Written evidence submitted by Professor Richard Penty (LEC0012)

Professor of Photonics, Head of School of Technology at University of Cambridge

Call for evidence for Low-Energy Computing

Silicon Photronics for Low Energy Computing

1. My expertise is in the field of Photonics, and not neuromorphic computing, and so I will confine myself to comments on Silicon Photonics and related areas in integrated photonics.
2. In my opinion there are three main areas where integrated photonics, of which silicon photonics is a major part, can contribute to reducing the energy burden of high performance computing
3. The first is already in existence. This is the use of photonics to provide the very high data capacity interconnect that is needed for high performance computing and data centres. Here each compute board is connected to (currently – see below) the electronic ethernet style switch multiplex to enable each server or compute resource to connect with the HPC/data centre and indeed to the outside world.
4. The interconnect capacity of modern day facilities is simply enormous and at the high end of the multiplex it is necessary to transport multi-Tb/s data. This is simply not possible using electronic means and the only possible way is using photonics over optical fibre.
5. Photonics based transceivers are key elements here. Many transceivers typically employ directly modulated lasers, as these are the cheapest form of optical source and cost is a key driver. They typically require electronics to provide e.g. spectrally efficient modulation and channel impairment mitigation. With advanced modulation techniques and laser linearisation, my group has demonstrated >100Gb/s per lane, so high performance is possible even with these simple photonic devices..
6. However for high end transceivers, where directly modulated lasers will soon no longer sufficient, Si Photonics will be very useful to integrate many components to achieve the advanced functions (e.g. wavelength de/multiplexers, coherent receivers) that will be necessary onto the same chip.
7. Currently there are companies that are targeting links with 1.6 Tbit/s capacity, with a roadmap to 3.2 Tb/s and beyond.
8. Secondly the way data centres and HPC interconnects have evolved, from the early days were the interconnect was not a significant bottleneck, is very energy inefficient. Data packets often pass through multiple electronic packet switches, undergoing twice the number of electrical<->optical, conversions as there are switches along the way. All of these conversions cost a lot of energy, as does the electrical switching. Whilst this has been known as a problem for at least 10 years, with my group amongst others advocating for optical switching to mitigate this problem, until recently there was not

sufficient motivation to move away from the conventional that servers communicated with each other within a data centre..

9. The rise of HPC and AI data centres, which are often operated by a single hyperscaler user of all of the resource unlike in conventional data centres, has given the opportunity for different interconnect architectures to be deployed. For instance Google have publicly announced that they are now using optical switching and other major players have told us privately that they are doing this.
10. These optical switches are currently MEMS based. This technology is currently the only one that can provide the large port count required to replace the large switches at the top of the data centre switching tree. They are also the most commercially developed.
11. On the other hand they are relatively slow (millisecond switching time) and so can only do functions such as optical circuit switching. They are also very expensive.
12. This is sufficient for the first generation of optical switching as it allows the data centres to switch between e.g. learning and inference modes, where the network architecture connecting the compute resources is different.
13. Second and later generations will require optical switching to become more granular and be done closer to the induvial resources, such as the GPUs and their memory banks. This switching will also likely need to be done at the packet or flow levels, and so need switching on the microsecond or even nanosecond timescale, which is not possible with MEMS.
14. Si Photonics has been shown to be able to realise space switches at the ~250x250 scale as well as space and wavelength switches. Hence these provide a means of placing optical switches much closer to the racks of servers/GPUs.
15. Most of the switch demonstrations have been realised using thermo-optic actuators, meaning micro-second switching is possible, though with electro-optic actuators nanosecond switching has been demonstrated by my group and others.
16. Finally there is a real opportunity to make use of optical computing to carry out some of the compute function optically.
17. Photonics technology is inherently analogues and doesn’t have the digital signal nature of CMOS. However, this means that it is possible to use it for analogue computing functions.
18. Whilst relatively early implementations, there have been many recent demonstrations of linear algebra processors, analogue Fourier Transform processors etc. Approaches based on these show great potential in being employed at the front end of machine learning / AI processors and if they were shown to be robust approaches, they would greatly cut the required energy consumption for AI.
19. Silicon Photonics would certainly be a key enabler of all of the above possible energy reduction approaches. It is a very successful technology and would be able to implement most of the optical function and circuits that are necessary.
20. It does suffer from some significant downsides. Firstly tit has no native light source (silicon is an indirect bandgap material) and so optical sources need to be provided

externally to the chip. The current researches into heterogenous integration would need to be successful to overcome this.

21. Integrated photonic switches in Silicon Photonics, at least at the scale that would be attractive for this application, require a large optical I/O, as well as extremely large set of electrical contacts, some of which would need to have gigahertz bandwidths for the high speed applications.

Other Opportunities

22. There are many other application opportunities for silicon photonics. These range from optical sensors to biomedical applications.
23. These really are far too numerous to list in detail but some examples are
24. Biomedical sensors which are sensitised with biomarkers so that they can detect specific chemicals and hence act as diagnostic tests
25. Integrated spectrometer which can be used in industrial monitoring applications and which can be made small enough to be included in wearable technology such as watches or in phones.
26. Beamforming technology which can be used for example in miniaturised LIDAR sensors for self-driving vehicles.

UK Positioning

27. In terms of research, the UK holds a good position. It has some excellent Si Photonics groups (Prof Reed's at Southampton is probably the longest established). My own group has worked on integrated photonics for 20+ years and about 7 years ago made a significant investment into Silicon Photonics and are now producing papers in optical switching and optical computing in very well regarded journals.
28. With the exception of Southampton, most of these groups rely on foundries to realise their designs. My own group uses commercial foundries across the world but also is in EU projects with major groups such as IMEC. We also collaborate with Cornerstone in Southampton.
29. Cornerstone has made great progress over the last 5 years or so, but does not yet have the state of the art performance that the top commercial or research foundries have.
30. The UK suffers from not having much in the way of prototype packaging line. Europe has invested heavily in this area via the Chips act and this has enabled packaging facilities, such as the Tyndall National Facility in Cork to be established.
31. There are some pockets of indigenous industry in the UK which have some Si Photonics expertise.
32. The main competitors, both in terms of university (though also industrial) research are in the USA, Japan and China, though there are some true centres excellence in Europe IMEC, LETI etc).

Possible policy interventions

33. There is certainly an opportunity for the UK to make some policy interventions to make sure that that nation remains competitive in this field.
34. It is important that the UK science base in this area is continued to be funded so that it can continue to operate at the world leading level where it currently stands.
35. In order for the UK to be able to compete on a European and world stage for integrated photonics for low energy computing, but also for many other important applications, a state of the art prototyping facility will be necessary.
36. This should include conventional Silicon Photonics capability, but to remain competitive it will also need to be able to deliver heterogenous integration.
37. Either as part or separately, a photonics chip packaging facility will be important so that the chips can be put into suitable packages to be useful for systems deployment.
38. Ideally for scale up for industry, these facilities should have pilot line capacity.

13 May 2026