VCSEL-Based Light Sources—Scalability Challenges for VCSEL-Based Multi-100-Gb/s Systems
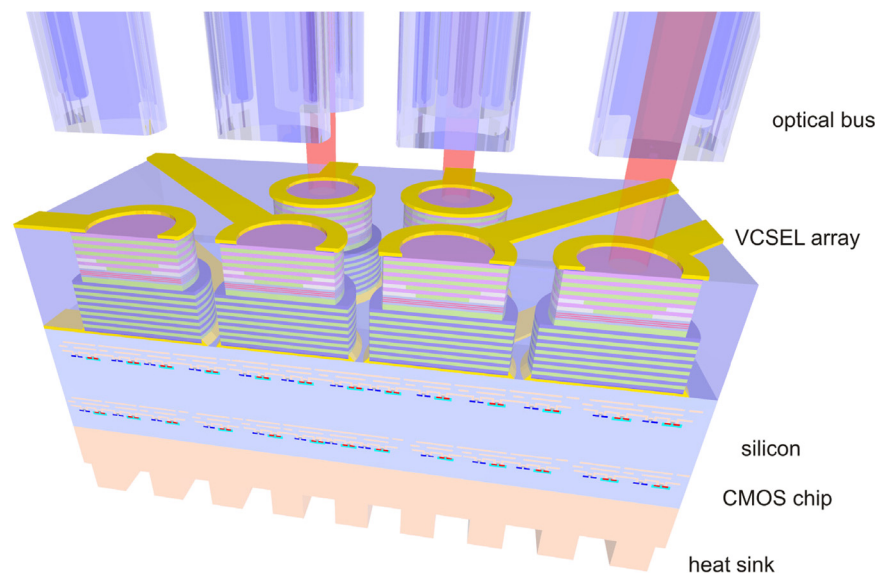
Werner Hofmann, Member, IEEE
Dieter Bimberg, Fellow, IEEE

VCSEL-Based Light Sources—Scalability Challenges for VCSEL-Based Multi-100-Gb/s Systems

Werner Hofmann,[1] *Member, IEEE*, **and Dieter Bimberg,**[1,2] *Fellow, IEEE*

[1]Institute of Solid State Physics and Center of Nanophotonics, Technical University of Berlin, 10623 Berlin, Germany
[2]King Abdulaziz University, Jeddah 22254, Saudi Arabia

Abstract: Future high-performance computers require optical interconnects with aggregated Exa-Byte/s data transport. Densely packed arrays of vertical-cavity surface-emitting lasers (VCSELs) might present the only feasible technical solution. The high-speed properties of semiconductor lasers, however, are strongly affected by their operating temperature. Thermal crosstalk becomes dominant when densely packed arrays of high-speed VCSELs are required. In this paper, we derive the maximum bandwidth of future VCSEL-based optical interconnects from the influence of device heating occurring in high-speed VCSEL arrays. Furthermore, we estimate the scalability of this technology and address the challenges. From our calculations we obtain, that VCSEL arrays are scalable from a bandwidth density of 100 Gbps/mm^2 with today's devices up to a technological limit of 15 Tbps/mm^2.

Index Terms: Vertical-cavity surface-emitting laser (VCSEL), optical interconnects, energy efficiency, high-performance computing, VCSEL arrays, scalability.

1. Bandwidth Demands in Systems

The exponential growth of computer performance and Internet traffic is taken as an indicator of development and progress of our communication society. The performance increase of computer systems is scaling with a factor 1000 for every ten years, and the global Internet protocol (IP) traffic grows even faster. These numbers are interdependent, as the Internet traffic is routed through larger and larger data centers (DCs). Currently, these DCs are built out of 100.000 s of servers operating in parallel. With the increase in parallelism, we are facing a transition from Moore's law, where the performance increase is scaling with the number of transistors that can be placed onto a die, to Amdahl's law of balanced and ideally parallel systems. Amdahl's law states that, for each floating point operation, one byte has to be transmitted through the system's interconnect. According to IBM, the performance of a single server is going to reach 10 Teraflop by 2020, leading to an aggregated Exa-Byte/s data transport in DCs and high-performance computers [1]. The trend of doubling the bandwidth requirements every 18 months is also a major challenge for network operators and system vendors, which are fighting against an exponential increase in technical effort and cost. Furthermore, without technological innovation, this also means that the power consumed by DCs and the grid continues to grow exponentially. "Green IT" has therefore become a crucial subject. According the International Technology Roadmap for Semiconductors (ITRS), lasers for future optical interconnects should be highly energy efficient. In 2015, energy-efficient high-speed

lasers operating at 100 mW/Tbps (100 fJ/bit) will be required [2], [3]. Each new generation of supercomputers is required to deliver a vast boost in computational speed but has to keep cost and energy consumption at a moderate level. This is no longer possible with copper-based interconnects. Already in 2008 IBM's Peta-Flop Supercomputer incorporated 48 000 optical links. Following the roadmap dictated by the market, 2020's Exa-Flop mainframes will need more than 320 million optical interconnects running at 5x the speed at 1/50 of the energy consumption and 1/400 of the price compared to the optical links used in 2008 [4]. The largest part of the energy consumed by these computers is then consumed by the interconnects. To give another example, today's most powerful high-performance computer provides only 0.01 Exa-Flop but consumes already 2/3 of the power budget (20 MW) reserved for future Exa-Flop mainframes.

Vertical-cavity surface-emitting lasers (VCSELs) are an answer to the question that optical technology will present a workhorse of future interconnects. VCSELs are capable to deliver highest modulation speeds beyond 40 Gbps [5], [6]. At the same time, they consume small amounts of power and can be mass fabricated at very low cost. While directly modulated edge-emitting laser diodes are a competing device technology with steady improvements in speed and efficiency, VCSELs are indispensable for 2-D arrays.

In order to be applied in short optical interconnects, VCSELs have to deliver high serial bandwidths with small footprints, allowing dense packaging and uncooled operation. Moreover, the scalability of a certain technology is also limited by the amount of links that can be connected. System designers from Google have stated in 2011 that 40 Gbps would be the bandwidth desired for their next generation of DCs [7]. IBM, on the other hand, envisions 25 Gb/s interconnect speed for their 2020's Exa-Flop mainframes [1].

For system scalability, very dense packaging of the optical chips is a necessity. The VCSEL footprint, being one order of magnitude smaller than edge emitters, is also in favor of this technology. Additionally, only surface-emitting devices can be easily fabricated in 2-D arrays. To ensure a compact hybrid package, bottom-emitting devices with the electrical fan-out on the one side, and the optical on the other have been proposed by IBM's TERABUS project [8]. Last but not the least, VCSELs have a much smaller energy consumption than other types of laser diodes and would, just for that reason, be the light source of choice in future applications as far as no excessive output power is in demand [9].

As the optical interconnect market requires a volume of multibillions of devices to be available, the mature GaAs-based technology seems to present the only basis for the years to come. For integrated optoelectronics on silicon, VCSELs at wavelengths beyond 1300 nm based on InP might be part of a later step of evolution [10]. Up to 1300 nm, quantum-dot VCSELs still rely on GaAs technology [11]. Further shrinking of footprint and energy consumption, together with higher bandwidth, are requirements favoring novel technologies like nanolasers or metal-clad surface-emitting lasers one day [12], [13].

In this paper, we evaluate the scalability limits of optical interconnects based on 2-D ultra high-speed VCSELs arrays. Thermal crowding turns out to be the limiting factor. As a result, energy-efficient VCSELs are advantageous for scaling up.

2. Scaling VCSEL-Based Optical Interconnects

Having the very challenging requirements of future optical interconnects with millions of channels in mind, we have to ask about the technological limits. Following IBM's visions for 2020, 80 Tb/s input/output (IO) bandwidth per server will be required [1]. Since the whole server will be on a chip, the interconnect should ideally not use more space than the server die. This translates into optoelectronic devices with a small footprint integrated into a 3-D stack on CMOS. One side of the chip will be occupied by the inevitable heat sink; the other side will be needed for the optical IO fan-out. This layout is schematically depicted in Fig. 1. As the VCSELs are sitting on top of their driver chip plus the server die, thermal stability is a crucial figure of merit. Furthermore, as the heat sinking of the optoelectronics goes via the silicon die and also affects the server chip, the heat generated by the devices should be as little as possible. Energy efficiency of the VCSEL devices is therefore a basic prerequisite.
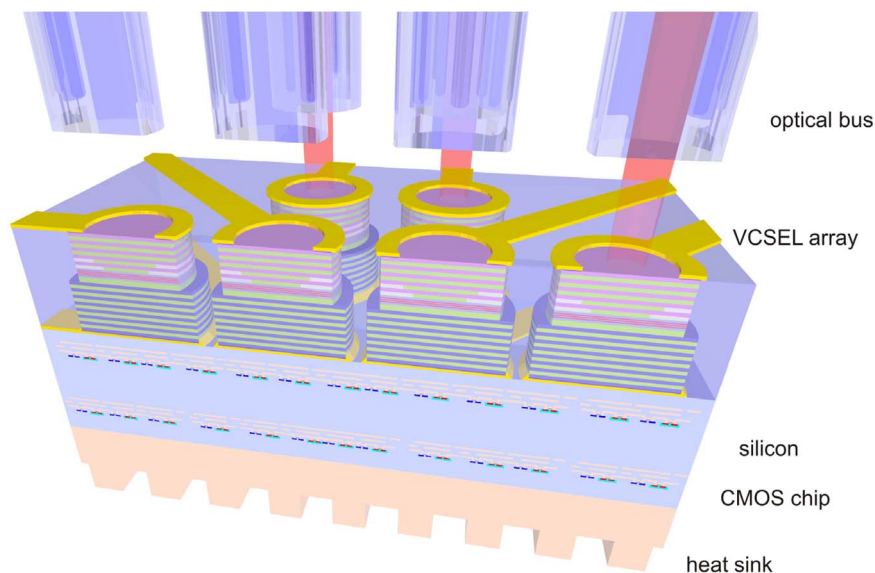
Fig. 1. Envisioned schematic layout of a future VCSEL-based optical interconnect based on 3-D stacking of optoelectronics on CMOS.

As these considerations seem trivial, we would like to derive a more quantitative statement addressing the questions about the scalability limit of VCSEL-based optical interconnects. Of particular interest are the tradeoffs that we have to face when scaling up. Additionally, we will identify characteristics of good device designs for scalable VCSELs within an array.

To start with, we can state that the bandwidth (BW) of an optical interconnect is aggregated from a number of n devices operating at a bitrate BR in parallel

$$BW = n \cdot BR. \tag{1}$$

Consequently, the system bandwidth can be limited by:
- the number of channels n;
- the footprint of the devices;
- the bitrate per channel;
- the thermal budget of each device;
- and the thermal crosstalk of large VCSEL arrays.

2.1. Limits Due to Thermal Crosstalk

Scaling up the total bandwidth of an optical interconnect means placing more and more devices more and more densely together. This leads inevitably to thermal crowding, which is a well-known effect for large VCSEL arrays [14]–[16]. Generally speaking, the heat-sink temperature T_{HeatSink} will be raised by a value ΔT due to thermal crowding within an array. On the other hand, any VCSEL device has a maximum operation temperature T_{max}, where it still works within its specifications. This basic physical relationship limits the scaling of optical interconnects. In order to get a more quantitative statement, we calculate the influence of this effect in a spherical coordinate system for round arrays in Cartesian pitch configuration on an infinite heat sink. The thermal distribution $\Delta T_i(s)$, with s as the distance from the heat source, at the surface of the heat sink caused by a single VCSEL is

$$\Delta T_i(s) = \frac{P_{\text{diss}}}{\lambda} \int_s^\infty \frac{dr}{2\pi r^2} = \frac{P_{\text{diss}}}{2\pi\lambda s}. \tag{2}$$
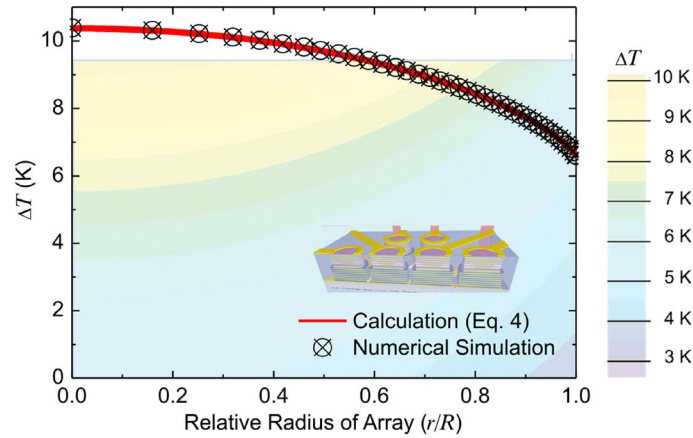
Fig. 2. Already in rather small arrays with less than 1000 VCSELs thermal crowding can no longer be neglected. Here, we depict results for an array of 900 devices; 50-μm pitch; $P_{\mathrm{diss}} = 12$ mW; directly on copper as calculated from (4) (solid line). From numerical finite element simulators we obtain the same results (scatters plus graphical output as background inset).

The integration path and the precise definition of all dimensional parameters can be found in [16]. As a next step, we integrate over all the thermal contributions of all VCSELs within the array of radius R, yielding the thermal distribution of the whole array $\Delta T(r)$. The boundary condition for the thermal distribution is $\Delta T(\infty) = 0$

$$\Delta T(r) = \frac{1}{d^2} \left(\int\limits_0^{R-r} \int\limits_{-\pi}^{+\pi} \Delta T_i(s)\, s\, d\varphi\, ds + \int\limits_{R-r}^{R+r} \int\limits_{-\varphi(s,r)}^{+\varphi(s,r)} \Delta T_i(s)\, s\, d\varphi\, ds \right) \tag{3}$$

$$= \frac{2}{\pi} \cdot \frac{P_{\mathrm{diss}} R}{d^2 \lambda} \cdot E\left(\frac{r}{R}\right) = 2\sqrt{\frac{1}{\pi^3}} \cdot \frac{\sqrt{n} \cdot P_{\mathrm{diss}}}{d \cdot \lambda} \cdot E\left(\frac{r}{R}\right) \tag{4}$$

with

$$E(k) \equiv \int\limits_0^1 \sqrt{\frac{1 - k^2 t^2}{1 - t^2}}\, dt \tag{5}$$

being the complete elliptic integral of the second kind. The thermal crowding from an array of $n = 900$ VCSELs placed at a distance d of 50 μm with 12-mW dissipated power P_{diss} each is depicted in Fig. 2. Even directly placed on an infinite copper heat sink with a good thermal conductivity λ of 400 W/(mK), we can predict a raise in temperature of more than 10 K.

For optical interconnects, still, the devices running at the highest temperatures in the center of the array have to be within the specifications

$$T_{\mathrm{HeatSink}} + \Delta T \leq T_{\mathrm{max}} \tag{6}$$

$$\Rightarrow \Delta T_{\mathrm{max}} = T_{\mathrm{max}} - T_{\mathrm{HeatSink}} \overset{!}{=} \Delta T(0). \tag{7}$$

With (4), we yield

$$\Delta T_{\mathrm{max}} = \Delta T(0) = \sqrt{\frac{1}{\pi}} \cdot \frac{\sqrt{n} \cdot P_{\mathrm{diss}}}{d \cdot \lambda}. \tag{8}$$

The dissipated power $P_{\mathrm{diss}} = UI - P_{\mathrm{opt}}$, with U and I as the VCSELs bias voltage and current and P_{opt} as optical output power, is dependent on the data rate, as VCSELs have to be operated at

sufficiently large current to achieve ultimate photon densities and high relaxation oscillation frequencies. To take this effect into account, we define the heat-to-bitrate ratio as a figure of merit for efficiency

$$HBR = \frac{P_{\text{diss}}}{BR}. \qquad (9)$$

Energy efficiency of VCSELs can be also benchmarked in terms of total electrical energy to bitrate, the so-called energy-to-data ratio (EDR) [17]. EDR and HBR are interconnected by the wall-plug efficiency (WPE)

$$WPE = 1 - HBR/EDR. \qquad (10)$$

The thermally limited maximum bandwidth BW_{max} of an optical interconnect based on VCSEL arrays is then

$$\Rightarrow BW_{\text{max}} = \sqrt{\pi} \cdot \Delta T_{\text{max}} \cdot \frac{d}{HBR} \cdot \sqrt{n} \cdot \lambda \qquad (11)$$

or, for a given area A

$$BW_{\text{max}} = \sqrt{\pi} \cdot \Delta T_{\text{max}} \cdot \sqrt{A} \cdot \frac{\lambda}{HBR} \qquad (12)$$

with the pitch d limited to

$$d_{\text{opt}} = \sqrt{\sqrt{\frac{A}{\pi}} \cdot BR \cdot \frac{HBR}{\lambda \cdot \Delta T_{\text{max}}}} \qquad (13)$$

or a bitrate limitation

$$BR_{\text{max}} = \sqrt{\pi} \cdot \Delta T_{\text{max}} \cdot \frac{\lambda}{HBR} \cdot \frac{d^2}{\sqrt{A}}. \qquad (14)$$

Equation (11) expresses the following dilemma: Scaling is not just putting more channels more densely together, as one might conclude from (1). The maximum achievable bandwidth is only scaling with the square root of the number n of devices, but the cost will at least scale with n if not n^2. High serial bandwidths of each channel (BR) are the way to keep the number of devices and the cost down. On the other hand, as we can see from (12), the maximum bandwidth for a given area does neither depend on the number of devices nor bitrate at which they are operating. These are quite interesting results, as this might seem to be counterintuitive to (1), where we started from. The derived scaling rules tell us the following:

- For large arrays, optical interconnect will suffer from thermal crowding.
- This limits the number of devices that we can put on the array and/or the maximum bandwidth at which we can drive them.
- For a given area, we can find an optimal pitch, limiting the thermal crosstalk to acceptable levels.
- If HBR would be a constant across all bitrates, we could also derive a maximum bitrate at given array geometries.

From (12), we can see that thermal stability and energy efficiency are the two figures of merit to scale bandwidth up. The total bandwidth of an interconnect is not directly dependent on the bitrate per channel, but the number of channels depends on this quantity. Therefore, the system complexity and cost can be kept down by utilizing fewer channels running at higher bitrates.

In the following, we would like to give some realistic numbers based on experimental evidence and discuss the scaling limits of VCSEL technology.

2.2. Limits of Energy Efficiency

In a first-order approximation, for a given directly modulated VCSEL, the resonance frequency f_R rises with the square root of the VCSEL power P_{out}, which is (at low drive currents near threshold I_{th}) proportional to the current I above threshold

$$f_R^2 \propto P_{out} \propto I - I_{th}. \tag{15}$$

Therefore, it is trivial to understand that high-speed VCSELs typically consume more energy per bit when they are operated at higher bitrates. However, VCSELs designed to work at ultra-high bitrates do not necessarily get more energy efficient just by simply reducing the pump current and the bitrate. In order to realize energy-efficient high-speed performance, large resonance frequencies must be achieved at a low drive current. Assuming linear behavior we can define a D-factor and a modified factor D_{BR} based on the bitrate

$$f_R \equiv D\sqrt{I - I_{th}} \propto BR \equiv D_{BR}\sqrt{I - I_{th}}. \tag{16}$$

To model the electrical power consumed by the VCSEL, we use an ideal diode with series resistance and define

$$U \equiv U_{th} + R_d I \tag{17}$$

with U_{th} as threshold voltage and R_d as the differential series resistance. Then, we can write the energy efficiency EDR as follows:

$$EDR = \frac{U \cdot I}{BR} = \frac{1}{BR}\left[\left(\left(\frac{BR}{D_{BR}}\right)^2 + I_{th}\right)U_{th} + \left(\left(\frac{BR}{D_{BR}}\right)^2 + I_{th}\right)^2 R_d\right] \tag{18}$$

$$= U_{th}\left(\frac{BR}{D_{BR}^2} + \frac{I_{th}}{BR}\right) + R_d\left(\left(\frac{BR^{3/4}}{D_{BR}}\right)^2 + \frac{I_{th}}{\sqrt{BR}}\right)^2. \tag{19}$$

Although, depending on the application, either the EDR or the HBR value appears to be a more relevant figure of merit, both figures can be used to judge the energy efficiency. Furthermore, a small value of EDR also translates into an even smaller value of HBR ($HBR < EDR$ as $WPE < 1$). In order to achieve a small value of EDR and an efficient laser, the differential resistance R_D and the laser threshold U_{th} have to be as small as possible. This means that we need to optimize the VCSEL targeting very low electrical losses. A large bitrate would be required to compensate high threshold currents. On the other hand, for a given D-factor, a higher bitrate translates into an inferior efficiency. Lasers with very small threshold currents are thus imperative. If we neglect the threshold current in (19), we obtain

$$EDR|_{I_{th}=0} = U_{th}\frac{BR}{D_{BR}^2} + R_d\frac{BR^3}{D_{BR}^4}. \tag{20}$$

From (20), we learn that low EDR-values become more and more difficult for higher bitrates. A high D-factor, on the other hand, is crucial for efficient lasers. A formula for the D-factor depending on intrinsic laser parameters can be derived from small-signal analysis

$$D = \frac{1}{2\pi}\sqrt{\frac{v_g}{e} \cdot \frac{\eta_i a}{V_{res}}} \tag{21}$$
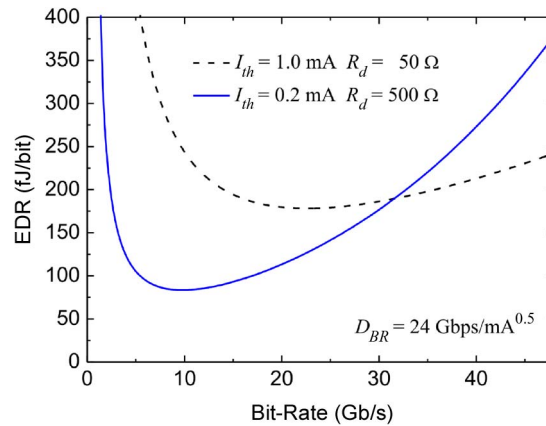
Fig. 3. Calculated efficiency versus bitrate plot based on (20). A typical VCSEL in ultra- energy-efficient design (solid line) and a typical VCSEL optimized for high data rates (dashed line) were modeled. No intrinsic bandwidth limitation is assumed here. Near the bandwidth limitation of real-world devices, the *EDR* raises very fast.
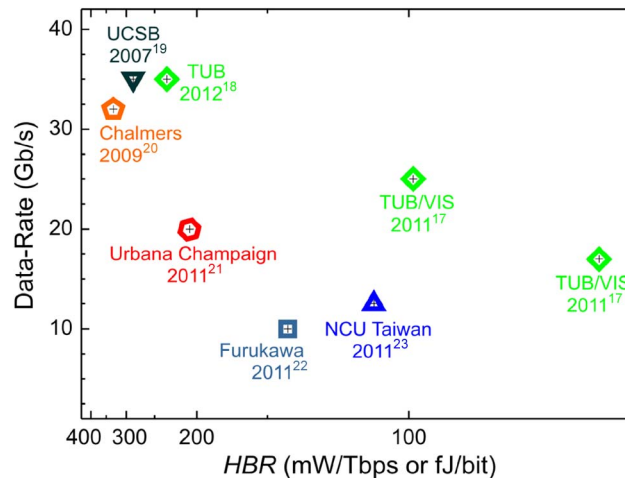

Fig. 4. Energy-efficient VCSEL data links. Data-transmission at efficiencies below 100 mW/Tbps have been demonstrated. Data-transmission at higher bitrates becomes less energy efficient.

with η_i as internal quantum efficiency; V_{res} as the volume of the optical resonator; and a, v_g, and e as differential gain, group velocity, and elementary charge, respectively. Obviously, small-aperture VCSELs are beneficial in many ways for energy-efficient links. Threshold is reached at low currents and the D-factor is larger due to smaller resonator and active volume. At the same time, the differential series resistance R_d will rise however. This tradeoff in design is depicted in Fig. 3. The efficiency versus the bitrate is plotted utilizing (19) and the device values shown in the inset. Extremely efficient devices obviously require ultra-low threshold currents. VCSELs optimized for energy-efficient operation at ultra-high speeds require more energy.

Recently, large progress on energy-efficient VCSELs was made at the Center of Nanophotonics at the Technical University of Berlin by Moser *et al.* [17], [18]. Devices were optimized for highest energy efficiency at different bitrates. The results are presented in Fig. 4 and compared with the results of other groups [19]–[23]. The best *HBR* values reported are 69-mW/Tbps at 17 Gb/s and 99 mW/Tbps at 25 Gb/s, respectively. At 35 Gb/s, 233 mW/Tbps is needed for error-free data transmission. These experimental results are in good agreement with our theoretical predictions presented in Fig. 3.

Data transmission at higher bitrates requires more energy and causes more dissipated heat limiting the maximum performance of large VCSEL arrays. However, we have to note that the energy efficiency of a VCSEL is not independent off the bitrate. The region of best energy efficiency depends on the device design, as shown in Fig. 3. Fortunately, the energy efficiency of VCSELs is quite tolerant toward temperature variations [7], [17], [18] if we keep a certain margin to the maximum device ratings.

2.3. Limits of Bitrates and Thermal Stability

From (12), we can learn that thermal stability and energy efficiency are equally important for VCSEL-based optical interconnects. Large bitrates per channel can keep the number of channels and therefore system complexity and cost down. Consequently, we have to ask about the limits of bitrates and thermal stability.

For applications in digital communication, the maximum speed at which information can be transmitted is depending on the laser's modulation bandwidth. Additionally, the transfer function of the laser should be free of peaking or resonances within the used frequency range. VCSELs, with their high intrinsic damping, are superior in this aspect compared to edge emitters. This is why VCSELs usually allow quite high data rates compared with their bandwidth.

To identify the limits of modulation, a rate-equation analysis is a useful instrument [24], [25]. Although many approximations and assumptions are necessary and numerical methods would give better results, the nonlinearity and uncertainties of real devices are so large that understandable and intuitive formulas are very helpful.

Damping and resonance are coupled by the K-factor giving the limits for overdamping. For easy modeling, parasitics can be added by a first-order low-pass filter. In various textbooks, we can find further calculations finding limits caused by parasitic, device heating, damping, and so on. Unfortunately, the assumptions made to keep these formulas easy might be valid for edge-emitting lasers operating around 2.5 Gb/s, however, do not hold for VCSELs beyond 10 Gb/s.

To identify intrinsic limitations, we have to look at the original equations gained from a small-signal rate-equation analysis

$$4\pi^2 \cdot f_R^2 = v_g^2 \cdot g_{th} \cdot \Gamma \cdot a \cdot S \tag{22}$$

$$K \equiv \frac{\gamma - \gamma_0}{f_R^2} \approx \frac{\gamma}{f_R^2} = \frac{4\pi^2}{v_g} \cdot \frac{1}{g_{th}} \left(\frac{1}{\Gamma} + \frac{a_p}{a} \right). \tag{23}$$

The resonance frequency f_R, as given in (22), is limited by group velocity v_g, threshold gain g_{th}, the confinement Γ of the optical wave, the differential gain a, and the photon density S. For large modulation speed, a large resonance frequency is preferential. By raising f_R, also the damping raises, causing a limitation of the performance. The figure of merit to judge how fast damping rises with increasing resonance frequency is the K-factor, which is also connected to device parameters. The relationship is given in (23) with a_p as the gain compression factor. Although the damping flattens out the response, which is beneficial in some cases [26], small K-factors are preferential for ultimate modulation speeds. Consequently, threshold gain, optical confinement, and differential gain have to be maximized. Additionally, high photon densities at low gain compression are needed. Last but not the least, a very good thermal design is crucial, as otherwise, a high photon density will never be achieved and all other figures of merit will degrade.

Therefore, we identify as real limits for high-speed VCSELs:
- Parasitics, which are directly limiting the overdamped responses;
- Damping, which gives a real physical limit and increases the effect of parasitics;
- Limitations of the relaxation oscillation frequency f_R by device heating, the active medium, and the laser cavity;
- Other nonlinear effects such as transport, spatial hole burning, current crowding, and modal properties.

TABLE 1

The impact of emission-wavelength dependent device technology on potential modulation speed. "+" means that high-speed devices can be realized with small effort utilizing state-of-the-art device technology. "−" means that physics can inhibit good device performance if no technological workaround is applied

	850 nm ...	980 nm ...	1310 nm ...	1550 nm		
Gain	*strained InGaAs on GaAs*			*strained InGaAs on InP*		
	o	+	++	o	+	++
Confinement	*DBR on GaAs*			*DBR on InP*		
	o	+	++	--	-	o
Parasitic	*low C oxide apertures*			*low R tunnel junction*		
	++	++	++	o	+	++
Heat extraction	*ternary*	*binary DBR*		*ternary DBR & Auger*		
	o	++	+	-	--	---

Although these considerations are, in general, independent of the laser wavelength, the desired emission wavelength also has an impact on the potential device speed and thermal properties. This is caused by two reasons. First, longer wavelength devices require lower band-gap materials, which provide inferior carrier confinement. Additionally nonradiative effects such as Auger are more pronounced, and free-carrier absorption rises with the second power of the wavelength [27]. Furthermore, adjustments made to accommodate the desired wavelength (e.g., devices grown on GaAs or InP substrates) also have side effects on the device performance. These facts in a nutshell can be found in Table 1, giving an overview of certain laser wavelengths favored by nature assuming state-of-the-art device technology. The line around 1310 nm represents the switching from GaAs to InP substrate. Longer emission wavelengths go together with higher material strain in the active medium and better gain properties. Fundamental absorption limits the material composition of high-index media at shorter wavelengths. This results in inferior vertical confinement [25]. Low-loss tunnel junctions require low band-gap materials and can therefore be realized more easily for long-wavelength devices. For heat extraction, binary alloys provide one order of magnitude better thermal conductivity than their ternary or quaternary counterparts. It turns out that emission wavelengths around 1 μm seem to be a very promising candidate for laser sources operating at ultimate speeds [19].

Achievements on the field of high-speed temperature-stable VCSELs were made at the Center of Nanophotonics at the Technical University of Berlin [28], [29]. Data-transmission experiments were carried out under ambient temperatures from −14 to +155 °C allowing us to identify error-free operation at bitrates from 12.5 to 49 Gb/s. These results are plotted together with the results of other groups in Fig. 5 [6], [23], [26], [30]–[40] and allow us to estimate the limits of the scalability of large VCSEL arrays for optical interconnects.

2.4. Limits of Scaling VCSEL Arrays

In order to meet the scalability challenge of future optical interconnects, we have to consider the tradeoffs between speed, temperature stability, and energy efficiency. We also have to estimate the technological limits of today's VCSEL technologies. From recently published results, we believe that commercial devices with serial nonreturn-to-zero (NRZ) coded data rates beyond 40 Gb/s are feasible. Devices with 12.5 Gb/s are standard, and 25-Gb/s VCSELs will be the next product offered by many companies in the very near future [26], [40]. Temperature stability can be extended up to 155 °C heat-sink temperature at data rates of 12.5 Gb/s, while heat-sink temperatures of 55 °C and 85 °C can be easily achieved by commercial devices. Highly energy-efficient devices can operate up to 25 Gb/s below 100 fJ/bit; 40 Gb/s below 250 fJ/bit is ambitious but not unphysical [18], [22], while efficiencies around 1000 fJ/bit are industry standard. For VCSEL arrays, a pitch of 250 μm is the industry standard for many years, which could be scaled down to around 25 μm then hitting technological limits.

Let us assume an optical interconnect in a 3-D-stack on CMOS as sketched in Fig. 1. Common die sizes are approximately 100 mm^2 for desktop CPUs and 500 mm^2 for server chips. Those die
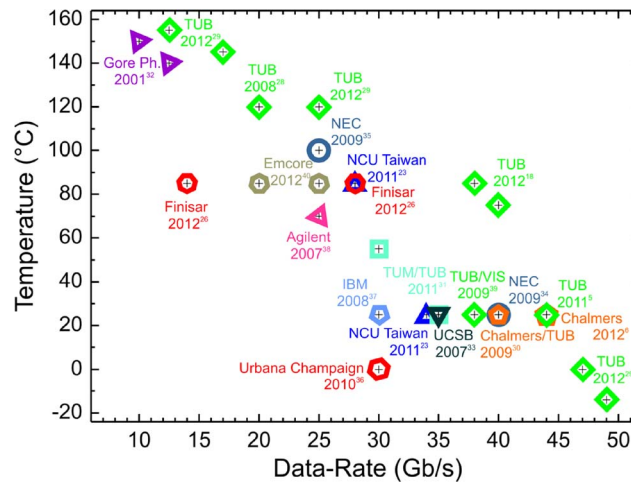
Fig. 5. High-speed temperature-stable VCSEL data links. Data-transmission up to 155 °C has been demonstrated. Data-transmission at highest bitrates requires active device cooling.

sizes were quasi-constant for more than a decade and seem therefore to be reasonable numbers for future projections. For heat sinking, we assume the VCSELs to directly sit on a silicon heat sink with a very good thermal conductivity λ of 120 W/(Km). If this heat sink contains a CPU and a driver chip, it will operate already at high temperature levels. Additionally, VCSELs need an optical fan-out blocking that side for heat extraction. Therefore, we sketch three scenarios: VCSELs that should not exceed 55 °C operating temperature should not suffer from thermal crowding exceeding 5 K, which corresponds to a quite cool die of 50 °C. Temperature-stable VCSELs operating at 85 °C could enable die temperatures of 75 °C and still keep a margin of 10 K allowing for larger arrays. An ultimate temperature scenario would be 155 °C for the VCSELs and a $\triangle T_{max}$ of 70 K corresponding to 85 °C original die temperature. Of course, one has to note that the die temperature will rise with the VCSEL array on top. However, highly temperature-stable silicon chips are already in the market for quite some time fulfilling rugged military standards.

Now, we are ready to answer the following questions:

- How many Gb/s can we pack on a mm^2?
- Are we limited by the *footprint* of the VCSELs?

These field-relevant numbers are now combined to several scenarios and evaluated by (11)–(13). The results are given in Table 2.

Scenario #1 demonstrates what could be achieved with easily available VCSELs. If these devices are integrated in a 3-D stack on CMOS, desktop-die-sized arrays can deliver bandwidths beyond 10 Tb/s. A bandwidth density of 100 Gbps/mm^2 can be achieved with quite sparse pitches around 350 μm, limiting thermal crowding to 5 K. The other scenarios show the potential scalability of this technology.

Scenario #2 is for large arrays of next-generation commercial VCSELs for high-speed operation at 25 Gb/s in large server-die-sized arrays. Since the bandwidth only scales with the square root of the array area, the solitary devices have to be twice as temperature stable to keep the bandwidth density at a constant level and allowing scaling up the total bandwidth. Very sparse pitches beyond 500 μm are required to keep thermal crowding at a moderate level.

Utilizing ultra-high-speed devices (scenario #3) reduce the number of channels and therefore the complexity of the optical bus, but neither deliver more bandwidth nor a larger bandwidth density. If we do not like to operate these devices at ultimate temperatures, which might be a wise decision having reliability in mind, energy efficiency is the way to scale up. Reducing the dissipated heat per bit by a factor of four quadruples total bandwidth and bandwidth density (scenario #4). If we need to scale up the bandwidth further, and we can handle tenth or hundreds of thousand optical channels, boosting energy efficiency up on the cost of per channel bitrate is the way to go. Scenario #5 shows

TABLE 2

Several scenarios of future VCSEL-based optical interconnects integrated in a 3-D stack on CMOS as depicted in Fig. 1 are evaluated by (11)–(13). These estimates demonstrate the potential of a certain technology and identify scalability challenges for VCSEL-based multi-100-Gb/s systems

Scenario # Bit-Rate	Die Size	Efficiency	Temperature ΔT_{max}	Array Size (# of VCSELs) ($\approx$ n x n Array)	minimum Pitch d_{opt}	Bandwidth Density	total Bandwidth BW_{max}
conservative, *standard* components							
1 12.5 Gb/s	100 mm^2	1000 fJ/bit	55°C 5 K	851 VCSEL 29 x 29 array	343 μm	106 Gbps/mm^2	**10.6 Tb/s**
next generation industry devices							
2 25 Gb/s	500 mm^2	1000 fJ/bit	85°C 10 K	1902 VCSEL 44 x 43 array	513 μm	95 Gbps/mm^2	**47.6 Tb/s**
ultra-high-speed devices							
3 40 Gb/s	500 mm^2	1000 fJ/bit	85°C 10 K	1189 VCSEL 35 x 34 array	649 μm	95 Gbps/mm^2	**47.6 Tb/s**
ultra-high-speed *energy-efficient* devices							
4 40 Gb/s	500 mm^2	250 fJ/bit	85°C 10 K	4756 VCSEL 69 x 69 array	324 μm	380 Gbps/mm^2	**190.2 Tb/s**
ultra-energy-efficient devices							
5 25 Gb/s	500 mm^2	100 fJ/bit	85°C 10 K	19 024 138 x 138	162 μm	951 Gbps/mm^2	**475.6 Tb/s**
ultra-temperature-stable, energy-efficient devices							
6 12.5 Gb/s	100 mm^2	100 fJ/bit	155°C 70 K	119 109 345 x 345	29 μm	14.9 Tbps/mm^2	**1.489 Pb/s**

that 1 Tbps/mm^2 can be achieved with recently published ultra-energy-efficient VCSELs. Finally, we would like to address the question whether we are facing a footprint limit of VCSELs soon. Scenario #6 is maximized to stress the footprint requirements. Squeezing ultra-temperature-stable ultra-energy-efficient VCSELs on a small array would allow reducing the pitch down to 29 μm. This is still in the technological range of VCSELs.

These considerations show that VCSEL-based optical interconnects can fulfill the requirements of future high-performance computers for about one decade. The bandwidth is scalable from the Terabyte/second to the Petabyte/second range. This gives potential successor technologies like electrooptically-modulated VCSELs [41], [42] or nanolasers [12], [13] time to mature. Today's VCSELs are limited by their still too large heat produced rather than their large footprint. Successor technologies have to show clear advantages in energy efficiency.

This paper focused on the emitter side. Similar estimates have to be made for receivers [5]. Devices that can both operate as sender or receiver might also be of interest [43]. Handling the optical fan-out with thousands of optical channels and electrical multichannel drivers are other big challenge. Furthermore, efficient VCSEL driver chips in CMOS are another open issue, which still needs to be solved.

3. Conclusion

Based on present experimental data, VCSEL-based optical interconnect technology can, in principle, deliver the highly ambitious 80-Tbps IO bandwidth envisioned for future 10-Teraflop server CPUs. The bandwidth density is scalable from reasonably 100 Gbps/mm^2 to a physical limit (based on today's device technology) around 15 Tbps/mm^2. More energy-efficient devices dissipating less heat are the best solution for scaling the bandwidth up. A high serial bandwidth of the individual channels just helps to keep the number of channels and the complexity of the optical output low but does not enable higher IO throughput.

References

[1] B. Offrein, "Silicon photonics packaging requirements," in *Proc. IBM Silicon Photon. Workshop*, Munich, Germany, May 23, 2011, pp. 1–14. [Online]. Available: http://www.siliconphotonics.eu/munich_slides/2_IBM.pdf

 [2] D. Miller, "Device requirements for optical interconnects to silicon chips," *Proc. IEEE*, vol. 97, no. 7, pp. 1166–1185, Jul. 2009.
 [3] *International Technology Roadmap for Semiconductors*, Jul. 2012. [Online]. Available: http://www.itrs.net/Links/2007ITRS/ExecSum2007.pdf
 [4] F. Doany, C. Schow, C. Baks, D. Kuchta, P. Pepeljugoski, L. Schares, R. Budd, F. Libsch, R. Dangel, F. Horst, B. Offrein, and J. Kash, "160 Gb/s bidirectional polymer-waveguide board-level optical interconnects using CMOS-based transceivers," *IEEE Trans. Adv. Packag.*, vol. 32, no. 2, pp. 345–359, May 2009.
 [5] W. Hofmann, P. Moser, P. Wolf, A. Mutig, M. Kroh, and D. Bimberg, "44 Gb/s VCSEL for optical interconnects," in *Proc. OFC/NFOEC*, 2011, pp. 1–3.
 [6] P. Westbergh, R. Safaisini, E. Haglund, B. Kögl, J. Gustavsson, A. Larsson, M. Green, R. Lawrence, and A. Joel, "High-speed 850 nm VCSELs with 28 GHz modulation bandwidth operating error-free up to 44 Gb/s," *Electron. Lett.*, vol. 48, no. 18, pp. 1145–1147, Aug. 2012.
 [7] A. Vahdat, H. Liu, X. Zhao, and C. Johnson, "The emerging optical data center," in *Proc. OFC/NFOEC*, 2011, pp. 1–3.
 [8] L. Schares, J. Kash, F. Doany, C. Schow, C. Schuster, D. Kuchta, P. Pepeljugoski, J. Trewhella, C. Baks, R. John, L. Shan, Y. Kwark, R. Budd, P. Chiniwalla, F. Libsch, J. Rosner, C. Tsang, C. Patel, J. Schaub, R. Dangel, F. Horst, B. Offrein, D. Kucharski, D. Guckenberger, S. Hegde, H. Nyikal, C.-K. Lin, A. Tandon, G. R. Trott, M. Nystrom, D. P. Bour, M. R. T. Tan, and D. W. Dolfi, "Terabus: Terabit/second-class card-level optical interconnect technologies," *IEEE J. Sel. Topics Quantum Electron.*, vol. 12, no. 5, pp. 1032–1044, Sep./Oct. 2006.
 [9] W. Hofmann, P. Moser, and D. Bimberg, "Energy efficient interconnects," *IEEE Photon. J.*, vol. 4, no. 2, pp. 650–656, Apr. 2012.
[10] W. Hofmann, "Evolution of high-speed long-wavelength vertical-cavity surface-emitting lasers," *Semicond. Sci. Technol.*, vol. 26, no. 1, p. 014011, Jan. 2011.
[11] M. Laemmlin, G. Fiol, M. Kuntz, F. Hopfer, A. Mutig, N. Ledentsov, A. Kovsh, C. Schubert, A. Jacob, A. Umbach, and D. Bimberg, "Quantum dot based photonic devices at 1.3μm: Direct modulation, mode-locking, SOAs and VCSELs," *Phys. Stat. Sol. (C), Appl. Res.*, vol. 3, no. 3, pp. 391–394, Mar. 2006.
[12] S. Chuang and D. Bimberg, "Metal-cavity nanolasers," *IEEE Photon. J.*, vol. 3, no. 2, pp. 288–292, Apr. 2011.
[13] A. Matsudaira, C.-Y. Lu, M. Zhang, S. Chuang, E. Stock, and D. Bimberg, "Cavity-volume scaling law of quantum-dot metal-cavity surface-emitting microlasers," *IEEE Photon. J.*, vol. 4, no. 4, pp. 1103–1114, Aug. 2012.
[14] D. Francis, H.-L. Chen, W. Yuen, G. Li, and C. Chang-Hasnain, "Monolithic 2D-VCSEL array with $>$ 2 W cw and $>$ 5 W pulsed output power," *Electron. Lett.*, vol. 34, no. 22, pp. 2132–2133, Oct. 1998.
[15] J-F. Seurin, C. Ghosh, V. Khalfin, A. Miglo, G. Xu, J. Wynn, P. Pradhan, and L. D'Asaro, "High-power high-efficiency 2D VCSEL arrays," in *Proc. SPIE*, 2008, vol. 6908, p. 690 808.
[16] M.-C. Amann and W. Hofmann, "InP-based long-wavelength VCSELs and VCSEL arrays," *IEEE J. Sel. Topics Quantum Electron.*, vol. 15, no. 3, pp. 861–868, May/Jun. 2009.
[17] P. Moser, W. Hofmann, W. P. Wolf, J. Lott, G. Larisch, A. Payusov, N. Ledentsov, and D. Bimberg, "81 fJ/bit energy-to-data ratio of 850 nm vertical-cavity surface-emitting lasers for optical interconnects," *Appl. Phys. Lett.*, vol. 98, no. 23, p. 231 106, Jun. 2011.
[18] P. Moser, P. Wolf, A. Mutig, G. Larisch, W. Unrau, W. Hofmann, and D. Bimberg, "85 °C error-free operation at 38 Gb/s of oxide-confined 980-nm vertical-cavity surface-emitting lasers," *Appl. Phys. Lett.*, vol. 100, no. 8, p. 081103, Feb. 2012.
[19] Y. Chang and L. Coldren, "Efficient, high-data-rate, tapered oxide-aperture vertical-cavity surface-emitting lasers," *IEEE J. Sel. Topics Quantum Electron.*, vol. 15, no. 3, pp. 704–715, May/Jun. 2009.
[20] P. Westbergh, J. Gustavsson, A. Haglund, A. Larsson, F. Hopfer, G. Fiol, D. Bimberg, and A. Joel, "32 Gb/s multimode fibre transmission using high-speed, low current density 850 nm VCSEL," *Electron. Lett.*, vol. 45, no. 7, pp. 366–368, Mar. 2009.
[21] F. Tan, C. H. Wu, M. Feng, and N. Holonyak, "Energy efficient microcavity lasers with 20 and 40 Gb/s data transmission," *Appl. Phys. Lett.*, vol. 98, no. 19, p. 191 107, May 2011.
[22] S. Imai, K. Takaki, S. Kamiya, H. Shimizu, J. Yoshida, Y. Kawakita, T. Takagi, K. Hiraiwa, H. Shimizu, T. Suzuki, N. Iwai, T. Ishikawa, N. Tsukiji, and A. Kasukawa, "Recorded low power dissipation in highly reliable 1060-nm VCSELs for 'Green' optical interconnection," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 6, pp. 1614–1620, Nov./Dec. 2011.
[23] J. Shi, W. Weng, F. Kuo, and J. Chyi, "Oxide-relief vertical-cavity surface-emitting lasers with extremely high data-rate/power-dissipation ratios," in *Proc. OFC/NFOEC*, Los Angeles, CA, 2011, pp. 1–3.
[24] L. Coldren and S. Corzine, "Diode lasers and photonic integrated circuits," in *Dynamic Effects*. New York: Wiley, 1995, pp. 184–212.
[25] W. Hofmann, "InP-based long-wavelength VCSELs and VCSEL arrays for high-speed optical communication," in *Laser Dynamics*. Munich, Germany: Technische Univ., 2009, pp. 25–40. [Online]. Available: http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss-20081119-679286-1-5
[26] L. Graham, H. Chen, D. Gazula, T. Gray, J. Guenter, B. Hawkins, R. Johnson, C. Kocot, A. MacInnes, G. Landry, and J. Tatum, "The next generation of high speed VCSELs at Finisar," in *Proc. SPIE*, 2012, vol. 8276, p. 827 602.
[27] J. Buus and M.-C. Amann, *Tunable Laser Diodes and Related Optical Sources*. Weinheim, Germany: Wiley-VCH, 2005.
[28] A. Mutig, G. Fiol, P. Moser, D. Arsenijevic, V. Shchukin, N. Ledentsov, S. Mikhrin, I. Krestnikov, D. Livshits, A. Kovsh, F. Hopfer, and D. Bimberg, "120 °C 20 Gb/s operation of 980 nm VCSEL," *Electron. Lett.*, vol. 44, no. 22, pp. 1305–1306, Jul. 2008.
[29] P. Wolf, P. Moser, G. Larisch, M. Kroh, A. Mutig, W. Unrau, W. Hofmann, and D. Bimberg, "High-performance 980 nm VCSELs for 12.5 Gbit/s data transmission at 155 °C and 49 Gb/s at −14 °C," *Electron. Lett.*, vol. 48, no. 7, pp. 389–390, Mar. 2012.
[30] P. Westbergh, J. Gustavsson, B. Kögel, A. Haglund, A. Larsson, A. Mutig, A. Nadtochiy, D. Bimberg, and A. Joel, "40 Gb/s error-free operation of oxide-confined 850 nm VCSEL," *Electron. Lett.*, vol. 46, no. 14, pp. 1014–1016, Jul. 2010.
[31] W. Hofmann, M. Müller, P. Wolf, A. Mutig, T. Gründl, G. Böhm, D. Bimberg, and M.-C. Amann, "40 Gb/s modulation of 1550 nm VCSEL," *Electron. Lett.*, vol. 47, no. 4, pp. 270–271, Feb. 2011.

[32] F. Peters and M. MacDougal, "High-speed high-temperature operation of vertical-cavity surface-emitting lasers," *IEEE Photon. Technol. Lett.*, vol. 13, no. 7, pp. 645–647, Jul. 2001.

[33] Y. Chang, C. Wang, and L. Coldren, "High-efficiency, high-speed VCSELs with 35 Gb/s error-free operation," *Electron. Lett.*, vol. 43, no. 19, pp. 1022–1023, Sep. 2007.

[34] N. Suzuki, T. Anan, H. Hatakeyama, K. Fukatsu, K. Yashiki, K. Tokutome, T. Akagawa, and M. Tsuji, "High speed 1.1 μm range InGaAs-based VCSELs," *IEICE Trans. Electron.*, vol. E92-C, no. 7, pp. 942–950, 2009.

[35] H. Hatakeyama, T. Akagawa, K. Fukatsu, N. Suzuki, K. Tokutome, K. Yashiki, T. Anan, and M. Tsuji, "25 Gb/s 100 °C operation of highly reliable InGaAs/GaAsP VCSELs," *Electron. Lett.*, vol. 45, no. 1, pp. 45–46, Jan. 2009.

[36] Z. Tian, C. Chen, K. Choquette, and D. Plant, "30 Gb/s direct modulation of holey VCSELs with thermoelectric cooling," presented at the Proc. Conf. Lasers Electro-Optics (CLEO/QELS), San Jose, CA, 2010, Paper CME3.

[37] R. Johnson and D. Kuchta, "30 Gb/s directly modulated 850 nm datacom VCSELs," presented at the Proc. Conf. Lasers Electro-Optics (CLEO/QELS), San Jose, CA, 2008, Paper CPDB2.

[38] C. Lin, A. Tandon, K. Djordjev, S. Corzine, and M. Tan, "High speed 985 nm bottom emitting VCSEL arrays for chip-to-chip parallel optical interconnects," *IEEE J. Quantum Electron.*, vol. 13, no. 5, pp. 1332–1339, Sep./Oct. 2007.

[39] S. Blokhin, J. Lott, A. Mutig, G. Fiol, N. Ledentsov, M. Maximov, A. Nadtochiy, V. Shchukin, and D. Bimberg, "Oxide confined 850 nm VCSELs operating at bit rates up to 40 Gb/s," *Electron. Lett.*, vol. 45, no. 10, pp. 501–503, May 2009.

[40] N. Li, C. Xie, W. Luo, C. Helms, L. Wang, C. Liu, Q. Sun, S. Huang, C. Lei, K. Jackson, and R. Carson, "Emcore's 1 Gb/s to 25 Gb/s VCSELs," *Proc. SPIE*, vol. 8276, p. 827603, 2012.

[41] V. Shchukin, N. Ledentsov, J. Lott, H. Quast, F. Hopfer, L. Karachinsky, M. Kuntz, P. Moser, A. Mutig, A. Strittmatter, V. Kalosha, and D. Bimberg, "Ultrahigh-speed electrooptically-modulated VCSELs: Modeling and experimental results," *Proc. SPIE*, vol. 6889, p. 688 90H-1, Feb. 2008.

[42] T. Germann, W. Hofmann, A. Nadtochiy, J.-H. Schulze, A. Mutig, A. Strittmatter, and D. Bimberg, "Electro-optical resonance modulation of vertical-cavity surface-emitting lasers," *Opt. Exp.*, vol. 20, no. 5, pp. 5099–5107, Feb. 2012.

[43] Q. Gu, W. Hofmann, M.-C. Amann, and L. Chrostowski, "Optically injection-locked VCSEL as a duplex transmitter/receiver," *IEEE Photon. Technol. Lett.*, vol. 20, no. 7, pp. 463–465, Apr. 2008.