Engineering 7 (2021) 1195–1196

Contents lists available at ScienceDirect

Engineering

journal homepage: www.elsevier.com/locate/eng

News & Highlights

Light-Based Chips Promise to Slash Energy Use and Increase Speed

Mitch Leslie

Senior Technology Writer

Computers already gobble a large share of the world's electrical power, and their energy demand will likely soar with the deployment of more and more power-hungry artificial intelligence (AI) systems [1]. In a step that might reduce AI's electricity use and environmental impact, Lightmatter, a startup based in Boston, MA, USA, announced that it has developed a microchip that performs calculations with light and requires about one-sixth the energy of a comparable electronic chip (Fig. 1) [2]. Other companies are developing similar photon-based chips for AI and for many other uses, including self-driving cars and quantum computing [3,4].

The energy consumption of computers has grown rapidly in recent decades. Researchers estimate that data centers now draw about 1% of the world's power [5]. Google alone uses over 12 TW·h of electricity annually, more than the country of Sri Lanka [6]. Mining for Bitcoin and other cryptocurrencies, an activity that only began in 2009, is also consuming increasing amounts of electricity [7], with the latest authoritative estimate putting Bitcoin's annual use at 121 TW·h [8]. AI is also a heavy power user [9]. In particular, training the deep learning algorithms necessary for functions such as facial recognition demands extensive data processing, which in turn requires large amounts of electricity and potentially produces large amounts of CO_2 [9]. A study estimated that training one type of deep learning algorithm takes the same amount of energy as an automobile uses in its lifetime [2].

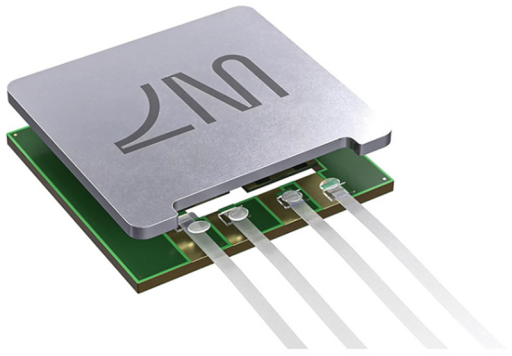

Fig. 1. The Envise photonic chip developed by Boston-based company Lightmatter is designed to speed deep learning for AI. The company says that a rack of servers that use these chips consumes dramatically less power while performing more than three times as many inferences per second as standard servers produced by a leading electronic rival. Credit: Lightmatter (public domain).

Companies have taken some steps to curb energy consumption and reduce the climate impact of computing. For instance, improved energy efficiency of data centers means that their electricity use rose only 6% between 2010 and 2018, whereas their computing power increased by six times during that period [5]. But photonic integrated circuits that operate with photons rather than electrons could yield much larger reductions.

These circuits can be so miserly because of the properties of light. Electrons run into resistance as they travel through the transistors, capacitors, and other components in a conventional integrated circuit, and the result is heat. As designers have packed more and more components onto chips, heat production has soared. It has become an obstacle to improving microchip performance [10] and is a major reason why computers consume so much energy [11]. About 40% of the power use by data centers goes toward cooling, for instance [12]. In a photonic chip, by contrast, electrical resistance is not an issue as photons generated by a laser speed through a similar array of components, including waveguides, modulators, and reflectors. Therefore, the chips generate less heat and require less power.

Photonic chips can also be much faster. Data moves within a photonic device at the speed of light, about ten times faster than electrons in a standard circuit. "There are huge gains that are allowed by the physics," said Dirk Englund, an associate professor of electrical engineering and computer science at the Massachusetts Institute of Technology in Cambridge, MA, USA. Photonic chips could increase processing speed by six or seven orders of magnitude, he said.

The idea of using light for circuits instead of electrons has a long history. Researchers began developing photonic chips in the 1960s and 1970s, and at the time some experts predicted they would follow the same trajectory of rapid miniaturization as conventional integrated circuits [13]. By 1990, AT&T Bell Labs in Murray Hills, NJ, USA, had created a prototype optical computer that relied on light to perform calculations [14]. But light-based chips never caught up to their electronic counterparts. One reason is that although engineers were able to shrink electronic components so that billions can fit on a single chip, they did not know how to do the same for optical components, said Englund. "Thirty years ago, people could not make them compact enough."

https://doi.org/10.1016/j.eng.2021.07.005

Since then, improved methods for manufacturing photonic devices have allowed the components to shrink. For instance, it is possible to make a miniaturized version of the Mach–Zehnder interferometer, a key component of Lightmatter's new chip that splits beams of light and enables the device to perform matrix multiplication [15]. Photonic chips still carry far fewer components than do electronic chips. The record is a little over 10 000, said John Bowers, a professor of electrical and computer engineering and materials at the University of California, Santa Barbara. But for more than a decade photonic integrated circuits have been incorporated into products such as the transceivers that allow communication over optical fibers.

Now, more powerful and capable light-based chips are starting to reach the market or are under development. Researchers are working on photonic integrated circuits that could crunch the numbers for the light detection and ranging (LIDAR) systems that help guide self-driving cars [3]. A photonic chip announced in 2021 by the Toronto, Canada-based company Xanadu could boost efforts to create quantum computers because, unlike competing designs, machines containing the chip do not have to be cooled to extremely low temperatures [4].

Photonic chips could be a boon for AI not just because of their speed and low energy demands. They easily perform matrix–vector multiplication, the linear algebra calculation that underpins deep learning and is difficult to perform with conventional integrated circuits, said Englund.

Several companies are developing light-based chips for AI, including Boston-based Lightelligence [16]. However, Lightmatter's chip, known as Envise, is the closest to reaching users. The company claims that the chip is up to ten times faster than its leading electronic competitor and uses only 15% as much energy, although these numbers have not been independently verified [2]. Lightmatter plans to incorporate 16 of the chips into a blade server, a specialized AI computer for data centers that it says will start reaching customers in late 2021 [2]. The device will not be purely photonic—it will contain electronic chips as well. But "it looks like what they have done is a significant advance," said Bowers.

Photonic chips might also reduce the power consumption in other areas of computing. For example, some experts have argued that the devices could curb the energy appetite of cryptocurrency mining [17]. The chips still have significant limitations, however. For one thing, creating a light-based memory is extremely difficult, said Englund. A conventional electronic chip provides Envise's memory [2]. In addition, the chips are analog, and their calculations do not have the precision of their electronic rivals [15]. For that reason, Lightmatter will sell its server blade primarily as an "accelerator" that works with AI algorithms that have already been trained [15].

Nonetheless, AI researchers say that they are excited that they will soon be able to put the light chips through their paces. "We will start to see the real benchmarking and how powerful these devices are," said Englund.

References

[1] Knight W. AI can do great things—if it doesn't burn the planet [Internet]. San Francisco: Wired; 2020 Jan 21 [cited 2021 Jun 20]. Available from: https://www.wired.com/story/ai-great-things-burn-planet/.

[2] Knight W. This chip for AI works using light, not electrons [Internet]. San Francisco: Wired; 2021 Mar 10 [cited 2021 Jun 20]. Available from: https://www.wired.com/story/chip-ai-works-using-light-not-electrons.

[3] Harris N. Light is the key to long-range, fully autonomous EVs [Internet]. San Francisco: TechCrunch; 2021 May 25 [cited 2021 Jun 20]. Available from: https://techcrunch.com/2021/05/24/light-is-the-key-to-long-range-fully-autonomous-evs/.

[4] Choi CQ. In the race to hundreds of qubits, photons may have "quantum advantage" [Internet]. New York City: IEEE Spectrum; 2021 Mar 5 [cited 2021 Jun 20]. Available from: https://spectrum.ieee.org/tech-talk/computing/hardware/race-to-hundreds-of-photonic-qubits-xanadu-scalable-photon.

[5] Lohr S. Cloud computing is not the energy hog that had been feared [Internet]. New York City: New York Times; 2020 Feb 27 [cited 2021 Jun 20]. Available from: https://www.nytimes.com/2020/02/27/technology/cloud-computing-energy-usage.html.

[6] Bruce R. How Google powers its "monopoly" with enough electricity for entire countries [Internet]. New York City: Forbes; 2020 Oct 21 [cited 2021 Jun 20]. Available from: https://www.forbes.com/sites/robertbryce/2020/10/21/googles-dominance-is-fueled-by-zambia-size-amounts-of-electricity.

[7] Leslie M. Will cryptocurrencies break the energy bank? Engineering 2020;6(5):489–90.

[8] Criddle C. Bitcoin consumes "more energy than Argentina" [Internet]. London: BBC; 2021 Feb 10 [cited 2021 Jun 20]. Available from: https://www.bbc.com/news/technology-56012952.

[9] Dickson B. AI cloud save the world if it doesn't ruin the environment first [Internet]. New York City: PC Mag; 2020 Apr 17 [cited 2021 Jun 20]. Available from: https://www.pcmag.com/news/ai-could-save-the-world-if-it-doesnt-ruin-the-environment-first.

[10] Waldrop MM. The chips are down for Moore's law. Nature 2016;530(7589):144–7.

[11] Morrison J. A new system for cooling down computers could revolutionize the pace of innovation [Internet]. Washington, DC: Smithsonian; 2020 Sep 16 [cited 2021 Jun 20]. Available from: https://www.smithsonianmag.com/science-nature/researchers-craft-revolutionary-invention-cool-computers-180975807/.

[12] Meyer D. Microsoft hails success of its undersea data center experiment—and says it could have implications on dry land, too [Internet]. New York City: Fortune; 2020 Sep 15 [cited 2021 Jun 20]. Available from: https://fortune.com/2020/09/15/microsoft-project-natick-undersea-datacenter-scotland/.

[13] Smit M, Williams K, van der Tol J. Past, present, and future of InP-based photonic integration. APL Photon 2019;4(5):050901.

[14] Markoff J. Laser device could be key to faster computer [Internet]. New York City: New York Times; 1990 Jan 30 [cited 2021 Jun 20]. Available from: https://www.nytimes.com/1990/01/30/business/laser-device-could-be-key-to-faster-computer.html.

[15] Schneider D. Deep learning at the speed of light [Internet]. New York City: IEEE Spectrum; 2021 Jan 8 [cited 2021 Jun 20]. Available from: https://spectrum.ieee.org/computing/software/deep-learning-at-the-speed-of-light.

[16] Knight W. Making AI algorithms crazy fast using chips powered by light [Internet]. Cambridge: MIT Technology Review; 2018 Nov 29 [cited 2021 Jun 20]. Available from: https://www.technologyreview.com/2018/11/29/138899/making-ai-algorithms-crazy-fast-using-chips-powered-by-light/.

[17] Can photonic chips save Bitcoin? [Internet]. Cambridge: MIT Technology Review; 2019 Nov 21 [cited 2021 Jun 20]. Available from: https://www.technologyreview.com/2019/11/21/131869/can-photonic-chips-save-bitcoin/.