

Ultra-fast VCSELs promise to turbocharge chip communication

The copper interconnects that route chip-to-chip data transfer are starting to reach their speed limit. But this looming bottleneck can be overcome by switching to ultra-fast VCSELs with tiny threshold currents, say **Yu-Chia Chang** and **Larry Coldren** from the University of California, Santa Barbara.

Finally the silicon folks might actually need us. The copper interconnects that they are using are running out of steam, and this is starting to obstruct their advances in interchip communication. It's a weakness that's even threatening to bump Moore's law off course because the electrical input and output functions needed in tomorrow's state-of-the-art microprocessors will dissipate more power and demand a greater share of a chip's real estate. Even today data transfer is a complicated task. Loss, dispersion and cross-talk have to be addressed with predistortion and sophisticated pulse recovery techniques in the transmit and receive stages, respectively.

Against this background of increasing complexity, optical interconnection schemes look like a good bet for the future. This technology offers several key advantages over copper, such as lower signal delay, higher bandwidth, reduced power consumption and freedom from electromagnetic interference. In addition, it has the potential to deliver intrachip communication.

However, optical interconnects will only offer a practical solution if improvements are made to devices and interface systems. Copper is a viable solution at speeds of up to 10 Gbit/s, so data rates must be at least 20 Gbit/s for optoelectronics to be considered an option. Low power consumption is another prerequisite, alongside the ability to be integrated with silicon electronics. But there is no point in trying to meet any of these goals unless the emitters and detectors can be manufactured in large volumes with high yields.

At the University of California, Santa Barbara (UCSB), we are working on a device that could be the key component in chip-to-chip optical interconnects – a miniature, high-speed, efficient VCSEL. It has evolved out of our part of the US Defense Advanced Research Projects Agency's chip-to-chip optical interconnects program, which also supports IBM's "Terabus" project. It aims to demonstrate complete optical links between chips using polymer waveguides (figure 1). It's a technology that's competing with silicon photonics, but which offers lower power dissipation – the primary limiter for

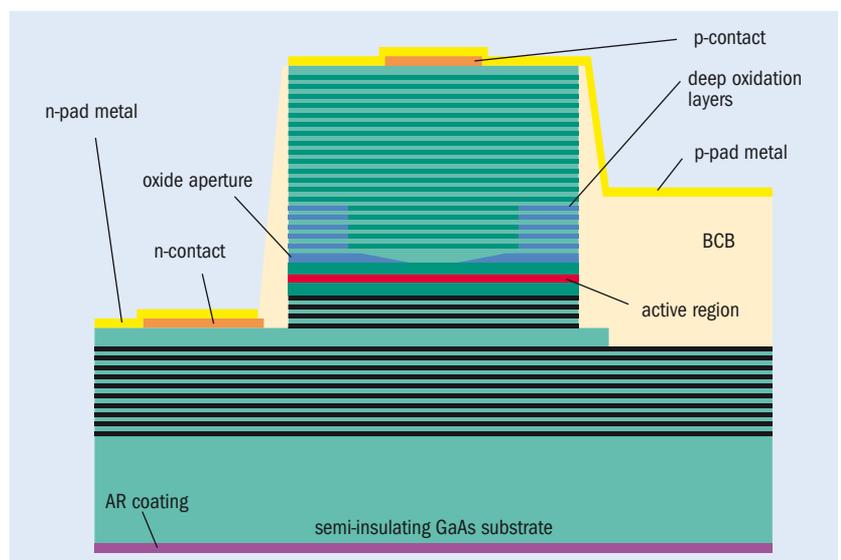
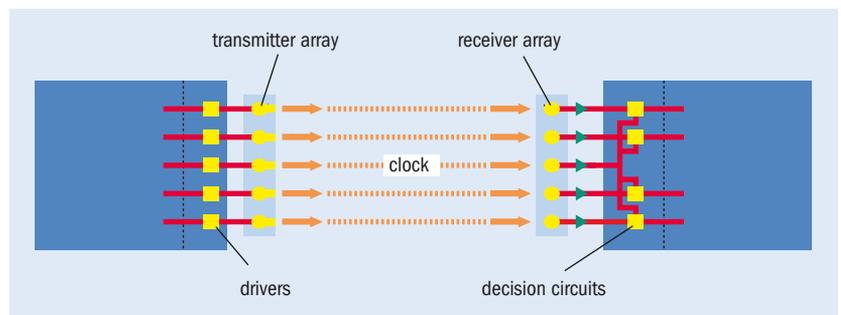


Fig. 1. (top) UCSB's VCSEL could be used to form an array of transmitters that provide chip-to-chip interconnects. **Fig. 2.** (bottom) The UCSB InGaAs/GaAs quantum-well design features a blunter tapered aperture that enables very high modulation speeds at acceptable optical losses.

continued scaling of processor speed.

There are several benefits that VCSELs offer over their edge-emitting cousins for optical interconnect applications: they are smaller; they are easier to fabricate in arrays; they support on-wafer testing; they offer high-speed operation at lower power consumption; and they are cheaper. However, many of the VCSELs that have been designed for optical interconnects have large diameter apertures – typically

Fig. 3. (right) This cross-sectional scanning electron image reveals key features in the UCSB VCSEL – five deep oxidation layers and a tapered oxide aperture. **Fig. 4.** (far right) UCSB’s high-speed VCSELs are fabricated with the following process steps: **(a)** mesa etch, **(b)** oxidation, **(c)** n-metal deposition, **(d)** p-metal deposition, **(e)** n-contact removal, **(f)** patterning with a low-dielectric resin, **(g)** pad metal deposition, and **(h)** adding an anti-reflection coating.



5–8 μm – that drive up the bias current required for a high modulation bandwidth. We avoid this pitfall by scaling down the VCSEL size. Although this leads to a power drop of 50%, it cuts power dissipation and increases the bandwidth, while maintaining acceptable optical losses.

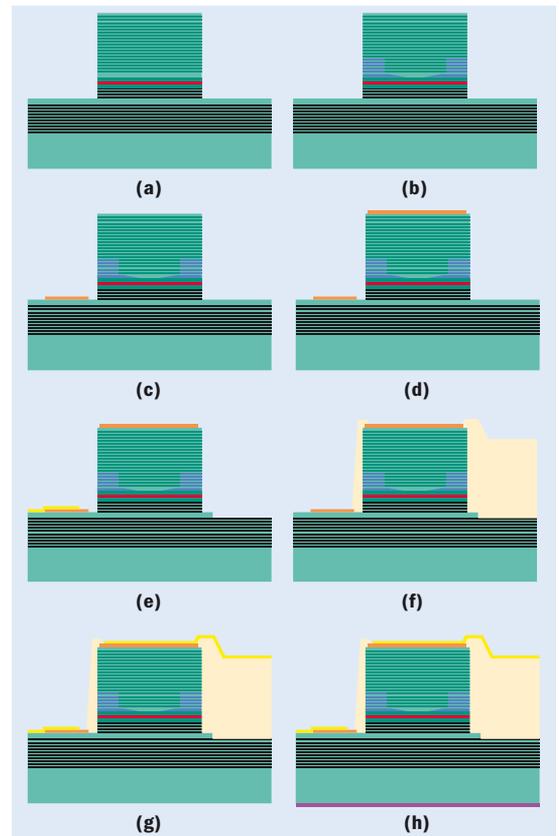
The devices that we have made are 980 nm bottom-emitting VCSELs with a tapered oxide aperture and a contact inserted within the n-type cavity (figure 2, p29). The standard emission wavelength of 850 nm was not selected because the strained InGaAs/GaAs quantum wells (QWs) that produce 980 nm emission deliver a higher intrinsic modulation bandwidth thanks to higher gain at lower carrier densities. Some studies suggest that InGaAs QW-lasers are more reliable. But more importantly, this particular design is compatible with bottom-emission thanks to GaAs’ transparency at 980 nm.

Bottom emission means that flip-chip bonding can be used to integrate our VCSELs with electronics. This eliminates wire bond inductance and the need for wire bonding. Backside microlenses can be added to collimate the output beam, which improves alignment tolerance and cuts packaging costs. This combination of advantages should equip our devices with the compatibility required for advanced computer and processor interconnection architectures and other datacom applications.

Our VCSEL’s key feature compared with prior art is a blunter tapered aperture and a larger oxide thickness next to the active region (figure 3). This provides a high degree of optical confinement and enables our laser to combine high speeds with acceptable optical losses. Employing a longer tapered section would reduce optical losses even further, but this results in larger mode volumes with an unacceptable penalty of much slower modulation speeds.

The parasitic capacitance of a relatively thin oxide aperture can restrict the modulation bandwidth of oxide-confined VCSELs. The capacitance is often reduced by proton implantation, which creates a thick, highly resistive region. However, this highly effective method requires additional process steps that threaten reliability, and increase fabrication costs and parasitic resistance.

We get around this problem by creating additional deep oxidation layers above the confining aperture, which lower the parasitic capacitance. By increasing the aluminum fraction of the AlGaAs layers in



the first several periods of the top distributed Bragg reflector, we can simultaneously form deep oxidation layers along with the oxide aperture.

It is easy to incorporate this simple approach into oxide-confined VCSELs with semiconductor mirrors. No modifications to the fabrication techniques are required, and it increases the refractive index contrast in the unoxidized region where optical modes exist, thanks to the layers with higher aluminum content. This means that mode volume is cut because the longitudinal mode is more confined. On top of this, reductions in parasitic resistance are delivered with minimal impact to optical loss, thanks to high-quality engineering of the band-structure and the p-doping profiles in the top mirror.

Our approach to VCSEL fabrication is compatible with existing manufacturing processes and begins with MBE growth of the epitaxial structure on semi-insulating (100) GaAs substrates (see figure 4 for an overview). Etching creates cylindrical mesas that expose the n-GaAs contact layer, before wet oxidation forms oxide apertures and deep oxidation layers. Metal evaporation adds the p-type Ti/Pt/Au and n-type Au/Ge/Ni/Au contacts.

Pad capacitance is reduced by removing the part of the n-GaAs contact layer (RF-ground) that lies beneath the p-pad metal (RF-signal). The low dielectric constant resin benzocyclobutene is inserted in its place before vias are opened to expose the contacts and Ti/Au is deposited for the pad metals. Finally, an anti-reflection coating is

“This device’s bandwidth makes it the fastest ever 980 nm VCSEL.”

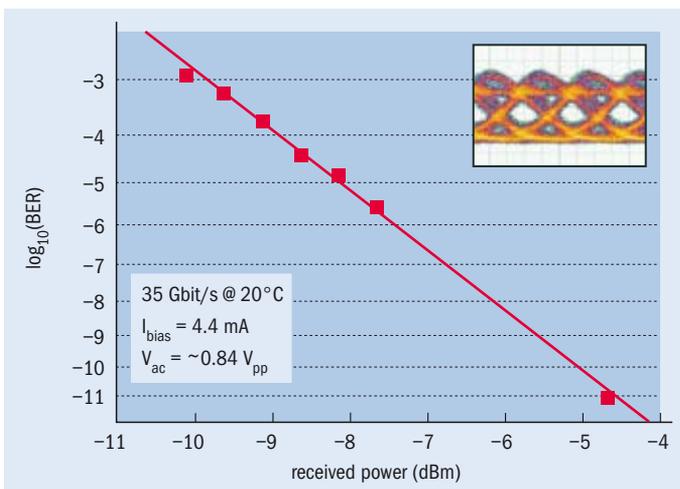
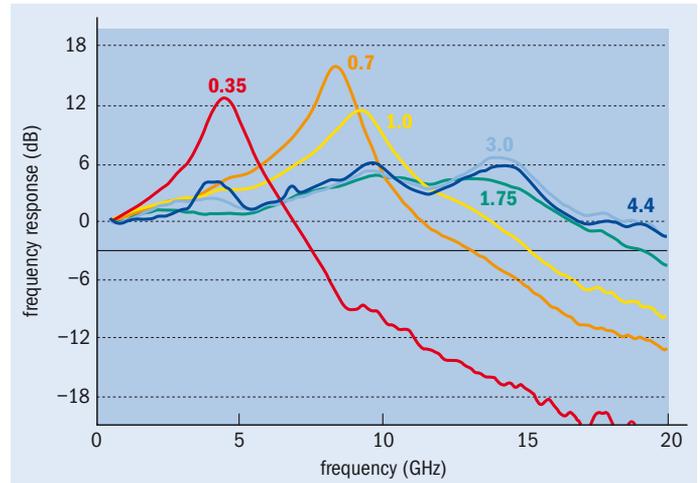
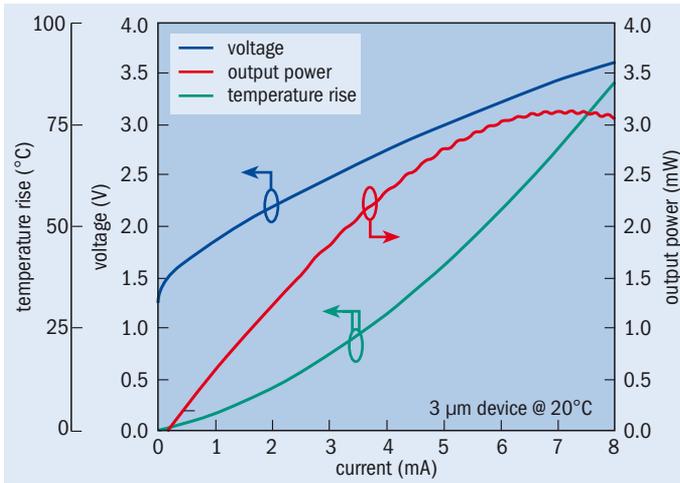


Fig. 5. (top left) UCSB's 3 μm diameter aperture VCSELs deliver up to 3 mW at 20 °C. This laser has a series resistance of approximately 250 Ω at 4.4 mA bias current, the conditions that are used for the large signal experiments. Thermal impedance is 3.3 °C/mW and the temperature rise is 33 °C at this current, but improvements may be possible through refinements in the design of the VCSELs.

Fig. 6. (top right) Adjustments in bias current from 0.35–4.4 mA alter the modulation speed of the 3 μm diameter device, which is highest at 4.4 mA. At this bias the modulation exceeds 20 GHz, which is a record for 980 nm VCSELs.

Fig. 7. (left) Bit error rates (BERs) were measured for UCSB's 35 Gbit/s, 3 μm diameter device at 20 °C. This VCSEL was biased at 4.4 mA and a peak-to-peak voltage swing of 0.84 V was used. The input was a non-return-to-zero signal with 2^7-1 word length. Eye diagrams are shown in the inset figure, which reveal that the eye is clearly open. The extinction ratio is 5.4 dB. In the BER curve, a variable optical attenuator (VOA) was used to take all but the lowest data point. Poor detector responsivity and a 3 dB insertion loss of the VOA prevented the measurement of BERs between 10^{-4} and 10^{-7} . To ensure measurement accuracy, the lowest data point was gated for 30 minutes with a total of 583 errors.

applied to reduce backside reflection.

We have produced 3 μm diameter aperture VCSELs with this process. These emitters have a slope efficiency of 0.67 W/A, which corresponds to a differential quantum efficiency of 54% (figure 5). Threshold current is just 0.144 mA, which is comparatively low for high-speed VCSELs as they usually have threshold currents of at least 0.4 mA. Peak wall-plug efficiency is 31%, maximum output power is 3.1 mW and the threshold voltage is just 1.47 V – very low for such a small device as it is only 220 meV larger than the photon energy.

Small-signal modulation measurements on our 3 μm aperture emitter revealed that it could respond to frequencies beyond 20 GHz, the upper detection limit for our instrument (figure 4). This device's bandwidth clearly exceeds that frequency, which makes it the fastest ever 980 nm VCSEL.

We evaluated this device's suitability for real-world systems with large-signal digital modulation experiments. At 20 °C and a modulation rate of 35 Gbit/s, our 3 μm diameter device produced bit error rates as low as 10^{-11} over a couple of meters of fiber, which demonstrates that it is suitable for high-quality interchip communication. The VCSEL

power dissipation, excluding RF circuitry, is only 10 mW, which corresponds to the highest data rate/power-dissipation ratio that has ever been obtained for any type of laser source.

We are now targeting data rates of 40 Gbit/s, which can be reached by simply extending our existing techniques. However, we are also engaged in another project that has the potential to get us to the 100 Gbit/s range with a single channel. If such high speeds could be achieved, then this will further enhance the attractiveness of employing VCSELs for chip-to-chip interconnects.

Further reading

T Anan 2007 *Proc. International Symposium on VCSELs and Integrated Photonics* E3.
 Y-C Chang *et al.* 2006 *Electron. Lett.* **42** 1281.
 Y-C Chang *et al.* 2007 *Electron. Lett.* **43** 1022.
 E Hegblom *et al.* 1997 *IEEE J. Select. Topics Quantum Electron.* **3** 379.
 A A L-Omari *et al.* 2004 *IEEE Photon. Technol. Lett.* **16** 969.
 L Schares *et al.* 2006 *IEEE J. Sel. Topics Quantum Electron.* **12** 1032.
 E Strzelecka *et al.* 1995 *Electron. Lett.* **31** 724.



About the authors

Larry Coldren (left) is the Fred Kavli professor of optoelectronics and sensors at UCSB, and the director of the university's Optoelectronics Technology Center. He joined UCSB in 1984 after spending 13 years as a researcher at Bell Laboratories. **Yu-Chia Chang** (right) is a PhD student in the electrical and computer engineering department at UCSB. His research interests include the design, growth and characterization of high-speed VCSELs.