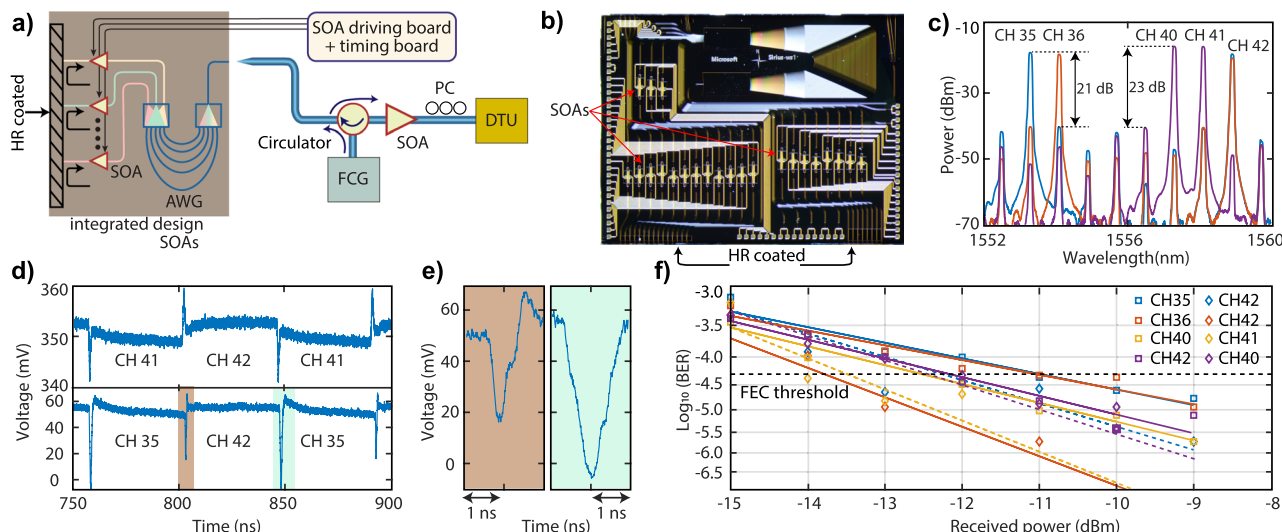
**Fig. 3 Experimental demonstration of burst mode NRZ and PAM-4 transmission using discrete SOAs while switching.** **a** A stream of multiple data packets showing data transmission along with sequential switching between the four comb channels. A single burst waveform sequence consists of header, payload, and guard zone containing 32, 1024, and 64 symbols respectively. **b** The signal after the frequency comb generator (FCG) and the optical switching unit (OSU) is amplified using a compact EDFA to compensate for losses of the 20-GHz Mach-Zehnder modulator (MZM). Then, it is filtered out using a wide-band optical bandpass filter (OBF) ( $\sim 20$  nm) to reject amplified spontaneous emission (ASE) noise from the SOAs and the EDFA. The data are encoded on the modulator using an arbitrary waveform generator (DAC). A fast photodiode (PD) is used to detect the signal. The electrical signal is amplified using a trans-impedance amplifier (TIA) and finally captured by an OSC. **c** The bit error ratio (BER) of four different comb channels while switching between them, using different modulation formats non-return to zero (NRZ) and four-level pulse amplitude modulation (PAM-4). A performance below forward error correction (FEC) threshold is achieved for NRZ and PAM-4 for a received optical power of  $\sim -12$  dBm and  $\sim -8$  dBm, respectively. CH-34: 1552.524 nm, CH-40: 1557.363 nm, CH-44: 1560.606 nm, and CH-48: 1563.863 nm.

floor for PAM4 at a ROP of  $> -6$  dBm emerges due to ASE and AWG crosstalk.

**Photonic integrated circuit based optical switching and data transmission.** Next, integrated Indium phosphide (InP) chip-based SOAs and AWG are used to show 25-Gbps NRZ burst mode data transmission along with fast OCS using a soliton microcomb. Figure 4b shows the photonic integrated circuit

(PIC) based wavelength selector PIC with a dimension of  $6 \text{ mm} \times 8 \text{ mm}$ . The reflection of the light from the high-reflection (HR) coated facet allows simultaneous utilization of an AWG with  $32 \times 50$  GHz separated channels as multiplexer and de-multiplexer. This simplifies the wavelength alignment procedure and reduces the footprint of the device. Nineteen SOAs are connected via integrated waveguides to AWG output channels. The wavelength alignment of the comb channels to the AWG is





**Fig. 4 Sub-ns optical circuit switching (OCS) and data transmission using on-chip SOAs and on-chip AWG along with soliton microcomb.** **a** Schematic of the setup used to perform the OCS and data transmission. The multi-wavelength optical carriers, generated via the frequency comb generator (FCG), are coupled to an InP chip containing an AWG and SOAs via an optical circulator. The coupled optical carriers are aligned to the AWG by changing the temperature of InP chip. The aligned carriers are transmitted to integrated SOAs; if one of them is biased, then the AWG channel (waveguide) connected to that particular SOA is reflected from the high-reflection coated facet of the chip while non-biased SOAs block the light. The reflected-back channels are coupled back via an anti-reflection coated optical fiber for encoding the information using the data transmission unit (DTU). **b** Microscope image of PIC showing the SOAs (red arrows) and AWG. **c** Optical spectrum of comb channels with different spacing while switching after the InP chip indicating more than 20 dB isolation with adjacent AWG channels (cross talk). The red curve shows CH 36 and 42 with 4.8 nm wavelength spacing, the blue curve CH 35 and 42 (5.6 nm) and the purple curve CH 40 and 41 (0.8 nm). **d** The sub-ns wavelength switching between two different comb channels using on-chip SOAs. The overshoot in the switching signal is due to impedance mismatch between the high-speed radio-frequency (RF) probes and the on-chip electrodes. This effect can be minimized by optimizing the drive signal. **e** The left and right figures show the zoomed-in view of switching signals between two different comb channels (CH35 and CH 42). **f** The bit error ratio (BER) performance of the 25 Gb/s NRZ PIC-based switching system for different combinations of two comb channels.

performed by changing the temperature of the PIC, resulting in seven comb channels matching with the AWG. This can also be realized by changing the temperature of the Si<sub>3</sub>N<sub>4</sub> chip. Figure 4c shows the optical spectrum of the AWG aligned comb channels indicating a >20 dB isolation with adjacent channels. Initially, the 10–90% rising and falling times of PIC is characterized by performing simultaneous switching between two comb channels. The maximum (minimum) experimentally observed switching time is ~820 ps (~375 ps). Moreover, the overshoot in the switching signal, as seen in Fig. 4d, arises due to impedance mismatch between the SOAs on-chip electrodes and the RF probes<sup>17</sup>. Then, a burst mode data transmission demonstration with 25-Gbps NRZ modulation is performed. A BER below FEC threshold is obtained when switching between two-channels with different combinations for a ROP >−11 dBm. While channels 40 and 41 have approximately the same optical power, channel 41 shows better BER performance as the crosstalk from channel 41 to 40 is about 8 dB lower than the crosstalk from channel 40 to 41. The PAM-4 burst mode transmission demonstration requires further improvement in the output power of the comb due to low in- and out-coupling in a packaged Si<sub>3</sub>N<sub>4</sub>. The main reason behind the low coupling efficiency (15%) is an additional 2 dB splicing loss between UHNA and SMF-28 fiber, which can be reduced to <0.2 dB by using state-of-art splicing instruments<sup>44</sup>. Similarly, the AWG crosstalk and insertion loss of the PIC can be further improved along with the utilization of lower FSR microcombs (25/50 GHz) to enhance the overall performance of this architecture. Nevertheless, the current results show the potential of a soliton microcomb as a suitable multi-wavelength source for 25 Gb/s NRZ burst mode transmission along with fast switching.

**Discussion**

Regarding the power consumption, the current multi-wavelength source consumes a total electrical power of ~30 W (cf. SI) providing more than 60 carriers, having an optical power >−20 dBm (~500 mW electrical power per carrier). The electrical power consumption can be improved down to <193 mW per carrier (15.5 W total) by reducing the splicing loss between UHNA and SMF fibers<sup>44</sup>, implementing on-chip actuators<sup>45</sup> instead of a bulk temperature controller and using a power-efficient, compact distributed feedback (DFB) laser as CW pump<sup>35,37</sup>. By optimizing the microresonator dispersion design, it is possible to generate 122 comb channels having an optical power >−14 dBm without needing any post-amplification<sup>30</sup>. Furthermore, an amplifier configuration with a separate amplifier for the C- and L-band would give a comb with an optical power per line ~13 dBm as mentioned in SI of ref. 30.

More importantly, this high power comb source can be shared among multiple racks by adding a hierarchy of passive optical splitters and amplifiers for better power and resource utilization (Fig. 1a). The soliton microcomb source distributed among 32 racks provides carriers with P<sub>opt</sub> ~−4 dBm while consuming ~2.57 W (1.115 W) electrical power per rack by using a state of the art commercial EDFA (on-chip amplifier<sup>46</sup>) making it a highly power-efficient and flexible solution for the data center (cf. SI). More broadly, the flexibility of sharing the comb across many racks, with the fast wavelength selection done on the rack itself, means that the overall electrical power efficiency per channel approaches the power efficiency of the EDFAs with the comb power as only a small contributor. This indicates that an optimized shared comb source would consume a comparable electrical power to other multi-wavelength source solutions, including recent techniques that use a bank of tunable lasers as a

multi-wavelength source<sup>18,47</sup>. Since the wavelength tuning is done on the bank of tunable lasers itself, sharing it between multiple racks would lead to an increased complexity when synchronizing the wavelength switching between the racks, due to the varying time delays between the bank of tunable lasers and the racks. As a matter of fact, one tunable laser bank per rack would be more optimal for a time-multiplexed solution<sup>47</sup> which would require at least  $2 \times 32$  tunable lasers for switching between 32 racks instead of a single amplified comb chip. Moreover, a comb does not require additional complex algorithms for fast switching and wavelength stabilization, thus offering an appealing division of functionality by leveraging a complex yet highly shared light source.

In this work, we demonstrate the possibility of achieving nanosecond OCS using a chip-based microcomb for future power-efficient and low latency data centers. More than twenty individual comb channels in C-band having a power  $> -20$  dBm are switched at  $< 520$  ps using discrete SOAs. The OCS system with 25-GBd NRZ and PAM-4 burst mode data transmission is shown while switching between different comb channels. Further, a PIC containing on-chip SOAs and an AWG is implemented to show sub-ns switching ( $< 900$  ps) and 25-GBd NRZ transmission. The current demonstration can provide a route for a fully integrated, fast-tunable transceiver providing dense carriers for wavelength switching to meet the power and latency requirements posed by future cloud workloads.

## Methods

**Switching architecture.** The link between two racks is only established via a single wavelength (comb tooth) in specific time slots ( $t_{64}, t_{128}, \dots$ ) as shown in Fig. 1a. The soliton microcombs provide many coherent optical carriers assigned to distinct racks. The switching operation is performed by applying a control signal on the SOAs, e.g., the switching from  $t_1$  to  $t_2$  data slot is done by applying an on signal to the second SOA and off signals to all other SOAs. A trigger signal from an external reference clock is used to align the switching control and data-encoding units.

**Si<sub>3</sub>N<sub>4</sub> microresonator.** The Si<sub>3</sub>N<sub>4</sub> microresonators are fabricated by using photonic Damascene reflow process enabling waveguides with ultra-smooth sidewalls and linear propagation loss  $\sim 1$  dB/m<sup>27,38</sup>. The designed microresonators are over-coupled ( $\kappa_{ex} \sim 4 \times \kappa_0$ ) with a FSR of 99.5 GHz and intrinsic linewidth of 15 MHz<sup>39</sup>. The waveguides have a dimension of  $1500 \times 900$  nm<sup>2</sup>. The double inverse nano tapers are used to facilitate the light coupling in- and out- of the chip<sup>48</sup>.

**Soliton microcomb generation.** A compact fiber laser (Koheras BASIK) is amplified using an EDFA. Then, ASE noise is filtered out using a narrow optical bandpass filter. The Si<sub>3</sub>N<sub>4</sub> chip is packaged by splicing the ultrahigh numerical aperture (UHNA) fiber with standard SMF-28 fiber with chip through (fiber-chip-fiber) 15% coupling efficiency<sup>49</sup>. A single soliton is initiated at an input power of  $\sim 450$  mW in the bus waveguide by applying a custom-designed ramp voltage. The deterministic soliton initiation and backward tuning are controlled via a computer interface. The strong pump line is filtered out using an OADM. Then the soliton is amplified using a low-noise EDFA before de-multiplexing. The power conversion from CW-pump to single soliton is  $\sim 2\%$  in the current device. An intrinsically coherent single soliton state is preferred due to low line to line power variation in 3 dB bandwidth (smooth spectral).

**Photonic integrated SOAs and AWG.** The InP-based wavelength selector PIC incorporates 23 SOAs of which 4 are used as references. The other 19 SOAs are connected to an  $1 \times 32$  AWG which acts as a multiplexer and demultiplexer. One of the PIC facets is high-reflection coated so that the light is reflected back through the SOAs and AWG to the input waveguide. The reflective single AWG design reduces the footprint compared to using two AWGs and avoids wavelength misalignment of the AWGs. The wavelength selector PIC was designed using the JEPPIX foundry and fabricated at Fraunhofer HHI.

**Switching control unit.** The switching control unit supplies the bias currents and electrical switching signals to the SOAs. A negative voltage is applied to drive the SOAs at the zero-level. The current flowing through the SOAs at the zero-level was around 1  $\mu$ A. To optimize the switching between two different channels, the upcoming channel starts switching on while the current channel is still switching off. The ER appears to be degraded since the signal intensity shows the sum of the channels and does not drop to the zero-level during the switching event. The ER of

the waveform in Fig. 2d is  $\sim 15$  dB. It is difficult to estimate actual ER as the signal was measured using a sampling scope and zero level of signal is buried in noise. This unit also controls the clock and time synchronisation of the switching signals. A time synchronization having  $< 100$  ps accuracy is achieved in current study<sup>18</sup>.

## Data availability

The data used to produce the plots within this paper are available at <https://doi.org/10.5281/zenodo.4588562>.

## Code availability

The code used to produce the plots within this paper are available at <https://doi.org/10.5281/zenodo.4588562>.

Received: 15 November 2020; Accepted: 20 July 2021;

Published online: 15 October 2021

## References

- Ballani, H. et al. Bridging the last mile for optical switching in data centers. In *Optical Fiber Communication Conference, W1C.3* (OSA, 2018).
- Cheng, Q., Bahadori, M., Glick, M., Rumley, S. & Bergman, K. Recent advances in optical technologies for data centers: a review. *Optica* **5**, 1354–1370 (2018).
- Barry, L. P., Wang, J., McArdle, C. & Kilper, D. in *Optical Switching in Next Generation Data Centers* (eds. Testa, F. & Pavesi, L.) 23–44 (Springer International Publishing, 2018).
- Han, S., Seok, T. J., Quack, N., Yoo, B.-w. & Wu, M. C. Monolithic  $50 \times 50$  mems silicon photonic switches with microsecond response time. In *Optical Fiber Communication Conference (2014), paper M2K.2, M2K.2* (Optical Society of America, 2014).
- Kim, J. et al. 1100 x 1100 port MEMS-based optical crossconnect with 4-dB maximum loss. *IEEE Photonics Technol. Lett.* **15**, 1537–1539 (2003).
- Clark, K. et al. Sub-nanosecond clock and data recovery in an optically-switched data centre network. In *2018 European Conference on Optical Communication (ECOC)*, 1–3 (2018).
- Tanaka, S. et al. Monolithically integrated 8:1 SOA gate switch with large extinction ratio and wide input power dynamic range. *IEEE J. Quantum Electron.* **45**, 1155–1162 (2009).
- Segawa, T., Nada, M., Nakamura, M., Suzuki, Y. & Takahashi, R. An  $8 \times 8$  broadcast-and-select optical switch based on monolithically integrated EAM-gate array. In *39th European Conference and Exhibition on Optical Communication (ECOC 2013)*, 1–3 (2013).
- Dupuis, N. et al. Nanosecond photonic switch architectures demonstrated in an all-digital monolithic platform. *Opt. Lett.* **44**, 3610–3612 (2019).
- Dupuis, N. et al. An  $8 \times 8$  silicon photonic switch module with nanosecond-scale reconfigurability. In *Optical Fiber Communication Conference Postdeadline Papers 2020, paper Th4A.6, Th4A.6* (Optical Society of America, 2020).
- Yin, Y. et al. LIONS: An AWGR-based low-latency optical switch for high-performance computing and data centers. *IEEE J. Sel. Top. Quantum Electron.* **19**, 3600409–3600409 (2013).
- Xiao, X., Zhang, Y., Zhang, K., Proietti, R. & Yoo, S. J. B. Multi-FSR on-chip optical interconnects using silicon nitride AWGR. In *2019 Conference on Lasers and Electro-Optics (CLEO)*, 1–2 (2019).
- Lavrova, O., Rau, L. & Blumenthal, D. 10-gb/s agile wavelength conversion with nanosecond tuning times using a multisection widely tunable laser. *J. Lightwave Technol.* **20**, 712–717 (2002).
- Blumenthal, D. et al. All-optical label swapping networks and technologies. *J. Lightwave Technol.* **18**, 2058–2075 (2000).
- Simsarian, J., Larson, M., Garrett, H., Xu, H. & Strand, T. Less than 5-ns wavelength switching with an SG-DBR laser. *IEEE Photonics Technol. Lett.* **18**, 565–567 (2006).
- Funnell, A. C. et al. Hybrid wavelength switched-TDMA high port count all-optical data centre switch. *J. Lightwave Technol.* **35**, 4438–4444 (2017).
- Shi, K. et al. System demonstration of nanosecond wavelength switching with burst-mode pam4 transceiver. In *45th European Conference on Optical Communication (ECOC 2019)*, 1–4 (2019).
- Ballani, H. et al. Sirius: A flat datacenter network with nanosecond optical switching. In *Proc. Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication, SIGCOMM '20*, 782–797 (Association for Computing Machinery, New York, NY, USA, 2020).
- Kippenberg, T. J., Gaeta, A. L., Lipson, M. & Gorodetsky, M. L. Dissipative Kerr solitons in optical microresonators. *Science* **361** (2018).

20. Herr, T. et al. Temporal solitons in optical microresonators. *Nat. Photonics* **8**, 145 (2013).
21. Trocha, P. et al. Ultrafast optical ranging using microresonator soliton frequency combs. *Science* **359**, 887–891 (2018).
22. Suh, M.-G. & Vahala, K. J. Soliton microcomb range measurement. *Science* **359**, 884–887 (2018).
23. Riemensberger, J. et al. Massively parallel coherent laser ranging using a soliton microcomb. *Nature* **581**, 164–170 (2020).
24. Torres-Company, V. & Weiner, A. M. Optical frequency comb technology for ultra-broadband radio-frequency photonics. *Laser Photonics Rev.* **8**, 368–393 (2014).
25. Wu, J. et al. Rf photonics: An optical microcombs' perspective. *IEEE J. Sel. Top. Quantum Electron.* **24**, 1–20 (2018).
26. Suh, M.-G. & Vahala, K. Gigahertz-repetition-rate soliton microcombs. *Optica* **5**, 65–66 (2018).
27. Liu, J. et al. Photonic microwave generation in the x- and k-band using integrated soliton microcombs. *Nat. Photonics* **14**, 486–491 (2020).
28. Ji, X. et al. Chip-based frequency comb sources for optical coherence tomography. *Opt. Express* **27**, 19896–19905 (2019).
29. Marchand, P. J. et al. Soliton microcomb based spectral domain optical coherence tomography. *Nat. Commun.* **12**, 1–9 (2021).
30. Marin-Palomo, P. et al. Microresonator-based solitons for massively parallel coherent optical communications. *Nature* **546**, 274 (2017).
31. Fülöp, A. et al. High-order coherent communications using mode-locked dark-pulse kerr combs from microresonators. *Nat. Commun.* **9**, 1598 (2018).
32. Corcoran, B. et al. Ultra-dense optical data transmission over standard fibre with a single chip source. *Nat. Commun.* **11**, 2568 (2020).
33. Moss, D. J., Morandotti, R., Gaeta, A. L. & Lipson, M. New cmos-compatible platforms based on silicon nitride and hydex for nonlinear optics. *Nat. Photonics* **7**, 597 (2013).
34. Blumenthal, D. J., Heideman, R., Geuzebroek, D., Leinse, A. & Roeloffzen, C. Silicon nitride in silicon photonics. *Proc. IEEE* **106**, 2209–2231 (2018).
35. Stern, B., Ji, X., Okawachi, Y., Gaeta, A. L. & Lipson, M. Battery-operated integrated frequency comb generator. *Nature* **562**, 401–405 (2018).
36. Pavlov, N. G. et al. Narrow-linewidth lasing and soliton kerr microcombs with ordinary laser diodes. *Nat. Photonics* **12**, 694–698 (2018).
37. Raja, A. S. et al. Electrically pumped photonic integrated soliton microcomb. *Nat. Commun.* **10**, 1–8 (2019).
38. Pfeiffer, M. H. P. et al. Ultra-smooth silicon nitride waveguides based on the damascene reflow process: fabrication and loss origins. *Optica* **5**, 884–892 (2018).
39. Liu, J. et al. Ultralow-power chip-based soliton microcombs for photonic integration. *Optica* **5**, 1347–1353 (2018).
40. Guo, H. et al. Universal dynamics and deterministic switching of dissipative kerr solitons in optical microresonators. *Nat. Phys.* **13**, 94–102 (2017).
41. NVIDIA A100 Tensor Core GPU Architecture. (2020).
42. Broadcom. 25.6 Tb/s StrataXGS Tomahawk 4 Ethernet Switch Series. (2020).
43. Ieee standard for ethernet amendment 2: Physical layer specifications and management parameters for 100 gb/s operation over backplanes and copper cables. IEEE Std 802.3bj-2014 (Amendment to IEEE Std 802.3-2012 as amended by IEEE Std 802.3bk-2013) 1–368 (2014).
44. Yin, P. et al. Low connector-to-connector loss through silicon photonic chips using ultra-low loss splicing of SMF-28 to high numerical aperture fibers. *Opt. Express* **27**, 24188–24193 (2019).
45. Liu, J. et al. Monolithic piezoelectric control of soliton microcombs. *Nature* **583**, 385–390 (2020).
46. Juodawlkis, P. W. et al. High-power, low-noise 1.5- $\mu\text{m}$  slab-coupled optical waveguide (SCOW) emitters: physics, devices, and applications. *IEEE J. Sel. Top. Quantum Electron.* **17**, 1698–1714 (2011).
47. Gerard, T. et al. SWIFT: Scalable ultra-wideband sub-nanosecond wavelength switching for data centre networks. Preprint at <https://arxiv.org/abs/2003.05489> (2020).
48. Liu, J. et al. Double inverse nanotapers for efficient light coupling to integrated photonic devices. *Opt. Lett.* **43**, 3200–3203 (2018).
49. Raja, A. S. et al. Chip-based soliton microcomb module using a hybrid semiconductor laser. *Opt. Express* **28**, 2714–2721 (2020).

## Acknowledgements

This work was supported by the Microsoft Research PhD scholarship program under MRL Contract No. 2019-034. This material is based upon work supported by the Air Force Office of Scientific Research under Award No. FA9550-19-1-0250, and by Swiss National Science Foundation under grant agreement No. 176563 (BRIDGE). This work was supported by funding from the EU H2020 research and innovation programme under grant agreement 862528 (PhoMEC). The  $\text{Si}_3\text{N}_4$  samples were fabricated and grown in the Center of MicroNanoTechnology (CMi) at EPFL. We would also like to thank Neon Photonics for their high-port-count AWGR.

## Author contributions

A.L. and K.S. designed the  $\text{Si}_3\text{N}_4$  and InP samples, respectively. J.L. fabricated the  $\text{Si}_3\text{N}_4$  microresonator. A.S.R., S.L., and M.K. characterized the samples. M.K., X.F., and A.S.R. contributed to breadboard based soliton generator. S.L., A.S.R., and K.S. carried out switching communication experiments. R.B., D.C., I.H., F.K., B.T., and K.J. contributed to the optical transmission system design and optimization. K. S., S.L., and A.S.R. analyzed and processed the data. A.S.R., S.L., H.B., P.C., and T.J.K. wrote the manuscript, with the input from others. H.B., T.J.K., and P.C. supervised the project.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41467-021-25841-8>.

**Correspondence** and requests for materials should be addressed to Tobias Jan Kippenberg or Hitesh Ballani.

**Peer review information** *Nature Communications* thanks David Moss and the anonymous reviewers for their contribution to the peer review of this work.

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

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



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021