



# Lighting up time

Silicon photonics is the favoured approach to moving vast amounts of data around tomorrow's exascale computing systems. By **David Boothroyd**.

**T**he 21st Century is experiencing a data avalanche. Not long ago, a terabyte ( $10^{12}$ byte) was seen as a vast amount of data. Today, a 1Tbyte hard drive for a desktop pc costs less than £50. Soon, much the same will happen at the petabyte ( $10^{15}$ byte) level. And plans are well underway for the exabyte, or  $10^{18}$ byte, generation.

But we are already dealing in exabytes – estimates suggest that around 21exabytes travel over the internet per month. And a similar trend can be seen in processing. For example, the US government's 2012 budget has allocated \$124million for the development of exabyte computers, which should arrive by around 2018 to 2020.

An exascale computer will be 1000 times more powerful than today's most powerful machine, Fujitsu's K Computer. Currently, the top 10 supercomputers all achieve petaflop/s performance, but perhaps the most revealing statistic featured in the TOP500 list (see box) of the most powerful computers is not their processing capability, but their power consumption.

There are 29 systems on the list that consume more than 1MW, with the K Computer using 9.89MW. Average power consumption of a top 10 system is 4.3MW, up from 3.2MW only six months ago. Figures like this explain why many people working on exascale computing think that achieving the processing capability is not the main challenge. Rather, it will be the need to

move vast amounts of information at enormous speed, which will consume too much power and generate too much heat.

"Some people say that flops are almost free, that really what you are paying for is moving the data," says Sudip Dosanjh, who leads an exascale project at the US Sandia and Oak Ridge laboratories.

Mario Paniccia, director of Intel's Photonics Technology lab, agrees. "Ultimately, achieving exascale computing is all about power. To get there, you have to connect lots of systems together and the connection itself will probably cost more than the processors and servers. The power consumption of the links is everything."

Ultimately, Paniccia and others believe there is only one solution: data movements in the exascale era must use light. Only by converting electronics into light will it be possible to handle the vast data bandwidth exascale requires. But until now, converting electronic signals into light for transportation, then back again for more processing, has been too expensive. What is needed is the kinds of integration and scaling up, with the associated dramatic cost reductions, that have been the hallmark of ics. Can that be achieved?

If it is, it will probably be thanks to a particular technology: silicon photonics (SP), the use of silicon as an optical medium. The hope is that, by integrating optics into silicon, such an approach will provide the economies

of scale, ease of manufacture and low costs that have been such critical features of the whole ic industry. The work done over the last few years by Paniccia and his team and others worldwide strongly suggests that SP can create a viable platform for exascale computing.

"I think everyone is now convinced that the path to high volume photonics is with SP and that it can be done," he said. "Ten years ago, that was far from being true: most people thought it was crazy, because silicon was seen as a poor optical material. So we had to start by proving it was possible by developing the basic building blocks."

Slowly, Intel and others did that. Components developed include the first continuous wave silicon Raman laser, a series of silicon modulators starting from 1Gbit/s and now reaching today's rate of 50Gbit/s, the first hybrid silicon laser and 40Gbit/s photo detectors.

Intel's 50Gbit/s Silicon Photonics Link is a fibre optic connection system designed to validate Paniccia's aim to 'siliconise' photonics. The prototype is the first silicon based optical data connection with integrated lasers. It consists of a silicon transmitter and a receiver chip, each integrating all the necessary building blocks. These include a hybrid silicon laser, which integrates the light emitting capabilities of indium phosphide with the light routing and low cost advantages of silicon, together with high speed optical modulators and photodetectors.

The transmitter chip is composed of four such lasers, whose light beams each travel into an optical modulator that encodes data onto them at 12.5Gbit/s. The four beams are then combined and output to a single



Paniccia: "Ultimately, achieving exascale computing is all about power. To get there, you have to connect lots of systems together and the connection itself will probably cost more than the processors and servers."

optical fibre for a total data rate of 50Gbit/s. At the other end of the link, the receiver chip separates the four optical beams and directs them into photo detectors, which convert data back into electrical signals. Both chips are assembled using familiar low cost semiconductor manufacturing techniques.


The aim is to create SP chips containing dozens – even hundreds – of hybrid silicon lasers, built using standard high volume, low cost silicon manufacturing techniques.

"The real value of SP is not about individual modulators or detectors, it's about integrating them together, as with ics," Paniccia says. "We have proved it's possible to create the building blocks; the next stage is about integration. We see this as integrating the photonic elements, but keeping the electronics separate. For the photonic elements, we need pc like assembly – high volume, pick and place, low cost packaging and efficient system level testing."

One crucial feature of SP for Paniccia is the potential for scaling. Currently, the 50Gbit/s link uses four 12.5Gbit/s channels. "But 12.5 is no magic number," he says. "Without changing anything else, we can scale the frequency of the modulators to 25Gbit/s, giving 100Gbit/s, or we can go to 40Gbit/s, giving 160. We can also 'scale out', increasing the number of channels to eight. At 12.5Gbit/s, that gives 100Gbit/s. By mixing and matching and scaling up and out, 25 lasers at 40Gbit/s reaches 1Tbit/s. I think the record so far is around 26.4Gbit/s, but that took a room full of equipment. We are talking about doing it with a chip the size of your fingernail. It's like going from the vacuum tube to the planar transistor."

"There is no way to do provide the kinds of bandwidth requirements with discrete components; it has to be integrated silicon photonics and the photonics has to be as close to the cpu as possible. But for exascale, the power consumption of processors will be enormous – 100 to 150W per processor – and photonics does not like high temperatures! Dealing with challenges like this – cooling, integration, manufacturability, testing, packaging, how to put SP into real systems – is what we are doing now. For SP, we have a path, it can be done. The question is, can we make it all work in reality?"

Another company convinced that SP is the road to exascale is IBM. At the



Intel's silicon photonics technology currently supports 50Gbit/s communications

end of last year, it announced development of what it calls CMOS Integrated Silicon Nanophotonics, again combining optical and electrical devices on chip and which can be produced on the front end of a standard CMOS manufacturing line, requiring no special tooling.

“Our CMOS Integrated Nanophotonics breakthrough promises unprecedented increases in silicon chip function and performance via ubiquitous low power optical communications between racks, modules, chips or even within a single chip,” says Dr Yuri Vlasov, manager of the Silicon Nanophotonics Department at IBM Research. “The next step in this advancement is to establishing manufacturability of this process in a commercial foundry.”


The technology means a range of SP components, like modulators, germanium photodetectors and wavelength division multiplexers, can be integrated with analogue and digital CMOS circuitry. IBM claims the density of optical and electrical integration the technology offers is unprecedented: a single transceiver channel with all accompanying optical and electrical circuitry occupies only  $0.5\text{mm}^2$  – 10 times smaller than previous devices.

“The technology is amenable for building single chip transceivers with area as small as  $4 \times 4\text{mm}$  that can receive and transmit terabits per second,” IBM says.


Another leading figure in SP is Britain’s Graham Reed, a professor of electronics engineering at Surrey University who has worked with Panicia. He sees SP as the most likely way of achieving exascale.

“Some are looking at ways to extend interconnect via other means, such as how to use novel materials for wide bandwidth electronic interconnect, but I think SP is probably now the leading contender. Power is a big issue. This is related to data rate because, with better device performance, there is usually an incremental increase in power. The usual metric is energy per bit (or power per bit/s which is the same thing). Thus, if one can increase the modulator data rate from 10Gbit/s to 40Gbit/s for little or no power increase, then the energy per bit has dropped by a factor of four.

“The main application are obviously optical interconnect, either for inter or intra chip or perhaps within data centres, and fibre to the home for high bandwidth internet, thanks to low cost transceivers. Other possibilities are lab on a chip, mid infrared applications for sensing or military and perhaps disposable biological or chemical sensors.”



Intel’s silicon photonics approach relies on two chips. The lower chip generates light and sends data, while the upper chip receives the light and converts it to the digital domain



Luxtera has recently announced a single chip 100Gbit/s optical transceiver

We will have to wait a few years before we reach the exascale, but SP is already available commercially from Luxtera. It has been shipping systems based on  $4 \times 10\text{Gbit/s}$  fibre optic transceivers for some time and recently announced the industry’s first single chip 100Gbit/s optical transceiver. The device includes four fully integrated 28Gbit/s transmit and receive channels powered from a single laser for an aggregate unencoded data rate of up to 112Gbit/s.

“Luxtera’s SP technology uses mainstream CMOS fabrication processes to deliver on-chip waveguide level modulation and photodetection along with associated electronics, resulting in a fully integrated single chip optical transceiver,” the company says. “Light from a single copackaged laser is used to power multiple optical transmitters on a chip, eliminating the need for multiple lasers and reducing transceiver cost and power consumption. This powerful combination makes SP an obvious choice for system designers over vertical cavity surface emitting lasers, providing key benefits in reliability, power consumption and signal integrity, which are critical to system design.”

The final question for the exascale era is: what are we going to do with such machines? In fact, there are lots of problems that will pose huge challenges, even for such computing behemoths: from nanoscale science to climate modelling, understanding proteins to astrophysics and many others.

Indeed, nature is so complex, we will probably never have enough computing power to master it. People in 2030 will be looking back at the exascale era with nostalgia and for them, all talk will be of the zettabyte or the yottabyte, just round the corner.



### TOP500 supercomputer list

The TOP500 list is compiled by Hans Meuer of the University of Mannheim, Erich Strohmaier and Horst Simon of NERSC/Lawrence Berkeley National Laboratory and Jack Dongarra of the University of Tennessee, Knoxville.

TOP500 lists the 500 most powerful computer systems of which details are known.

For more, go to [www.top500.org](http://www.top500.org)