AI-optimised tuneable sources for bandwidth-scalable, sub-nanosecond wavelength switching

THOMAS GERARD,[1,*] CHRISTOPHER PARSONSON,[1] ZACHARAYA SHABKA,[1] BENN THOMSEN,[2] POLINA BAYVEL,[1] DOMANIÇ LAVERY,[1,3] AND GEORGIOS ZERVAS[1]

[1]*Optical Networks Group, Department of Electronic and Electrical Engineering, University College London, London, WC1E 7JE, UK*
[2]*Microsoft Research, 21 Station Road, Cambridge, CB1 2FB, UK*
[3]*Infinera Canada Inc, Ottawa, ON K2K 2X3, Canada*
*uceetmh@ucl.ac.uk

Abstract: Wavelength routed optical switching promises low power and latency networking for data centres, but requires a wideband wavelength tuneable source (WTS) capable of sub-nanosecond switching at every node. We propose a hybrid WTS that uses time-interleaved tuneable lasers, each gated by a semiconductor optical amplifier, where the performance of each device is optimised using artificial intelligence. Through simulation and experiment we demonstrate record wavelength switch times below 900 ps across 6.05 THz (122×50 GHz) of continuously tuneable optical bandwidth. A method for further bandwidth scaling is evaluated and compared to alternative designs.

1. Introduction

Data traffic within cloud data centres is forecast to grow exponentially, and demand for internet and data centre services has further multiplied during the Covid-19 pandemic [1,2]. This growth has led to renewed interest in optical switching for data centres, which can be used to create scalable networks with high bandwidth and low latency. Drawing on the principles of wavelength routed optical burst switching (WROBS) [3], many optically-switched architectures have been proposed that use fast wavelength tuneable sources (WTSs) to transmit data through a wavelength routed fabric [4–12]. The WTS can be used to create highly dynamic optical circuit switching connections, permitting flat, all-optical, fixed-latency, power-efficient networking. The number of wavelength channels supported by the WTS can determine network scalability and bandwidth per channel. Architectures often require full C-band WTSs (~100 wavelengths with 50 GHz spacing) [5,6,8,9,12], or even support for the S- and L-bands [4,9]. Similarly, the bandwidth of passive star networks continuously improves with WTS channel count [7,11]. We note that wavelength routed networks are often designed around the constraint of a limited number of wavelengths, indicating the importance of bandwidth-scalable WTSs. Furthermore, the WTS reconfiguration time sets the bandwidth granularity of the wavelength routed network. Real data centre traffic is dominated by bursty, destination diverse packets that persist for tens of nanoseconds; sub-ns switching has, therefore, been determined a necessity for fine-granularity scheduling and to maintain network throughputs above 95% [5,11,13,14].

However, a full C-band WTS capable of sub-ns switching has not previously been demonstrated. Tuneable lasers (TLs) based on distributed Bragg reflector (DBR) gratings are space and power efficient but require exact current control and take several ns to switch, even when optimised

[15,16]. An electro-optic laser demonstrated sub-ns 10:90% rise times, but required ~2 ns for complete 90:90% wavelength switching, and supported <20 nm of tuneable bandwidth [17].

Disaggregated WTSs have also been proposed and select static wavelength sources using on/off optical gates. For example, an array of laser diodes (LDs) can be individually gated using semiconductor optical amplifiers (SOAs), providing sub-ns rise times [13]. A similar approach used an optical comb source, where each wavelength was filtered using an arrayed waveguide grating (AWG) and then gated by an SOA [18]. Though fast, disaggregated WTSs scale linearly with channel count; this makes their required power and defect-free fabrication area impractical for scaling to hundreds of channels. In addition, SOAs and their associated drive systems exhibit an oscillatory response after a switch event, such that few-ns guard zones were still required between each wavelength-to-wavelength switching event.

Hybrid designs combine the best features of TLs and disaggregated sources. An example is a duplex transmitter, composed of two time-interleaved tuneable lasers gated by an electro-optic switch [19]. In this scheme the laser tuning time determines the minimum dwell time on each wavelength; this decides the minimum packet burst size. An early integrated demonstration achieved 60 ps switching between 6 wavelength channels, though with 30 ms dwell times [20]. A more recent demonstration targeted 70 channels and 45 ns dwell times using DBR lasers and SOAs as optical gates [21]. However, the unpredictable switching response of the DBR lasers (25 to 90 ns) resulted in over 25% of switching events exceeding the permitted dwell time, such that the complete WTS supported just 53 channels. Furthermore, the non-ideal response of the SOAs limited 10:90% rise times to 9 ns and 90:90% wavelength switch times to 13 ns, with significant intensity fluctuations throughout each dwell period. Idealising the response of each optical component is, therefore, a key challenge in constructing hybrid WTSs appropriate for data centre networks. A summary of the reported rise times, wavelength switch times, and channel counts of few-ns WTS demonstrations are shown in Fig. 1.
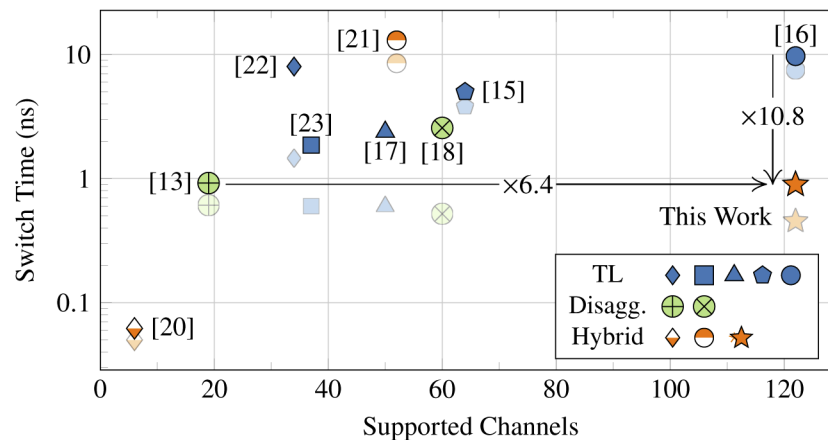

Fig. 1. Summary of ns and sub-ns WTS demonstrations, comparing switch time against number of supported 50 GHz spaced channels. 10:90% intensity rise times are shown faded, while 90:90% wavelength switch times are solid [13,15–18,20–23].

In this work, we propose a hybrid WTS that is capable of wideband, sub-nanosecond switching. The hybrid WTS contains multiple optical components, the performance of which are idealised through artificial intelligence (AI) methods. These methods require no prior device knowledge, and ensure fast and stable wavelength and power switching for every WTS deployed throughout a wavelength routed data centre. Moreover, the hybrid WTS is shown to be bandwidth scalable, alleviating the long-held constraint of finite wavelength count in wavelength routed architectures.

We experimentally demonstrate the AI-optimised hybrid WTS switching over 6.05 THz of continuous bandwidth with 90:90% wavelength switching times of 900 ps and below. Our design scales well in terms of on-chip area, yield, and implementation practicality, and, therefore, represents the first demonstration of bandwidth scalable, sub-nanosecond wavelength switching.

2. Operational principle

The principle of the hybrid WTS is shown in Fig. 2(a). Two fast tuneable lasers (TLs) are driven out of phase, such that one is lasing while the other is tuning. Each TL is optically gated by an SOA that absorbs the spurious laser output during tuning and amplifies the steady-state operation. The SOA outputs are then time-interleaved, providing sub-ns wavelength switching across the full TL bandwidth. To give the TLs time to prepare, wavelength requirements must be known at least one burst in advance of transmission.
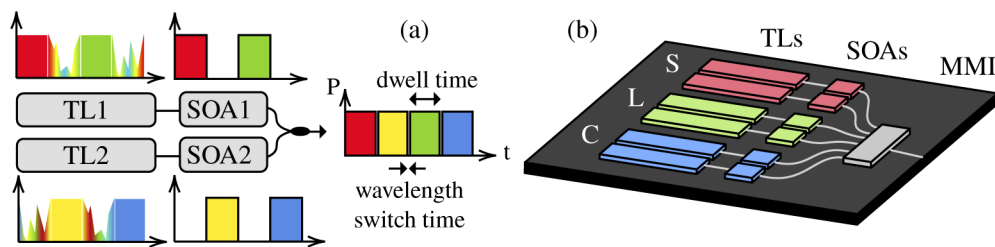

Fig. 2. (a) Operational principle of the hybrid WTS. (b) Principle of bandwidth scaling. Devices operating in different communication bands are multiplexed using an MMI coupler. To minimise loss, off-chip thin film multiplexers could also be used.

The shortest *dwell* time that can be supported by the hybrid WTS is determined by the worst-case TL tuning time; this is because the first laser must hold steady until the second has finished tuning. The *wavelength switch* time of the hybrid WTS is set by the rise time, fall time and alignment of the SOAs, determining the minimum inter-packet gap. Here, we measure the wavelength switch time as the 90:90% intensity transition time between wavelengths. To illustrate the timing requirements of different network traffic combinations, the dwell times and wavelength switch times supported by the hybrid WTS are shown in Fig. 3(a) and (b), respectively, for a range of line rates and data packet sizes. The results show that our AI-optimised system demonstration allows the hybrid WTS to support a wide variety of line rates with realistic packet sizes. Examples of potential network traffic demands are overlaid as markers. These include existing (Ethernet) and proposed (Stardust) inter-rack data centre traffic, but also GPU clusters for machine learning, as well as high speed memory transfer from solid state devices. More details on these traffic types can be found in Appendix A. This highlights the wide potential impact of the hybrid WTS, which can support the networking requirements of data centre racks, machine learning clusters, and even resource disaggregation in ambitious next-generation designs.

To scale the optical bandwidth of the hybrid WTS, TLs and SOAs operating in other communication bands (C,S,L) can be included, as shown in Fig. 2(b). This permits individual components to be optimised for their unique communication bands, offering greater combined bandwidth and lasing quality than that achievable using single devices [24]. Multiplexing can be achieved using a wideband multimode interference coupler (MMI) coupler [25]. The inherent gain of the SOAs, serving as fast optical gates, compensates for the optical loss of the MMI coupler. This design permits $N : 1$ multiplexing with reduced integration area, complexity, and optical loss compared to the electro-optic gates applied in previous hybrid sources [19,20].

As just four active devices are needed for each band, the hybrid WTS can offer greater practical bandwidth scalability than alternative sub-ns WTSs. To quantify this, the on-chip area and yields

of three sub-ns WTS designs were modelled and compared. The comparison WTSs were selected on the criteria that they have been shown to switch in under 1 ns, that they are (in principle) able to continuously increase the supported channel count. The first WTS, shown in Fig. 4(a), is an array of laser diodes each gated by an SOA, which are multiplexed using an AWG as proposed in [13]. The second WTS is shown in Fig. 4(b). The individual comb lines of an integrated comb source are separated using an AWG, selected for using an SOA array, then recombined using a second AWG; see [18]. Here the comb source is assumed to be composed of a seed laser and microring resonator, as described in [26], followed by a low-noise, high-gain booster SOA, as described in [27]. The scaling principle of the hybrid WTS is shown in Fig. 4(c). Note that, in contrast to the designs shown in Fig. 4(a) and (b), the hybrid WTS can be multiplexed using an MMI coupler for maximum chip-space efficiency. This is because the SOA gain can compensate for the loss of the low order coupler (joining 2-6 devices). However, a coarse thin-film filter could also be used multiplex each communication band off-chip, reducing excess loss to <1 dB [28]. We note that these three WTS designs are in consideration by major data centre operators for deployment in wavelength routed hyperscale networks [12].
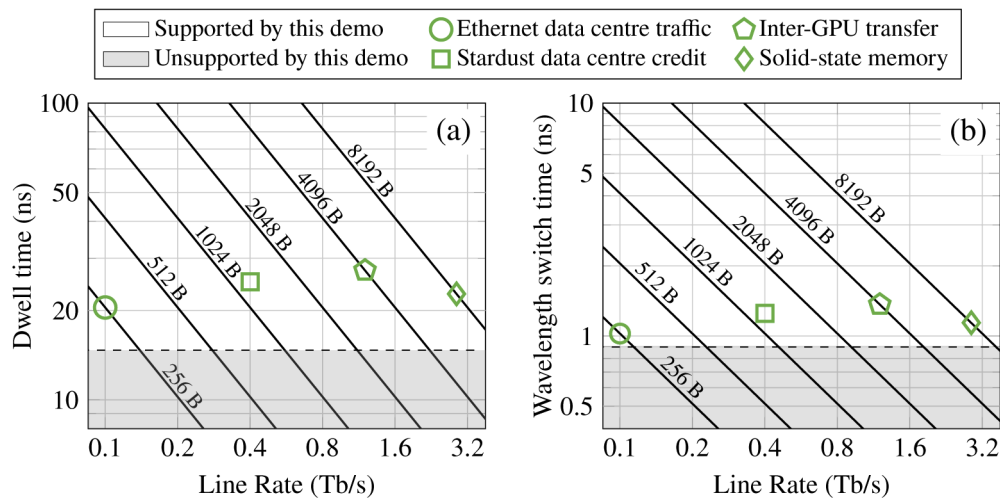

Fig. 3. (a) Required wavelength dwell time for different packet sizes at varying line rates. (b) Wavelength switch time for different packet sizes at varying line rates, allowing a 5% overhead. This system demo supports packets and line rates above the dotted lines.
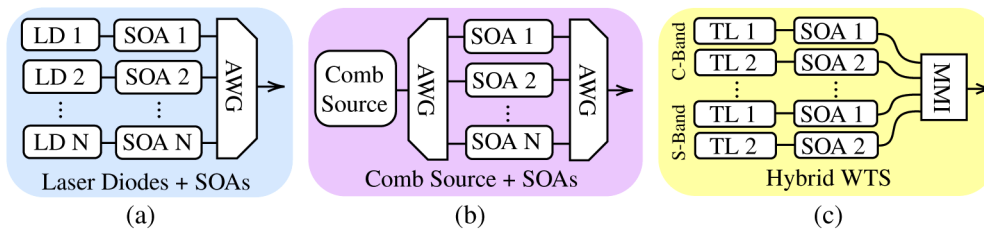

Fig. 4. Principle of channel scaling for three sub-ns WTSs. (a) An array of laser diodes (LD) gated by SOAs [13]. (b) An integrated comb source (seed laser, microring resonator and booster SOA) gated by SOAs [18]. (c) The proposed hybrid WTS, composed of 2 TLs and 2 SOAs per transmission window.

The device sizes used for each component when fabricated in indium phosphide (InP) are given in Table 1. These are as reported in the referenced publication, or measured from the published images and scale bars. It should be noted that the passive components considered in this model could potentially be fabricated in silicon to benefit from the higher refractive index contrast and mature fabrication platforms, reducing component footprint and defect rate. These could then be co-integrated with the optically-active InP components using III-V heterogeneous integration techniques [29–31]. In this work, analysis was restricted to the case where all components are fabricated monolithically in InP for ease of evaluation. The modeled WTS areas are given in Fig. 5(a). The hybrid WTS is observed to scale as a staircase function with the greatest area efficiency, and is always the most space efficient design.

Table 1. On-chip area occupied by different integrated InP devices.

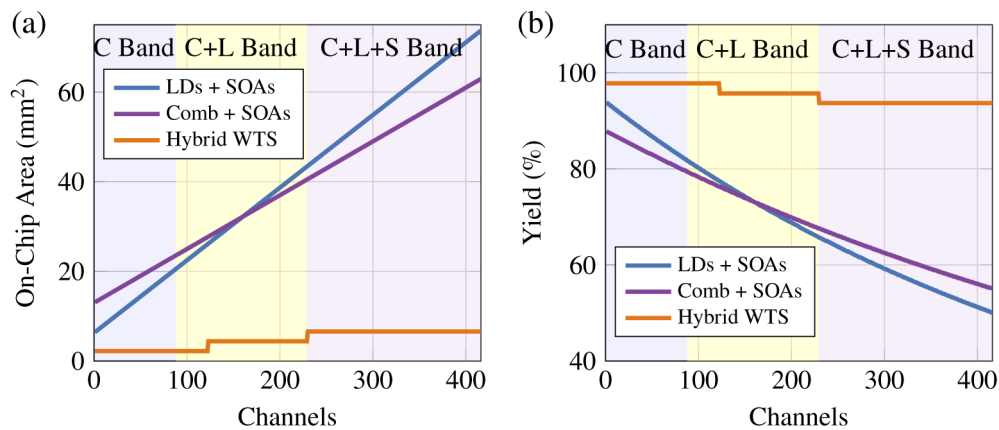| Device	Area (mm^2)	Ref.	Device	Area (mm^2)	Ref.
Laser diode	0.04	[32]	Booster SOA	6.3×10^{-3}	[27]
Switching SOA	0.12	[13]	AWG	6.27	[13]
Microring resonator	0.38	[26]	Tuneable laser	0.98	[33]
Seed laser	0.05	[26]	MMI coupler	3.9×10^{-3}	[34]

Fig. 5. Comparison of the integration area required to scale three sub-ns WTS designs, assuming ideal device packing and device areas given in Table 1.

Greater space efficiency permits (1) more complete devices per wafer, decreasing the cost-per-device, and (2) reduces the probability of encountering a major fabrication error, increasing device yield. This can be quantified using a Murphy's yield model for so called 'killer' random fabrication defects, such that the device yield is given by $Y_r = [(1 - \exp(-AD))/AD]^2$, where A is the device area and D is the defect rate per unit area [35]. Y_r was calculated for each of the WTS designs under consideration using the device areas calculated in Fig. 5(a). A state-of-the-art killer defect density for indium phosphide of $1\,\mathrm{cm}^{-2}$ was assumed [35]. The resulting yield vs WTS channel count is presented in Fig. 5(b). The yield of the WTS designs shown in Fig. 4(a) and (b) decreases quickly, such that the full C+L+S band devices have modeled yields of $\sim 53\%$. In practice, lower yields are expected, as the assumption of ideal-packing weakens as the component count increases. In contrast, the hybrid WTS has a yield of 94% for the same channel count. Furthermore, the assumption of ideal device packing is more accurate, as the C+L+S band hybrid WTS uses just 12 active devices and, therefore, the absolute impact of requiring additional space per device is low. Overall, the yield model shows that the hybrid WTS offers superior scalability

compared to the digital WTS designs, and is, therefore, the most practical design for widescale data centre integration and deployment.

3. Tuneable laser simulation

The unreliable wavelength tuning speed of wideband DBR lasers has been identified as a key challenge in constructing a hybrid WTS suitable for packet-timescale switching. Fast wavelength reconfiguration can be achieved by applying pre-emphasis to the drive sections of an integrated semiconductor laser, overshooting the target current values to accelerate the charge carrier settling time [36]. Previously we have demonstrated a regression optimiser that can autonomously calculate the pre-emphasis values required for reliable wavelength switching in a digital supermode DBR (DS-DBR) laser [16]. In this work, we present a finite difference time domain model of the DS-DBR laser to give physical insight to the operation of the optimiser. The model applies a simplified transfer matrix approach to simulate the frequency and time domain response of the laser over tens of nanoseconds when driven with time-varying injection currents. An illustration of the DS-DBR laser is shown in Fig. 6. An eight-section chirped front grating coarsely tunes over ~ 48 nm of bandwidth in ~ 7 nm steps, while the rear and phase sections provide fine tuning.
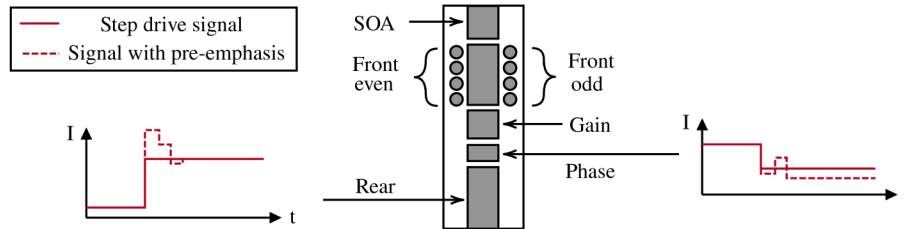

Fig. 6. Illustration of the DS-DBR laser. Fast synchronous arbitrary current sources can be applied to the front, rear and phase sections to achieve fast tuning. Pre-emphasis can be used to modify the drive signals to idealise wavelength tuning time.

3.1. Frequency domain model

The number of charge carriers within a semiconductor laser section is described by the laser rate equation, given by

$$\frac{dN}{dt} = \frac{I}{qV} - AN - BN^2 - CN^3, \tag{1}$$

where N is the carrier density, I is the applied current, q is the fundamental electron charge, V is the section volume, and A, B and C are constants that scale the linear recombination rate, bimolecular radiation rate, and Auger rate, respectively [37]. Here, we apply Eq. (1) to model the carrier density in each of the DBR gratings in the DS-DBR laser. The Bragg wavelength of each laser tuning section is given by $\lambda_{\text{Bragg}} = 2\Lambda n_{\text{eff}}$, where Λ is the grating pitch period and n_{eff} is the effective refractive index across the grating [37]. Fast wavelength tuning over tens of nanoseconds is principally achieved via the plasma effect; thermal effects acting on longer timescales are considered in Appendix C. When tuning using the plasma effect, it can be shown that n_{eff} decreases linearly with N [38], such that the change in laser frequency in section 'sec', df_{sec}, (where 'sec' is one of 'front', 'rear' or 'phase') can be modelled as

$$\frac{df_{\text{sec}}}{dN} = \frac{\Delta\lambda_{\text{max,sec}}c}{\lambda_{\text{c}}^2} \cdot \frac{1}{N_{\text{max,sec}}}, \tag{2}$$

where $\Delta\lambda_{\text{max,sec}}$ is the maximum possible change in wavelength in section 'sec', c is the speed of light, λ_{c} is the central reference wavelength, and $N_{\text{max,sec}}$ is the maximum carrier density

allowed by the laser section. $N = N_{\mathrm{max,sec}}$ can be calculated for each laser section by applying the maximum permitted current injection for that section, $I_{\mathrm{max,sec}}$, to Eq. (1) and numerically solving for N.

The DS-DBR laser uses a seven-peak frequency comb reflector in the rear section to achieve wideband tuning [33]. Here, the rear section frequency transfer function is modelled by adapting the equation for the electric field given in [39], such that

$$H_{\mathrm{rear}}(f, df_{\mathrm{rear}}) = \exp\left[-\left(\frac{f - df_{\mathrm{rear}}}{2\sigma}\right)^{2p}\right] \sum_{h_i=1}^{h} \cos\left(2\pi h_i f_{\mathrm{FSR,rear}}(f - df_{\mathrm{rear}})\right)^2,\tag{3}$$

where f is the frequency to be characterised and df_{rear} is the frequency shift applied to the rear section as calculated by Eq. (2). In Eq. (3), the summed cosine describes a frequency comb, where h is the number of harmonics and $f_{\mathrm{FSR,rear}}$ is the free spectral range of the comb lines, set equal to $1/\Delta\lambda_{\mathrm{max,rear}}$. The exponential term is a super Gaussian that serves as a top-hat filter for the frequency comb, where σ is the standard deviation and p is the order. σ specifies the width of the passband, defined as

$$\sigma = \frac{k + 1}{2} \cdot \frac{f_{\mathrm{FSR,rear}}}{2\sqrt{2\ln 2}},\tag{4}$$

where k is the number of comb lines inside the passband. The parameters h, $\Delta\lambda_{\mathrm{max,sec}}$, p and k are set to match experimentally reported values, and are given in Appendix B.

To select one of the seven rear section frequency comb lines, the DS-DBR laser uses a chirped front grating composed of eight DBR sections that can combine to create seven enhanced reflectance peaks. The corresponding transfer function, H_{front}, can be modelled as eight neighbouring Gaussian frequency windows, such that

$$H_{\mathrm{front}}(f, df_{\mathrm{front}}) = \left(\sum_{s=-S/2}^{S/2} \exp\left[-\left(\frac{f - df_{\mathrm{front,s}} - sf_{\mathrm{FSR,rear}}}{2\sigma}\right)^2\right]\right)^2,\tag{5}$$

where S is the number of front sections, df_{front} describes the frequency shifts applied to each front section, and $df_{\mathrm{front,s}}$ is the frequency shift applied specifically to front section s, calculated from Eq. (2). σ is again described by Eq. (4), though with $k = 1$. This approach permits a front section grating to be tuned via carrier injection until its reflectivity spectrum fully overlaps with that of its neighbouring front section, creating a enhanced reflectivity peak. The enhanced peak selects for one of the rear reflector peaks, permitting a single laser mode to dominate [33].

Together, Eq. (3) and Eq. (5) select for one of the longitudinal modes that the laser supports; this is fine tuned using the laser phase grating. The longitudinal transfer function is, therefore, given by

$$H_{\mathrm{long}}(f, df_{\mathrm{phase}}) = \cos\left(2\pi f_{\mathrm{FSR,long}}(f - df_{\mathrm{phase}})\right)^2,\tag{6}$$

where $f_{\mathrm{FSR,long}}$ is the longitudinal mode spacing, given as $1/\Delta\lambda_{\mathrm{max,long}}$, and df_{phase} is the tuning applied by the laser phase section, calculated using Eq. (2). Having established the transfer functions for the laser rear section, front section and longitudinal modes, the full description of the supported frequencies of the DS-DBR laser is given by

$$H_{\mathrm{laser}}(f, df_{\mathrm{rear}}, df_{\mathrm{front}}, df_{\mathrm{phase}}) = H_{\mathrm{rear}}(f, df_{\mathrm{rear}}) \cdot H_{\mathrm{front}}(f, df_{\mathrm{front}}) \cdot H_{\mathrm{long}}(f, df_{\mathrm{phase}}),\tag{7}$$

where it is assumed that the frequency with the highest amplitude in H_{laser} dominates.

To test the model, a simulated tuning map was generated in MATLAB by calculating H_{laser} for a range of injection currents. The injection currents for the rear, front and phase laser sections were varied from 0 mA up to their maximum current thresholds in increments of ΔI. For each

combination of injection currents, df_{rear}, df_{front} and df_{phase} were calculated, and then H_{laser} was evaluated for f centered at 1550 nm (193.41 THz) spanning 8.9 THz of bandwidth over 4096 frequency increments, giving a frequency resolution of 2.2 GHz. The key parameters used within this model are given in Appendix B, and are sourced from the reported results in [33,39,40].

The resulting tuning map is shown in Fig. 7(a). An experimental tuning map is shown in Fig. 7(b). The first 122×50 GHz spaced frequency channels are labelled on both maps as crosses. The maps show good agreement, following the same general shape and featuring a higher density of 50 GHz channels at lower rear section currents. Exact channel positions will be unique to each real laser based on its manufacturing, packaging, temperature and age. This result demonstrates the model can accurately translate laser section injection currents into dominant laser frequency.
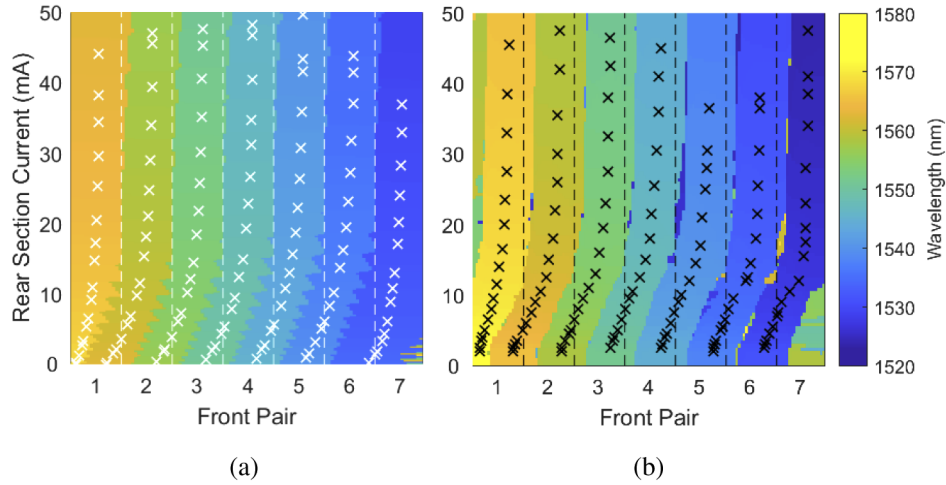

(a)   (b)

Fig. 7. (a) Simulated, and (b) experimental DS-DBR laser tuning maps.

3.2. Time domain model

The temporal dependence of laser frequency in response to time-varying injection currents can be modelled by numerically solving the laser rate equation. To achieve this, a step-change in injection current was considered within all 10 (8 front, phase, and rear) DS-DBR laser sections. The injection currents were considered to be applied by a digital to analogue converter (DAC) with 10 channels, N_{ch}, each of sample rate $F_{s,DAC}$, applied over $N_{s,DAC}$ discrete samples. The injection currents are described by the matrix $\mathbf{x}$ with dimensions $N_{ch} \times N_{s,DAC}$. $\mathbf{x}$ is populated such that

$$\mathbf{x}(i,j) = \begin{cases} x_{i,\text{initial}} & \text{for } j < N_{s,\text{switch}} \\ x_{i,\text{final}} & \text{for } j \geq N_{s,\text{switch}}, \end{cases} \tag{8}$$

where $x_{i,\text{initial}}$ is the initial current applied by each DAC channel, $x_{i,\text{final}}$ is the corresponding final current, i indexes the DAC channels from 1 to N_{ch}, $N_{s,\text{switch}}$ is the sample on which the step change occurs, and j indexes the samples from 1 to $N_{s,DAC}$.

The carrier densities in each laser section are described by the matrix $\mathbf{N}$, with dimensions $N_{ch} \times m \cdot N_{s,DAC}$, where m is an integer oversampling factor to be discussed shortly. The initial carrier densities in each section i were considered to be steady state values, and were calculated directly from $x_{i,\text{initial}}$ using Eq. (1). The evolution of $\mathbf{N}$ over time was simulated by numerically

solving Eq. (1) using Euler's method, such that

$$N(i, j_m + 1) = N(i, j_m) + \frac{1}{m} \Delta N(i), \qquad (9)$$

where j_m is the oversampled sample index, counting from 1 to $m \cdot N_{s,\text{DAC}}$, and $\Delta N(i)$ is Eq. (1) evaluated for a discrete simulation step in laser section i. Equation (9) shows that the oversample rate m determines the simulation step size, such that larger values of m result in more accurate approximations of N. By solving Eq. (9) for all values of j_m across all laser sections i, the full progression of N over time was found. Finally, the instantaneous laser frequency was then calculated at every time sample j_m using the frequency domain model given in section 3.1.

3.3. Drive current optimisation

An AI-optimiser was implemented to automatically calculate the multi-sample (overshoot and stability correction) pre-emphasis weights required for fast and reliable wavelength switching [16]. This method uses the laser frequency offset as the error term within a regression optimiser to iteratively update the applied pre-emphasis weights. The measured error at each time sample was integrated into a bin that matches the time window over which $N_{s,\text{DAC}}$ was applied ($1/F_{s,\text{DAC}}$). This method allowed individual pre-emphasis weights to push past positions of local minima where changes to an individual sample weight resulted in an increase in the summed error across the switching event. The final pre-emphasis sample weights can then be saved in a lookup table with the laser to recall the optimised switching response. The effect of pre-emphasis on the laser drive signals is illustrated in Fig. 6. This optimisation method brings the most benefit to many-section lasers with complex tuning maps like the DS-DBR laser considered in this work, but can in principle be used with any semiconductor tuneable laser to automate the calculation of the pre-emphasis sample weights required for fast and stable switching.

To demonstrate the operational principle, the DS-DBR laser model was set to simulate a decreasing current step-change of 50 mA to 1 mA in the rear section. This is shown as the 'initial' case in Fig. 8(a). The evolution of carrier density over time, solved numerically using Eq. (9), is shown in Fig. 8(b). The carrier density is observed to exponentially decay towards the steady-state value. The time-evolution of the laser frequency, determined at each time instance by N, is shown in Fig. 8(c). The laser frequency tracks the carrier density nonlinearly as the cavity adopts discrete frequency modes. In this case, the carrier density and laser frequency take 39.25 ns to reach their steady state values; this is typical of tuneable lasers driven with unoptimised step currents [7,21].

The regression optimiser was then applied to the laser model and the simulation repeated. The number of sample weights to be optimised was set to $K = 4$. The converged value for the optimised driving currents is shown as the 'final' case in Fig. 8(a). It is observed that the optimiser applied pre-emphasis to the first two samples after the switch, overshooting the target steady-state currents. A small corrective counter weight is also present on the third sample, while the fourth sample was left unchanged. The overall effect of the pre-emphasis was to accelerate the rise time and stability of the carrier density within the laser section, as shown within Fig. 8(b). Note that, although the optimiser does not have knowledge of the carrier density, it is still able to push this towards an ideal position. The resulting laser frequency offset is shown in Fig. 8(c), with the error feedback bins labelled. The optimiser has reduced the wavelength switch time of the laser from 39.25 ns to 9.5 ns after 256 updates. The improvement in the overall error (the absolute difference between the ideal frequency step and the final laser switch shown in Fig. 8(c)) is shown in Fig. 8(d). The bin-based feedback method is observed to have successfully pushed passed positions of local minima achieve an improved overall response.

To understand the aggregate performance of the optimiser and the laser response, the simulation procedure shown in Fig. 8(a) to (d) was repeated for 100 wavelength switching events. The data is presented alongside the experimental results in section 4.2.
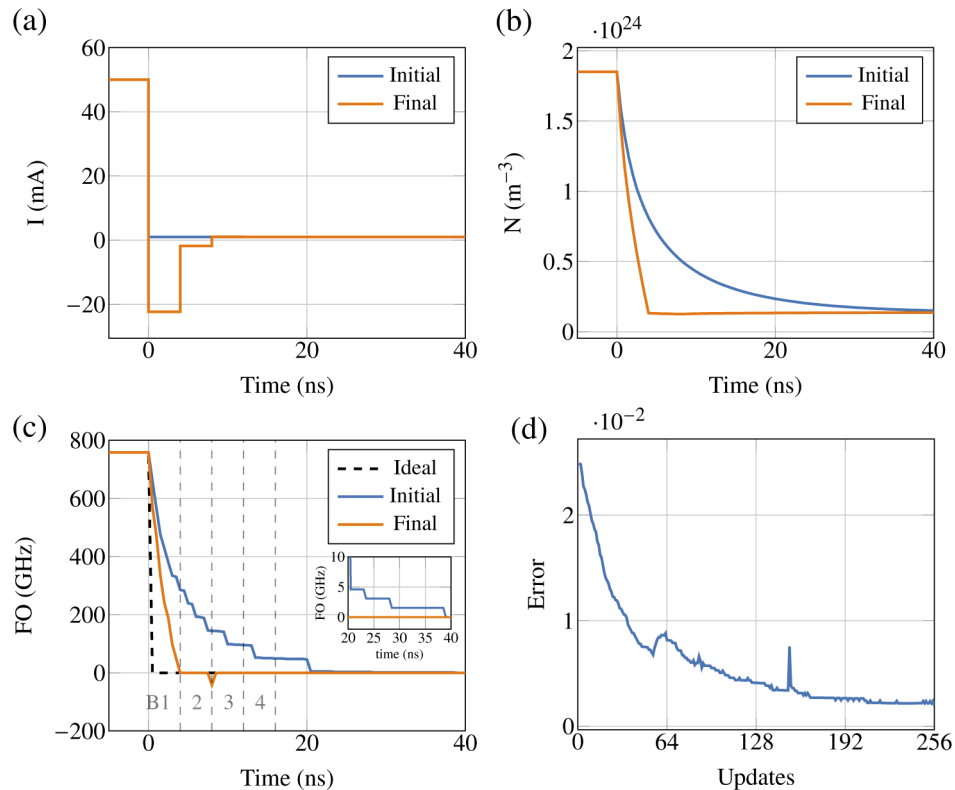
Fig. 8. Time domain simulation of a DS-DBR laser responding to a falling step in rear injection current, before and after regression optimisation. (a) Laser rear section injection current. (b) Carrier density inside the rear section. (c) Instantaneous laser frequency, measured as the frequency offset from the target frequency. B1-4 indicate the bins over which error is integrated for feedback. (d) Absolute error between the laser frequency offset and the ideal frequency step change shown in (c).

4. Experimental demonstration

4.1. Setup

The setup used to demonstrate the hybrid WTS is shown in Fig. 9. A pair of commercial DS-DBR lasers [33] were driven by 250 MS/s 12-bit DACs with 125 MHz bandwidth. A pair of digital-select analogue-route electrical multiplexers were used to reduce the number of DAC channels required to drive each laser from 10 to 4, as described in [16]. The output of each laser was connected to an off-the-shelf SOA, supporting 69 nm of bandwidth with typical characteristics of 7 dB noise figure, 20 dB gain, and 10 dBm saturation power. Each SOA was biased at 45 mA using a low noise ILX Lightwave current source and a bias tee. The biased current was modulated through the RF input of the bias tee using a 12 GS/s DAC with 4.8 GHz bandwidth and ±1 V output, amplified to ±4 V using a 12 GHz RF amplifier with 18 dB gain, saturation power of 22 dBm, and noise figure of 6 dB. This arrangement permitted the SOA to be switched between −30 mA 'off' states and +120 mA 'on' states. The SOA outputs were coupled together and passed to a digital coherent receiver, sampling at 50 GS/s with 22 GHz bandwidth. This was used to measure the frequency offset and provide optimisation feedback to the DS-DBR lasers. A digital sampling oscilloscope (DSO), sampling at 50 GS/s with 30 GHz bandwidth, was used to measure the hybrid WTS output intensity, providing optimisation

feedback to the SOAs. This setup was used so that the DS-DBR laser and SOA optimisation routines could be run concurrently; in practice only the coherent receiver is required, where the intensity information can be recovered through the sum-of-squares of the coherent beat signals.
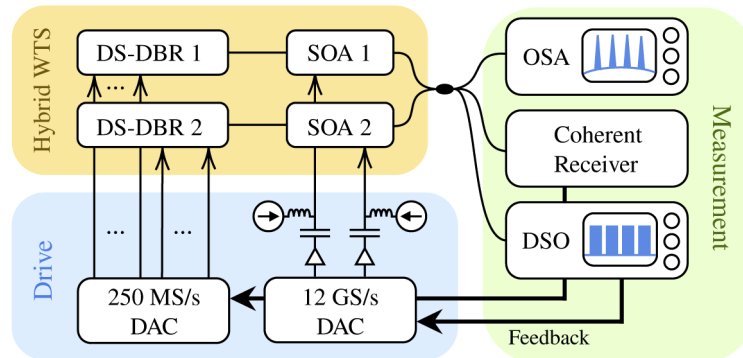

Fig. 9. Experimental setup. Feedback is provided by the coherent receiver and the digital sampling oscilloscope (DSO) to optimise the DS-DBR laser and SOA response.

4.2. Tuneable laser optimisation

The 250 MS/s DACs were set to drive one of the DS-DBR lasers with initially-square 5 MHz waveforms, producing dwell times of 100 ns dwell on each wavelength. Extending on our previous work [16], a fine-frequency correction procedure was applied to the laser phase section prior to pre-emphasis calculation to improve frequency accuracy. The procedure applied a one-tap LMS updater to the phase section drive current in order to minimise the average frequency offset measured after 50 ns. This shifts the target current of the laser phase section, as illustrated in Fig. 6. The regression optimiser demonstrated in section 3.3 was then applied to accelerate the laser wavelength switch and stability time.

The regression optimiser with fine-phase correction was applied to 330 wavelength switching events, testing the extremes of lasing bandwidth (1524.11-1572.48 nm, 6.05 THz), drive current (0-48 mA), and all lasing sections. The frequency offset over time for three example optimised wavelength switching events are shown in Fig. 10(a), measured using the coherent receiver. For each event, the DS-DBR laser enters the receiver bandwidth after ~ 7 ns and settles to within ± 1 GHz over the following tens of nanoseconds. The extremities of the frequency offset at all time positions for all measured wavelength switching events are indicated by the shaded profile. The profile shows that all the wavelength switching events have been brought within ± 5 GHz of their targets before the 20 to 40 ns time window used by the hybrid WTS. The cumulative distribution of the time taken for all wavelength switch events to reach ± 5 GHz of their target wavelengths is given in Fig. 10(b). The optimiser has successfully reduced the worst case tuning time from 78 ns to 14.7 ns, comfortably within the 20 ns target for hybrid WTS. The distribution of frequency offsets 20 and 40 ns after the laser switch is shown in Fig. 10(c). We report a worst-case offset of -4.6 GHz after 20 ns, which settles to within ± 2 GHz after 40 ns. Comparison with our laser simulation suggests that these results are limited by the sample rate and bandwidth of the DACs, and that frequency offset convergence to within ± 1 GHz is possible in under 10 ns if the DAC bandwidth is not limited. Despite this, our achieved results indicate that the WTS is potentially suitable for burst mode coherent detection, as 28 GBd dual-polarisation quadrature phase shift keying is tolerant of frequency offsets up to ± 7 GHz [41]. This is also sufficiently accurate for wavelength locking techniques to maintain frequency accuracy long after the initial wavelength switch event, see Appendix C for details. It should be noted that the recovery of higher-order

quadrature amplitude modulation formats is challenging for DS-DBR lasers due to their high intrinsic linewidth of ~ 1 MHz which is effectively increased during fast switching operation [42]. This can be corrected for (along with residual frequency offset) using feed-forward compensation techniques [43,44].
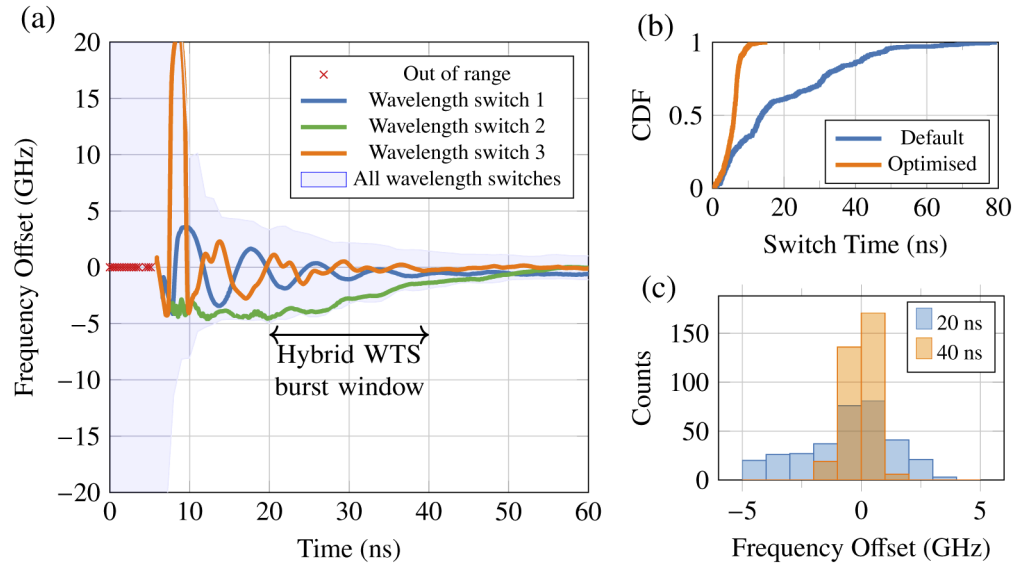

Fig. 10. Experimental performance of the optimised DS-DBR laser. (a) Frequency offset over time for all measured wavelength switch events. (b) Cumulative distribution of the time taken for to reach ±5 GHz of the target wavelength. (c) Histogram of the instantaneous frequency offset after optimisation, 20 and 40 ns after the start of each wavelength switch event.

To examine the pre-emphasis requirements, 100 wavelength switching events were simulated and optimised for a range of increasing and decreasing rear currents using the process described in section 3.3. The simulated and experimental pre-emphasis sample weights for increasing rear current steps are shown in Fig. 11(a) and (b), respectively. The results show good qualitative agreement, applying significant pre-emphasis to the first sample and second sample that decreases as the rear current step increases. The simulated and experimental results for decreasing rear current steps are shown in Fig. 11(c) and (d), respectively. Again, good qualitative agreement is observed, with sample weight decreasing as the magnitude of the step change increases. The experimental results of sample 2 in Fig. 11(d) are larger than simulated, indicating the carrier lifetime of this laser is slightly higher than modelled. However, the regression procedure has successfully found the required values, reflecting the importance of switch-specific optimisation. Figure 11(d) also shows that the optimiser has used sample 3 to help stabilise the experimental switch events for changes in rear current of >30 mA; this has been excluded from the simulated results where values close to zero were calculated. Overall, we note that the decreasing sample weights are lower than their increasing counterparts; this is explained by Eq. (2) which shows that the negative terms help accelerate decreasing current swings, but act against increasing current steps. All of the presented results show a progression but are not explicitly predictable, and in some cases no pre-emphasis is required at all. This is reflective of the complex interaction of the packaging, multiple laser sections and common grounding plane, and highlights the importance of flexible and reliable AI-optimisation.

The size of the lookup table needed to save the pre-emphasis values required for fast laser tuning can be estimated using the above results. We found that increasing current steps required

2 pre-emphasis sample weights for the rear section, 4 for the phase section, and 1 sample for the long-term phase current to ensure frequency accuracy, such that increasing current steps required 7 samples. The decreasing current steps required 3 sample weights for the rear section, 4 for the phase section, and 1 sample for the long-term phase current, totalling 8 samples. Assuming 12-bit values as used in this experiment, each fast tuneable laser requires

$$\left(\frac{121 \cdot 122}{2} \text{ combinations}\right) \times (7 + 8 \text{ samples}) \times (12\text{-bits})/(8\text{-bits per Byte}) = 166 \text{ kB},$$

such that the two-laser Hybrid WTS demonstrated in this paper requires a 332 kB lookup table. This could be stored in the flash memory that is already included with tuneable lasers, where the injection currents required for each wavelength channel are saved. In our tests, the saved pre-emphasis values remained valid for the duration of the experimental testing (~ 4 weeks) so long as the device temperature was maintained by a temperature controller. Pre-emphasis calculation could, therefore, be performed by the manufacturer prior to deployment, and periodically updated within the network (to account for e.g. device aging) using shared optical monitoring equipment.
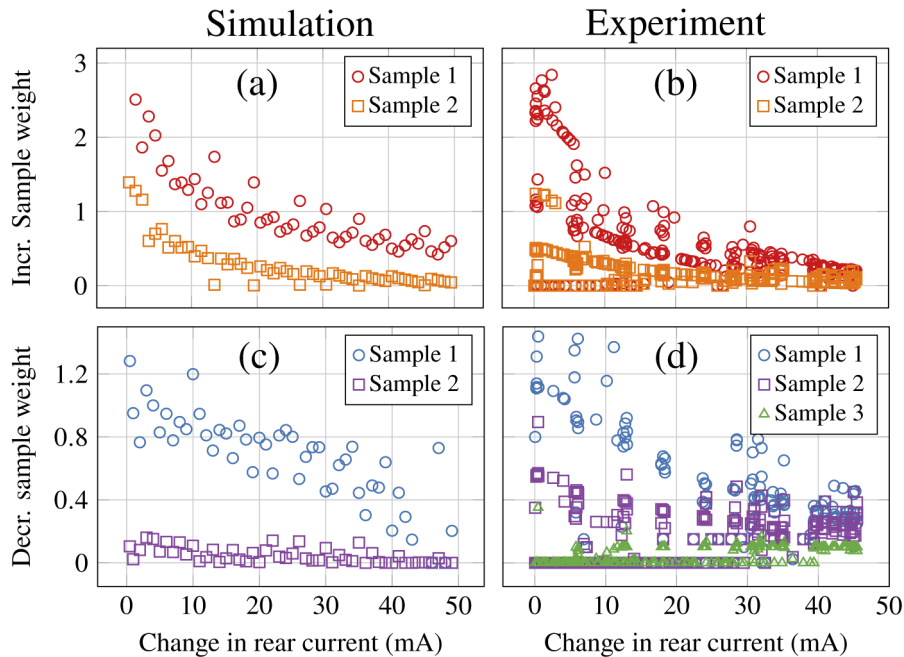

Fig. 11. Simulated (a)/(c) and Experimental (b)/(c) pre-emphasis sample weights for increasing/decreasing changes in rear current. The third tap in (c) was always set close to zero so is excluded for clarity.

4.3. SOA optimisation

The SOA driving signals must also be individually optimised to overcome the intrinsic oscillatory impulse response of the SOA and approach the theoretical rise/fall times of ~100 ps. Even if pre-emphasis overshoot is applied to the SOA driving signal [45], the switch settling time (the time taken for the SOA to settle to within ±5% of its target intensity) can still take several ns. To demonstrate this, SOA 1 in Fig. 9 was driven with a simple step function, shown in Fig. 12(a). The intensity response of the SOA, measured using the DSO, is shown in Fig. 12(b). The 10-90% rise time was measured as 694 ps, but the ±5% settling time was 4.5 ns. This non-ideal SOA gating can delay the sub-ns wavelength switching response needed by the hybrid WTS.

To solve this, we applied particle swarm optimisation (PSO) to achieve automated, sub-ns switch settling times. PSO is a population-based metaheuristic that combines swarm theory with evolutionary programming to optimise continuous nonlinear functions. A detailed simulated and experimental investigation of the application of PSO, and other AI-optimisation methods, for SOA optimisation can be found in [46]. $n = 160$ particles (driving signals) were initialised in an $m = 240$ (number of samples) hyperdimensional search space. Each particle was iteratively 'flown' through the space by evaluating that particle's position with a fitness function f. f is defined as the mean squared error between the optical output generated by a given drive signal, measured on the DSO, and an ideal target set point (SP). Figure 12(a) shows the PSO drive signal, which now features overshoot and stability correction throughout the waveform. The resulting PSO optimised SOA response is shown in Fig. 12(b). The 10-90% rise time was reduced to 454 ps while the ±5% settling time was achieved in just 547 ps; this represents a near order-of-magnitude improvement. We note the PSO signal has introduced some small oscillations to the SOA 'on' state; this is expected behaviour caused by the intrinsic response of the SOA and has been estimated to increase the effective noise figure by 8% [46]. Despite this, the SOA is now able to approach an ideal step response, providing superior optical gating and enabling sub-ns wavelength switching for the hybrid WTS. Furthermore, the PSO routine requires no prior knowledge of the SOA, providing a flexible, automated and practical method for optimising SOA gating. Comparison with simulation results in [46] suggests that the 12 GS/s, 4.8 GHz DAC used to drive the SOAs in this work exceeds the specifications required to achieve idealised sub-ns gating. Future work will investigate the minimum specifications needed to achieve the same result.
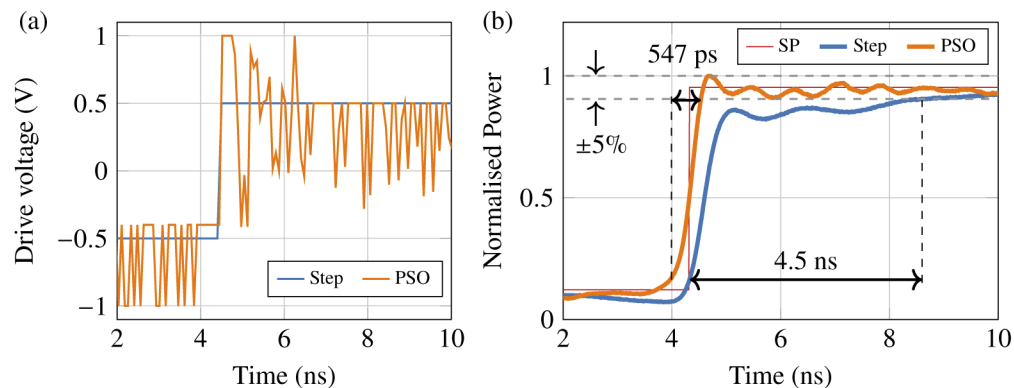

Fig. 12. (a) SOA drive signals using a step function and a waveform optimised using particle swarm optimisation (PSO). (b) Experimental comparison of the SOA response to the step function drive signal, and the PSO drive signal, which was optimised to approach the ideal set point (SP) performance.

4.4. Demonstration of the hybrid WTS

After component optimisation, the operation of the complete hybrid WTS was characterised. The DS-DBR lasers were driven out of phase with 12.5 MHz regression-optimised signals, resulting in dwell times of at least 25 ns on each wavelength. The SOAs were driven by 25 MHz PSO-optimised signals, resulting in 20 ns gates, and aligned to block the first 15 ns and last 5 ns of each laser switching event. Figure 13(a) shows the DSO output for one of the most difficult switching instances, where DS-DBR laser 1 switched from 1572.48 nm to 1524.11 nm, incurring a large rear current swing of 45 mA. Note the rectangular leading edge of DS-DBR 1+ SOA 1, especially compared to SOA 2 which has been left un-optimised for comparison. Packet-to-packet

power variations are due to variations in laser wavelength power, and are visible here because the SOAs were not driven to saturation in order to improve the extinction ratio. For the 22 wavelength channels tested in this work, the hybrid WTS had a mean instantaneous wavelength power of 2.8 dBm, maximum 4.6 dBm (at 1566.31 nm) and minimum 0.8 dBm (at 1572.48 nm). The mean extinction ratio was measured as 20.5 dB (best 22.9 dB, worst 17.1 dB) where the worst case result also occurred at 1572.48 nm. These measurements were limited by the noise floor of our DSO; in principle SOA gating can achieve extinction ratios in excess of 60 dB [47]. To achieve constant output power and high extinction ratios simultaneously, the SOA input power should be limited and the SOAs driven to extinction. In this case the hybrid WTS should approach its ideal output power of 7 dBm for all wavelengths, which is set by the saturation power of the SOAs minus the loss associated with the 2:1 optical coupler.

The wavelength switch time of the hybrid WTS is determined by a combination of the rise time, fall time, and alignment of SOA 1 and SOA 2 while gating. To assess the wavelength switch time of our demonstration system, four wavelength transitions are shown in more detail in Fig. 13(b) with their 90-90% intensity transition times overlaid. All these transitions are below 900 ps, indicating that, thanks to the PSO-optimised SOA gating, our system supports sub-ns wavelength switching. As the SOA operation is independent of the laser tuning, these wavelength switch times are consistent across all wavelength reconfiguration combinations.
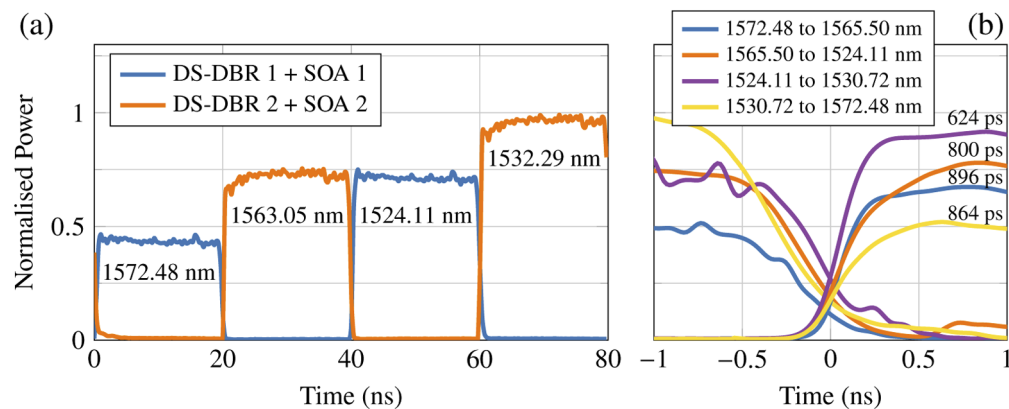

Fig. 13. (a) Output intensity of the hybrid WTS. SOA 1 is driven with a PSO signal. For comparison, SOA 2 has been left unoptimised. (b) Detailed view of the intensity transitions of each SOA. 90-90% wavelength switch times are overlaid.

To measure the underlying laser operation, the output of the coherent receiver for four wavelength slots, both with and without SOA gating, is shown in Fig. 14. The high ±25 GHz frequency deviations of both lasers are observed to be fully suppressed by the SOAs, such that the gated outputs are always within ±5 GHz of their target values. The observed frequency ripples are a result of the low sample rate of the 250 MS/s DAC used in this experiment. The DAC introduces Fourier components to the driving signal creating oscillations in the laser injection currents and the resulting lasing frequency.

Having established the principal characteristics of the hybrid WTS, all of the 22 channels under test were measured, in groups of four, while undergoing sub-ns gated wavelength switching. The optical spectra for all 22 channels are shown in Fig. 15, as measured using the optical spectrum analyser with a resolution of 0.1 nm. The sidemode suppression ratio (SMSR) was measured for each group of four wavelengths, giving a mean result of 38.9 dB (best 41.0 dB, worst 36.8 dB). All the collected spectra were then added together (giving the image in Fig. 15) to measure
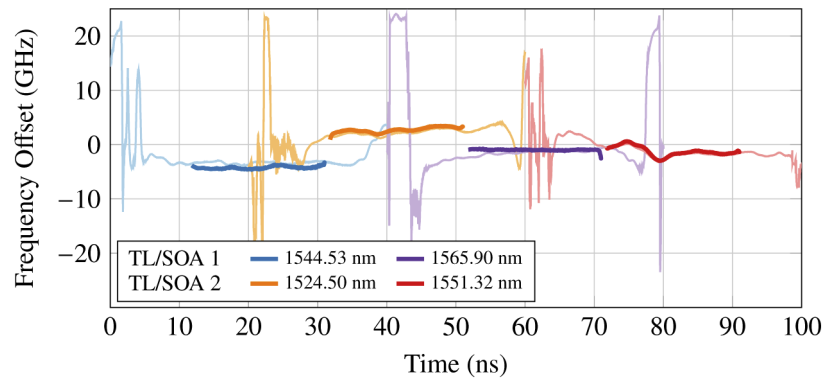
Fig. 14. Instantaneous frequency offset data of the hybrid WTS without (faint) and with (bold) SOA gating. All lasers are within ±5 GHz of their target when gated open.

the worst-case broadband mode suppression ratio of 35.7 dB. This gives an idea of the expected mode suppression when operating many hybrid WTSs on the same optical network.
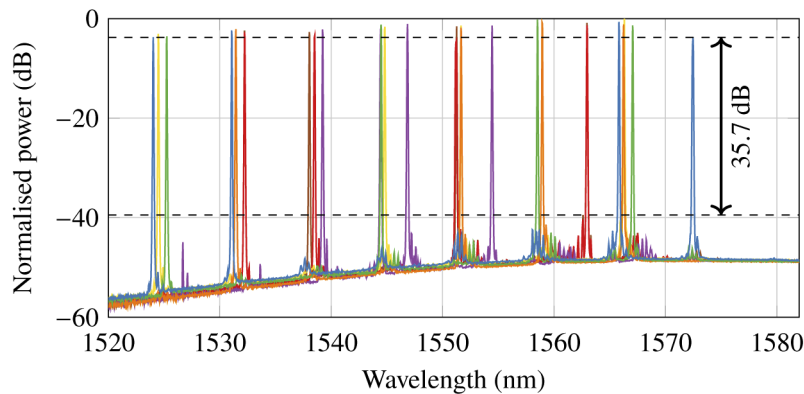

Fig. 15. Spectra of all 22 channels under test undergoing ultra-fast gated switching. The continuously tuneable subsystem is able to support 122×50 GHz channels.

5. Conclusion

We have proposed an AI-optimised, hybrid wavelength tuneable source (WTS) composed of time-interleaved tuneable lasers, each gated by an SOA. The design allows for bandwidth-scalable sub-ns wavelength switching with superior on-chip space efficiency compared to alternative WTS designs. By optimising the driving signals of two tuneable lasers and two SOAs, we have experimentally achieved record performance, demonstrating <900 ps wavelength switching across 6.05 THz (122×50 GHz) of continuously tuneable bandwidth. Our demonstration system supports wavelength dwell times of 20 ns and greater, suitable for fine-granularity, destination-diverse, packet-switched networking. This result shows the hybrid WTS outperforms the channel count of the previously published sub-ns WTS by a factor of ×6.4 [13], and the speed of the previously published wideband WTS by ×10.8 [16]. The AI-methods proposed here are autonomous and require no prior device knowledge, providing a practical method of guaranteeing the performance of every component within every WTS required within a wavelength routed optical burst switching system. The hybrid WTS can scale to hundreds of wavelength channels, accommodating for the high bandwidth per channel and intraconnectivity of hyperscale cloud

data centres, and is the first demonstration of bandwidth-scalable, sub-nanosecond wavelength switching.

Appendix A

Examples of potential line rates and data packet sizes of different types of network traffic were presented in Fig. 3. Here, justifications for these values are given. WTS dwell times are set to match data burst times, which are calculated based on transfer size and line rate. The wavelength switch time is set as 5% of the dwell time.

Ethernet data centre traffic refers to the most common data traffic reported in modern hyperscale data centres. Measurements from both Microsoft and Facebook indicate that hyperscale data centres are dominated by short packets of ∼ 256 B, propagated over 100 Gb/s links [48,49]. Hyperscale networks are expected to remain at 100G for the near future [50]. 256 B at 100G persists for 20.48 ns, requiring a wavelength switch time of 1.02 ns.

Stardust data centre credit considers a proposed data centre fabric, Stardust, in which data packets are chopped into fixed-sized cells, then optimally packed into longer credits for transmission [51]. For strict dwell time requirements, here we consider a Stardust Fabric Adapter that generates a credit every clock cycle, processed using an application specific integrated circuit (ASIC) with 10 Tb/s bandwidth, 400G per port, and (assumed) clock frequency 1 GHz [52]. Referring to section 4.1 of [51], this Stardust Fabric Adapter has a minimum credit size of 1250 B. This credit persists for 25 ns in a 400G link, requiring a wavelength switch time of 1.25 ns.

Inter-GPU transfer references local networks of graphics processing units (GPUs) used for machine learning and high performance computing. Discrete GPUs can be interconnected using Infiniband over PCIe switches for transmission of up to 4 kB packets [53], however, transfer speeds are limited just a few GB/s [54]. To improve on this, NVIDIA has introduced a 16 GPU cluster interconnected using a custom ASIC called the NVSwitch, supporting 300 GB/s of bi-directional bandwidth to each GPU [55]. Generically, this represents a non-blocking 16-node network with uni-directional line rates of 1.2 Tb/s/node. A 4 kB packet persists for 27.3 ns in a 1.2 Tb/s link, requiring a wavelength switch time of 1.37 ns.

Solid-state memory considers a possible future type of network traffic, where dynamic random access memory (DRAM) can be accessed remotely by many devices in a resource disaggregated network [56]. Today, a DRAM technology named high bandwidth memory 2 (HBM2) can be co-integrated with a GPU to provide 180 GB/s of uni-directional traffic with a minimum transfer size of 4 kB [57]. As an example, two HBM2 units could be co-packaged and networked using the hybrid WTS for 8 kB packets at 2.88 Tb/s. These packets persist for 22.76 ns, requiring a wavelength switch time of 1.14 ns.

Appendix B

The parameters used for the frequency and time domain simulations of the DS-DBR laser are given in Table 2. Parameters used for modelling thermal effects are given in Appendix C.

Table 2. DS-DBR laser simulation parameters, sourced from [33,39,40].

Parameter	Value	Parameter	Value
Laser width	2 μm	Front increment, ΔI_{front}	$I_{\text{max,front}}/16$
Laser height	0.4 μm	Front max shift, $\lambda_{\text{max,front}}$	6.8 nm
Linear recombination, A	$1\times10^{-8}\,\text{s}^{-1}$	Front sections, S	8
Bimolecular radiation, B	$1\times10^{-16}\,\text{m}^3\,\text{s}^{-1}$	Front comb lines, k	1
Auger rate, C	$4\times10^{-41}\,\text{m}^6\,\text{s}^{-1}$	Phase length	125 μm
Centre wavelength, λ_{c}	1550 nm	Phase max, $I_{\text{max,phase}}$	5 mA
Rear length	500 μm	Phase increment, ΔI_{phase}	$I_{\text{max,phase}}/16$
Rear current max, $I_{\text{max,rear}}$	50 mA	$\lambda_{\text{max,phase}}$	0.2 nm
Rear increment, ΔI_{rear}	$I_{\text{max,rear}}/256$	Frequency increments	4096
Rear max shift, $\lambda_{\text{max,rear}}$	6.8 nm	DAC sample rate, $F_{s,\text{DAC}}$	250 MHz
Harmonics, h	5	DAC samples, $N_{s,\text{DAC}}$	32
Super Gaussian order, p	8	DAC channels, N_{ch}	10
Rear comb lines, k	7	Oversample rate, m	16
Front length	300 μm	Update scalar, μ	0.2
Front current max, $I_{\text{max,front}}$	5 mA	Simulated bandwidth	8.88 THz

Appendix C

The research presented in sections 3.3 and 4 has focused on how to approach and overcome the fundamental limits of tuneable laser reconfiguration time for application in highly dynamic OCS systems. However, fast tuning using carrier injection can cause thermal heating within the laser, resulting in frequency drifts of several GHz over millisecond timescales. The correction of long term frequency drifts has been the subject of extensive research: proposed solutions include compensatory counter-heating [58–61], gain section current correction [62], and active feedback frequency-lock loops [63–66]. If every hybrid WTS were to be deployed along with a simplified coherent receiver, as proposed in [7,11], the coherent receiver could be used to (1) locally optimise the laser in-situ for few-ns switching, as performed in section 4.2, and (2) provide fine frequency measurements and active feedback correction over long time periods. To demonstrate this, time varying frequency drift was included within the laser model presented in section 3.1 to emulate the presence of thermal heating. Adapting the method presented in [67], the frequency offset over time was modelled as

$$\Delta f(dI, t) = D \exp(-t/T_D) + dI\Big[E\Big(1 - \exp(-t/T_E)\Big) + F\Big(1 - \exp(-t/T_F)\Big)\Big], \qquad (10)$$

where dI is the change in rear section injection current and t is the time after the switch event. The first term of Eq. (10) describes the fast convergence of ΔF over tens of nanoseconds, achieved through pre-emphasis optimisation. The second and third term model the subsequent frequency drift over micro and milliseconds, respectively, and scale with the magnitude of the change in injection current [67]. For this simulation, values of $D = 20$ GHz and $T_D = 14$ ns were chosen to match the worst-case fast switching experimental results reported in section 4.2. Values of $E = 0.148$ GHz/mA, $T_E = 5.13$ μs, $F = 0.212$ GHz/mA and $T_F = 0.204$ ms were selected to match experimental measurements of DS-DBR laser thermal frequency drift reported in [62].

The modelled frequency offset for $dI = 50$ mA is shown in Fig. 16(a). The laser settles on the target frequency over the first tens of nanoseconds via the plasma effect, then increases by up to ∼ 18 GHz over the next 10 ms via thermal effects. To correct for the drift, a simple active locking loop can applied to the laser phase section. As a proof of principle, a normalised least means squared (LMS) filter was implemented that measured the frequency offset and update

the injection current applied to the laser phase section. Figure 16(b) shows the phase section current with and without locking enabled, for a current update frequency of 5 MHz and a loop bandwidth of 10 MHz. With locking enabled, the injection current successfully tracked the frequency drift such that the laser frequency offset was suppressed, as shown in Fig. 16(a). Stable locking required a loop bandwidth at least equal to the phase current update frequency; we found an update rate of 5 MHz was appropriate for maintaining frequency offsets below 0.5 GHz.
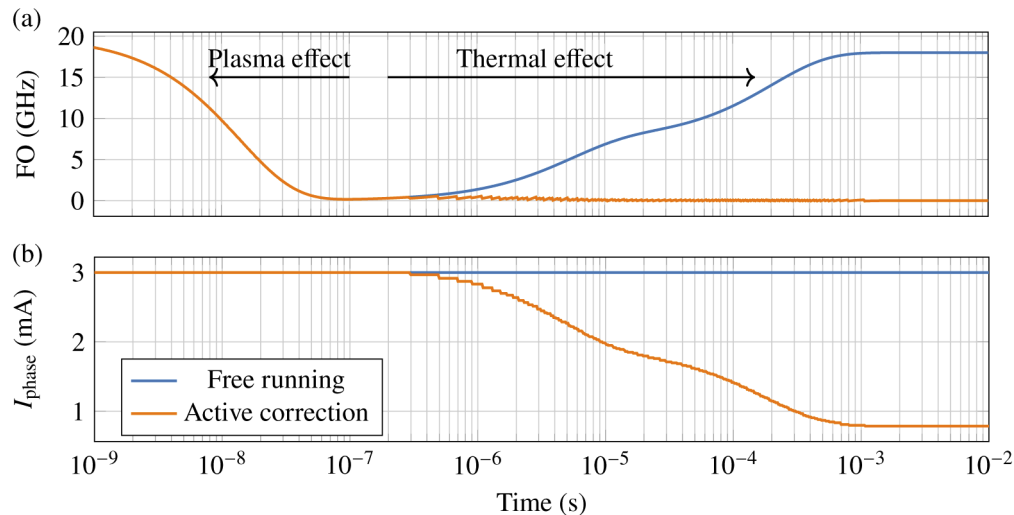

Fig. 16. Simulation of frequency drift in a DS-DBR laser, with and without correction. (a) Frequency offset (FO) over time. (b) Current applied to the laser phase section.

An advantage of the active locking solution presented here is that it is agnostic to the origin of thermal drift. It could, therefore, be used not only to correct for laser self-heating, but also cross-heating introduced by co-integrated lasers and SOAs [68]. Generally, however, the presence of a coherent receiver at every node is not a requirement for the hybrid WTS as the pre-emphasis values needed for few-ns reconfiguration are pre-calculated. This could be performed by the manufacturer prior to deployment, and periodically updated within the network using shared optical monitoring equipment. In this case, the correction of thermal frequency drift is better achieved using an existing wavelength locking technologies based on etalon feedback [63–66].

References

1. A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis, "The lockdown effect: Implications of the covid-19 pandemic on internet traffic," in *Proceedings of the ACM Internet Measurement Conference*, (Association for Computing Machinery, New York, NY, USA, 2020), IMC'20, p. 1–18.
2. Microsoft, "Update #2 on Microsoft cloud services continuity," [Online] https://azure.microsoft.com/en-us/blog/update-2-on-microsoft-cloud-services-continuity/ (2020).
3. M. Duser and P. Bayvel, "Analysis of a dynamically wavelength-routed optical burst switched network architecture," J. Lightwave Technol. **20**(4), 574–585 (2002).
4. Y.-K. Yeo, Z. Xu, D. Wang, J. Liu, Y. Wang, and T.-H. Cheng, "High-speed optical switch fabrics with large port count," Opt. Express **17**(13), 10990–10997 (2009).
5. K. Xi, Y.-H. Kao, and H. J. Chao, *A Petabit Bufferless Optical Switch for Data Center Networks* (Springer New York, New York, NY, 2013), pp. 135–154.

6. R. Proietti, Y. Yin, R. Yu, C. J. Nitta, V. Akella, C. Mineo, and S. J. B. Yoo, "Scalable Optical Interconnect Architecture Using AWGR-Based TONAK LION Switch With Limited Number of Wavelengths," J. Lightwave Technol. **31**(24), 4087–4097 (2013).
7. A. C. Funnell, K. Shi, P. Costa, P. Watts, H. Ballani, and B. C. Thomsen, "Hybrid wavelength switched-TDMA high port count all-optical data centre switch," J. Lightwave Technol. **35**(20), 4438–4444 (2017).
8. K. Sato, "Realization and application of large-scale fast optical circuit switch for data center networking," J. Lightwave Technol. **36**(7), 1411–1419 (2018).
9. M. Xu, C. Liu, and S. Subramaniam, "PODCA: A passive optical data center network architecture," J. Opt. Commun. Netw. **10**(4), 409–420 (2018).
10. R. Proietti, X. Xiao, K. Zhang, G. Liu, H. Lu, P. Fotouhi, J. Messig, and S. J. B. Yoo, "Experimental demonstration of a 64-port wavelength routing thin-CLOS system for data center switching architectures," J. Opt. Commun. Netw. **10**(7), B49–B57 (2018).
11. J. L. Benjamin, T. Gerard, D. Lavery, P. Bayvel, and G. Zervas, "PULSE: Optical circuit switched data center architecture operating at nanosecond timescales," J. Lightwave Technol. **38**(18), 4906–4921 (2020).
12. H. Ballani, P. Costa, R. Behrendt, D. Cletheroe, I. Haller, K. Jozwik, F. Karinou, S. Lange, K. Shi, B. Thomsen, and H. Williams, "Sirius: A flat datacenter network with nanosecond optical switching," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, (Association for Computing Machinery, New York, NY, USA, 2020), SIGCOMM'20, p. 782–797.
13. K. Shi, S. Lange, I. Haller, D. Cletheroe, R. Behrendt, B. Thomsen, F. Karinou, K. Jozwik, P. Costa, and H. Ballani, "System demonstration of nanosecond wavelength switching with burst-mode PAM4 transceiver," in *European Conference on Optical Communication*, (2019).
14. K. Clark, D. Cletheroe, T. Gerard, I. Haller, K. Jozwik, K. Shi, B. Thomsen, H. Williams, G. Zervas, H. Ballani, P. Bayvel, P. Costa, and Z. Liu, "Synchronous subnanosecond clock and data recovery for optically switched data centres using clock phase caching," Nat. Electron. **3**(7), 426–433 (2020).
15. J. E. Simsarian, M. C. Larson, H. E. Garrett, H. Xu, and T. A. Strand, "Less than 5-ns wavelength switching with an SG-DBR laser," IEEE Photonics Technol. Lett. **18**(4), 565–567 (2006).
16. T. Gerard, H. Dzieciol, J. Benjamin, K. Clark, H. Williams, B. Thomsen, D. Lavery, and P. Bayvel, "Packet Timescale Wavelength Switching Enabled by Regression Optimisation," IEEE Photonics Technol. Lett. **32**(8), 477–480 (2020).
17. Y. Ueda, T. Shindo, S. Kanazawa, N. Fujiwara, and M. Ishikawa, "Electro-optically tunable laser with ultra-low tuning power dissipation and nanosecond-order wavelength switching for coherent networks," Optica **7**(8), 1003–1006 (2020).
18. S. Lange, A. Raja, K. Shi, M. Karpov, R. Behrendt, D. Cletheroe, I. Haller, F. Karinou, X. Fu, J. Liu, A. Lukashchuk, B. Thomsen, K. Jozwik, P. Costa, T. J. Kippenberg, and H. Ballani, "Sub-nanosecond optical switching using chip-based soliton microcombs," in *Optical Fiber Communication Conference (OFC)*, (2020).
19. Y. Tada, N. Shibata, O. Ishida, K. Nosu, and H. Ishii, "Duplex transmitter configuration for bit-error-free optical FDM cross-connect system," IEEE Photonics Technol. Lett. **4**(9), 1051–1053 (1992).
20. H. Okamoto, H. Yasaka, K. Sato, Y. Yoshikuni, K. Oe, K. Kishi, Y. Kondo, and M. Yamamoto, "A wavelength-tunable duplex integrated light source for fast wavelength switching," J. Lightwave Technol. **14**(6), 1033–1041 (1996).
21. N. Ryan, M. Todd, T. Farrell, A. Lavin, P.-J. Rigole, B. Corbett, B. Roycroft, and J.-P. Engelstaedter, "A 10Gbps optical burst switching network incorporating ultra-fast (5ns) wavelength switched tunable laser sources," in *International Conference on Space Optics (ICSO)*, (2008), pp. 289–297.
22. S. Matsuo, S.-H. Jeong, T. Segawa, H. Okamoto, Y. Kawaguchi, Y. Kondo, H. Suzuki, and Y. Yoshikuni, "Stable and fast wavelength switching in digitally tunable laser using chirped ladder filter," IEEE J. Sel. Top. Quantum Electron. **13**(5), 1122–1128 (2007).
23. J. P. Engelstaedter, B. Roycroft, F. H. Peters, and B. Corbett, "Fast wavelength switching in interleaved rear reflector laser," in *2010 22nd International Conference on Indium Phosphide and Related Materials (IPRM)*, (2010), pp. 1–3.
24. A. J. Ward, N. D. Whitbread, P. J. Williams, A. K. Wood, and M. J. Wale, "Extending the tuning range of DS-DBR lasers," in *2008 IEEE 21st International Semiconductor Laser Conference*, (2008), pp. 147–148.
25. A. Maese-Novo, R. Halir, S. Romero-García, D. Pérez-Galacho, L. Zavargo-Peche, A. O.-M. nux, I. Molina-Fernández, J. G. Wangüemert-Pérez, and P. Cheben, "Wavelength independent multimode interference coupler," Opt. Express **21**(6), 7033–7040 (2013).
26. A. Raja, A. Voloshin, H. Guo, S. Agafonova, J. Liu, A. Gorodnitskiy, M. Karpov, N. Pavlov, E. Lucas, R. Galiev, A. Shitikov, J. Jost, M. Gorodetsky, and T. Kippenberg, "Electrically pumped photonic integrated soliton microcomb," Nat. Commun. **10**(1), 680 (2019).
27. K. Morito, S. Tanaka, S. Tomabechi, and A. Kuramata, "A broad-band MQW semiconductor optical amplifier with high saturation output power and low noise figure," IEEE Photonics Technol. Lett. **17**(5), 974–976 (2005).
28. N. Keil, Z. Zhang, C. Zawadzki, C. Wagner, A. Scheibe, H. Ehlers, D. Ristau, J. Wang, W. Brinker, and N. Grote, "Ultra low-loss 1 x 2 multiplexer using thin-film filters on polymer integration platform," Electron. Lett. **45**(23), 1167–1168 (2009).
29. G. Roelkens, A. Abassi, P. Cardile, U. Dave, A. De Groote, Y. De Koninck, S. Dhoore, X. Fu, A. Gassenq, N. Hattasan, Q. Huang, S. Kumari, S. Keyvaninia, B. Kuyken, L. Li, P. Mechet, M. Muneeb, D. Sanchez, H. Shao, T. Spuesens, A. Z. Subramanian, S. Uvin, M. Tassaert, K. Van Gasse, J. Verbist, R. Wang, Z. Wang, J. Zhang, J.

Van Campenhout, X. Yin, J. Bauwelinck, G. Morthier, R. Baets, and D. Van Thourhout, "Iii-v-on-silicon photonic devices for optical communication and sensing," Photonics **2**(3), 969–1004 (2015).

30. T. Komljenovic, D. Huang, P. Pintus, M. A. Tran, M. L. Davenport, and J. E. Bowers, "Photonic integrated circuits using heterogeneous integration on silicon," Proc. IEEE **106**(12), 2246–2257 (2018).

31. J. M. Ramirez, H. Elfaiki, T. Verolet, C. Besancon, A. Gallet, D. Néel, K. Hassan, S. Olivier, C. Jany, S. Malhouitre, K. Gradkowski, P. E. Morrissey, P. O'Brien, C. Caillaud, N. Vaissiére, J. Decobert, S. Lei, R. Enright, A. Shen, and M. Achouche, "Iii-v-on-silicon integration: From hybrid devices to heterogeneous photonic integrated circuits," IEEE J. Sel. Top. Quantum Electron. **26**(2), 1–13 (2020).

32. C. Chase, Y. Rao, W. Hofmann, and C. J. Chang-Hasnain, "1550 nm high contrast grating VCSEL," Opt. Express **18**(15), 15461–15466 (2010).

33. A. J. Ward, D. J. Robbins, G. Busico, E. Barton, L. Ponnampalam, J. P. Duck, N. D. Whitbread, P. J. Williams, D. C. J. Reid, A. C. Carter, and M. J. Wale, "Widely tunable DS-DBR laser with monolithically integrated soa: design and performance," IEEE J. Sel. Top. Quantum Electron. **11**(1), 149–156 (2005).

34. T. Fujisawa, T. Kitoh, K. Watanabe, K. Suzuki, M. Kotoku, and H. Takahashi, "Wide-bandwidth, low-waveguide-width-sensitivity InP-based multimode interference coupler designed by wavefront matching method," IEICE Electron. Express **8**(24), 2100–2105 (2011).

35. F. Kish, V. Lal, P. Evans, S. W. Corzine, M. Ziari, T. Butrie, M. Reffle, H. Tsai, A. Dentai, J. Pleumeekers, M. Missey, M. Fisher, S. Murthy, R. Salvatore, P. Samra, S. Demars, N. Kim, A. James, A. Hosseini, P. Studenkov, M. Lauermann, R. Going, M. Lu, J. Zhang, J. Tang, J. Bostak, T. Vallaitis, M. Kuntz, D. Pavinski, A. Karanicolas, B. Behnia, D. Engel, O. Khayam, N. Modi, M. R. Chitgarha, P. Mertz, W. Ko, R. Maher, J. Osenbach, J. T. Rahn, H. Sun, K. Wu, M. Mitchell, and D. Welch, "System-on-chip photonic integrated circuits," IEEE J. Sel. Top. Quantum Electron. **24**(1), 1–20 (2018).

36. P. J. Rigole, M. Shell, S. Nilsson, D. J. Blumenthal, and E. Berglind, "Fast wavelength switching in a widely tunable GCSR laser using a pulse pre-distortion technique," in *Proceedings of Optical Fiber Communication Conference)*, (1997), pp. 231–232.

37. J. E. Carroll, J. Whiteaway, and D. Plumb, *Distributed feedback semiconductor lasers*, vol. 10 (IET, 1998).

38. S. Adachi, *Physical Properties of III-V Semiconductor Compounds: InP, InAs, GaAs, GaP, InGaAs, InGaAsP* (John Wiley & Sons, 1992), chap. 8, pp. 1–27.

39. A. J. Ward, D. J. Robbins, D. C. J. Reid, N. D. Whitbread, G. Busico, P. J. Williams, J. P. Duck, D. Childs, and A. C. Carter, "Realization of phase grating comb reflectors and their application to widely tunable DBR lasers," IEEE Photonics Technol. Lett. **16**(11), 2427–2429 (2004).

40. A. J. Ward, G. Busico, N. D. Whitbread, L. Ponnampalam, J. P. Duck, and D. J. Robbins, "Linewidth in widely tunable digital supermode distributed bragg reflector lasers: Comparison between theory and measurement," IEEE J. Quantum Electron. **42**(11), 1122–1127 (2006).

41. J. E. Simsarian, J. Gripp, S. Chandrasekhar, and P. Mitchell, "Fast-tuning coherent burst-mode receiver for metropolitan networks," IEEE Photonics Technol. Lett. **26**(8), 813–816 (2014).

42. R. Maher and B. Thomsen, "Dynamic linewidth measurement technique using digital intradyne coherent receivers," Opt. Express **19**(26), B313–B322 (2011).

43. I. Kim, O. Vassilieva, P. Palacharla, and M. Sekiya, "Feed-forward optical carrier phase noise compensation in m-qam transmitter," in *2012 17th Opto-Electronics and Communications Conference*, (2012), pp. 148–149.

44. R. Maher, D. Lavery, M. Paskov, P. Bayvel, S. J. Savory, and B. C. Thomsen, "Fast wavelength switching 6 gbd dual polarization 16qam digital coherent burst mode receiver," IEEE Photonics Technol. Lett. **26**(3), 297–300 (2014).

45. C. M. Gallep and E. Conforti, "Reduction of semiconductor optical amplifier switching times by preimpulse step-injected current technique," IEEE Photonics Technol. Lett. **14**(7), 902–904 (2002).

46. C. W. F. Parsonson, Z. Shabka, W. K. Chlupka, B. Goh, and G. Zervas, "Optimal control of SOAs with artificial intelligence for sub-nanosecond optical switching," J. Lightwave Technol. **38**(20), 5563–5573 (2020).

47. K. Taguchi, K. Asaka, S. Kimura, and N. Yoshimoto, "High output power and burst extinction ratio λ-tunable ONU transmitter using burst-mode booster SOA for WDM/TDM-PON," J. Opt. Commun. Netw. **7**(1), 1–7 (2015).

48. Q. Zhang, V. Liu, H. Zeng, and A. Krishnamurthy, "High-resolution measurement of data center microbursts," in *Proceedings of the 2017 Internet Measurement Conference*, (Association for Computing Machinery, New York, NY, USA, 2017), IMC'17, p. 78–85.

49. K. Clark, H. Ballani, P. Bayvel, D. Cletheroe, T. Gerard, I. Haller, K. Jozwik, K. Shi, B. Thomsen, P. Watts, H. Williams, G. Zervas, P. Costa, and Z. Liu, "Sub-nanosecond clock and data recovery in an optically-switched data centre network," in *European Conference on Optical Communication (ECOC)*, (2018).

50. A. E. A. Andreyev and X. Wang, "Reinventing facebook's data center network," [Online] https://engineering.fb.com/data-center-engineering/f16-minipack/ (2019).

51. N. Zilberman, G. Bracha, and G. Schzukin, "Stardust: Divide and conquer in the data center network," in *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, (USENIX Association, Boston, MA, 2019), pp. 141–160.

52. Broadcom, "BCM88690 StrataDNX 10 Tb/s scalable switching device, product brief," [Online] https://docs.broadcom.com/doc/88690-PB100 (2018).

53. Mellandox, "Mellanox OFED GPUDirect RDMA, GPU-GPU Communication Acceleration Software, product brief," [Online] https://www.mellanox.com/related-docs/prod_software/PB_GPUDirect_RDMA.PDF (2018).

54. L. Oden, H. Fröning, and F. Pfreundt, "Infiniband-Verbs on GPU: A Case Study of Controlling an Infiniband Network Device from the GPU," in *2014 IEEE International Parallel Distributed Processing Symposium Workshops*, (2014), pp. 976–983.

55. NVIDIA, "NVIDIA NVSWITCH, The World's Highest-Bandwidth On-Node Switch," Tech. rep., NVIDIA (2018).

56. G. Zervas, H. Yuan, A. Saljoghei, Q. Chen, and V. Mishra, "Optically disaggregated data centers with minimal remote memory latency: Technologies, architectures, and resource allocation [invited]," J. Opt. Commun. Netw. **10**(2), A270–A285 (2018).

57. NVIDIA, "NVIDIA Tesla P100, The Most Advanced Datacenter Accelerator Ever Built, white paper," Tech. rep., NVIDIA (2016).

58. N. Fujiwara, H. Ishii, H. Okamoto, Y. Kawaguchi, Y. Kondo, and H. Oohashi, "Suppression of thermal wavelength drift in super-structure grating distributed bragg reflector (SSG-DBR) laser with thermal drift compensator," IEEE J. Sel. Top. Quantum Electron. **13**(5), 1164–1169 (2007).

59. N. Nunoya, H. Ishii, Y. Kawaguchi, R. Iga, T. Sato, N. Fujiwara, and H. Oohashi, "Tunable distributed amplification (TDA-) DFB laser with asymmetric structure," IEEE J. Sel. Top. Quantum Electron. **17**(6), 1505–1512 (2011).

60. T. Kanai, N. Nunoya, T. Yamanaka, R. Iga, M. Shimokozono, and H. Ishii, "High-accuracy, sub-μs wavelength switching with thermal drift suppression in tunable distributed amplification TDA-DFB laser array," in *2013 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC)*, (2013), pp. 1–3.

61. M. Gay, L. Bramerie, C. Peucheret, J. O'Carroll, R. Phelan, M. Gleeson, D. Byrne, B. Kelly, M. Nawrocka, and F. Saliou, "Frequency drift reduction in a four-laser array for TWDM PON applications," IEEE Photonics Technol. Lett. **30**(14), 1345–1348 (2018).

62. S. Yoo, J. K. Lee, and K. Kim, "Suppression of thermal wavelength drift in widely tunable DS-DBR laser for fast channel-to-channel switching," Opt. Express **25**(24), 30406–30417 (2017).

63. Analog Devices, "Tunable laser reference design for designers with the ADuC832/ADN8830/ADN2830, AN-655 application note," [Online] https://www.analog.com/media/en/reference-design-documentation/reference-designs/5763882AN655_b.pdf (2004).

64. J. E. Simsarian and L. Zhang, "Wavelength locking a fast-switching tunable laser," IEEE Photonics Technol. Lett. **16**(7), 1745–1747 (2004).

65. A. Bianciotto, B. J. Puttnam, B. Thomsen, and P. Bayvel, "Optimization of wavelength-locking loops for fast tunable laser stabilization in dynamic optical networks," J. Lightwave Technol. **27**(12), 2117–2124 (2009).

66. B. Puttnam, B. C. Thomsen, R. Muckstein, A. Bianciotto, and P. Bayvel, "Nanosecond tuning of a DS-DBR laser for dynamic optical networks," in *CLEO/Europe - EQEC 2009 - European Conference on Lasers and Electro-Optics and the European Quantum Electronics Conference*, (2009), p. 1.

67. A. A. Saavedra, R. Passy, and J. P. von der Weid, "Thermal drift in wavelength-switching DFB and DBR lasers," Electron. Lett. **33**(9), 780–781 (1997).

68. M.-C. Lo, Z. Zhou, S. Pan, G. Carpintero, and Z. Liu, "Characterisation of thermal crosstalk-induced wavelength shift in monolithic InP dual DFB lasers PIC," in *Integrated Photonics Platforms: Fundamental Research, Manufacturing and Applications*, vol. 11364 R. G. Baets, P. O'Brien, and L. Vivien, eds., International Society for Optics and Photonics (SPIE, 2020), pp. 190–195.