



CORNING

Artificial intelligence and the impact on our data centers.

By Tony Robinson, Global Marketing Applications Manager
Corning Optical Communications

**It might help discover new levels of efficiency,
but the trade-off is a massive increase in demand
for bandwidth**

It never ceases to amaze how filmmakers are able to introduce concepts that at the time seem so far from reality, but in time

those concepts make it into our daily lives. In 1990, the Arnold Schwarzenegger movie “Total Recall,” showed us the “Johnny Cab,” a driverless vehicle that took them anywhere they wanted to go. Now, most major car companies are investing millions into bringing this technology to the masses. And thanks to “Back to the Future II,” where Marty McFly evaded

the thugs on a hoverboard, our kids are now crashing into the furniture (and each other) on something similar to what we saw back in 1989.

It was way back in 1968 (which some of us can still remember) when we were introduced to Artificial Intelligence (AI) with HAL 9000, a sentient computer on board the Discovery One Spaceship in “2001: A Space Odyssey.” HAL was capable of speech and facial recognition, natural language processing, lip reading, art appreciation, interpreting emotional behaviors, automated reasoning, and, of course, Hollywood’s favorite trick for computers, playing chess.

Fast forward to the last couple of years and you can very quickly identify where AI has become an essential part of our daily lives. You can ask your smartphone what the weather will be like at your next travel destination, your virtual assistant can play your favorite music, and your social media account will provide news updates and adverts tailored to your personal preferences. And without insulting the tech companies, this is AI 101.

But there is so much more happening in the background that we don’t see that is helping to improve, and even save, lives. Language translation, news feeds, facial recognition, more accurate diagnosis of complex diseases, and accelerated drug discovery are just some of the applications on which companies are developing and deploying AI. According to Gartner, AI-derived business value is forecast to reach \$3.9 trillion by 2022.

Thoughtful servers

So how does AI impact the data center? Well, back in 2014 Google deployed Deepmind AI (using machine learning, an application of AI) in one of their facilities. The result? They were able to consistently achieve a 40 percent reduction in the amount of energy used for cooling, which equated to a 15 percent reduction in overall PUE overhead after accounting for electrical losses and other non-cooling inefficiencies. It also produced the lowest PUE the site had ever seen. Based on these significant savings, Google looked to deploy the technology across their other sites and suggested other companies will do the same.

Facebook’s mission is to “give people the power to build community and bring the world closer together,” outlined in their white paper Applied Machine Learning at Facebook: A Datacenter Infrastructure Perspective. It describes the hardware and software infrastructure that supports machine learning at a global scale.

To give you an idea of how much computing power AI and ML needs, Andrew Ng, chief scientist at Baidu’s Silicon Valley Lab, said training one of Baidu’s Chinese speech recognition models requires not only four terabytes of training data, but also 20 exaFLOPS of compute, or 20 billion, billion math operations across the entire training cycle.

But what about our data center infrastructure? How does AI impact the design and deployment of all of the different-sized and -shaped facilities that we are looking to build, rent, or refresh to accommodate this innovative, cost-saving, and life-saving technology?

Machine learning (ML) can be run on a single machine, but thanks to the incredible amount of data throughput is typically run across multiple machines, all interlinked to ensure continuous communication during the training and data processing phases, with low latency and absolutely no interruption to service at our fingertips, screens, or audio devices. As a human race, our desire for more and more data is driving exponential growth in the amount of bandwidth required to satisfy our most simple of whims.

This bandwidth needs to be distributed within and across multiple facilities using more complex architecture designs where spine-and-leaf networks no longer cut it – we are talking about super-spine and super-leaf networks to provide a highway for all of the complex algorithmic computing to flow between different devices and ultimately back to our receptors.

Technology deployment options in the data center

This is where fiber plays such a pivotal role in ensuring your picture or video of that special (or stupid) moment is broadcast to the whole world to see, share, and comment. Fiber has become the de facto transmission media across our data center infrastructures thanks to its high speed and ultra-high-density capabilities compared to its copper cousins. As we migrate to higher network speeds, we also introduce a whole new complexity into the mix – which technology to adopt?

Traditional 3-tier networks used core, aggregate, and edge switching to connect the different servers within the data center where inter-server traffic travels in a North-South direction through the active devices to talk with each other. Now however, and greatly thanks to the high computational requirements and interdependency that AI and ML bring to the game, more of these networks are implemented using a 2-tier

spine-and-leaf network, where servers talk to each other in an East-West direction due to the ultra-low latency demanded by production and training networks.

Since the IEEE-approved of 40G and 100G back in 2010, there have been a number of competing proprietary solutions which have somewhat clouded the judgment of users who are not certain which path to follow. To explain, before 40G and the others we had SR, or short reach, for multimode and LR, or long reach, for single-mode. Both used a single pair of fibers to transmit a signal between two devices. It didn't matter whose equipment you used or which transceiver was installed in that device, it was a simple data transaction over two fibers.

But the IEEE-approved solutions at 40G and beyond, and its competing brethren changed the game. Now we are looking at either two fibers using standards-approved or proprietary, noninteroperable WDM techniques, and standards-approved, or multisource agreements (MSAs) and engineered techniques for parallel optics using eight fibers (four to transmit and four to receive) or 20 fibers (10 to transmit and 10 to receive).

- If you want to continue with standards-approved solutions and keep optics costs down because you don't need the distance capabilities of single-mode fiber, you select multimode parallel optics which also enables you to break out higher-speed 40 or 100G switch ports into smaller 10 or 25G server ports. I'll cover this in more detail a little further in this article.
- If you want to extend the life of your installed duplex fiber and don't mind staying with your preferred hardware vendor without the option of interoperability and again don't need longer distances, you select one of the multimode WDM solutions.

Now I'll tell you what the majority of the tech companies who are deploying AI on a massive scale are designing into their networks for today and tomorrow ... single-mode parallel optics. And here are three simple reasons why.

1. Cost and distance

The current market trend is that parallel optic solutions are developed and released first, with WDM solutions following suit a couple of years later, so volumes in parallel are much higher driving a lower manufacturing cost. They also support lesser distances than the 2 km and 10 km WDM solutions, so you don't need as many of the complex components to

cool the lasers and multiplex and demultiplex the signal at either end. And while we've seen the size and scale of these "hyperscale" facilities explode into buildings the size of 3-4 football pitches within huge campuses, our own data shows that the average deployed length over single-mode fiber is yet to exceed 165 m in these facilities, so there is no need to pay for a more expensive WDM transceiver to drive a distance that they don't need to support.

Parallel single-mode also uses less power than a WDM variant. As we saw from the Google example earlier regarding their power usage, anything that can be done to reduce the single biggest operating cost of a data center has to be a good thing.

2. Flexibility

One of the major advantages of deploying parallel optics is the ability to take a high-speed switch port, say 40G, and break that out into 4x10G server ports. Port breakout offers great economies of scale because breaking out into lower speed ports can significantly reduce the number of chassis or rack-mount units for the electronics from 3:1 (and data center real estate is not cheap) and uses less power, which requires less cooling so lowers the energy bill further, with our data showing this to equate to a 30 percent savings on a single-mode solution. The transceiver vendors also confirm that a large proportion of all shipped parallel optic transceivers are deployed to take advantage of this port breakout capability.

3. Simple and clear migration

The technology road map of the major switch and transceiver vendors shows a very clear and simple migration path for customers who deploy parallel optics. I mentioned earlier that the majority of the tech companies have followed this route, so when the optics are available and they migrate from 100G to either 200 or 400G, their fiber infrastructure remains in place with zero upgrades required. Those companies who decide to stay with a duplex, 2-fiber infrastructure may find themselves wanting to upgrade beyond 100G, but the WDM optics may not be available within the time frame of their migration plans.

Impact on data center design

From a connectivity perspective, these networks are heavily meshed fiber infrastructures to ensure that no one server is more than two network hops from each other. But such is the bandwidth demand that even the traditional 3:1 oversubscription ratio from the spine switch to the leaf switch is not sufficient and

is more typically used for distributed computing from the super spines between the different data halls.

Thanks to the significant increase in switch IO speeds, network operators are striving for better utilization, higher efficiencies, and the ultra-low latency we mentioned by designing their systems using a 1:1 subscription ratio from spine to leaf, an expensive but necessary requirement in today's AI crunching environment.

Additionally, we have another shift away from the traditional data center design following the recent announcement from Google of their latest AI hardware, a customized ASIC called Tensor Processing Unit (TPU 3.0) which, in their giant pod design, will be eight times more powerful than last year's TPUs with over 100 petaflops. But packing even more computing power into the silicon will also increase the amount of energy to drive it, and therefore the amount of heat, which is why the same announcement said they are shifting to liquid cooling to the chip because the heat generated by TPU 3.0 has exceeded the limits of its previous data center cooling solutions.

In conclusion

AI is the next wave of business innovation. The advantages it brings from operational cost savings, additional revenue streams, simplified customer interaction, and much more efficient, data-driven ways of working are just too attractive – not just to your CFO and shareholders but also your customers. This was confirmed in a recent panel discussion when the moderator talked about websites using Chatbots and claimed that if it wasn't efficient and customer focused enough he would drop the conversation and that company would never receive his business again.

So we have to embrace the technology and use it to our advantage, which also means adopting a different way of thinking about data center design and implementation.

Thanks to the significant increase in performance at the ASICs we will ultimately see an increase in the IO speeds driving connectivity even deeper. Your data centers will need to be super-efficient, highly fiber-meshed, ultra-low latency, East-West spine-and-leaf networks that accommodate your day-to-day production traffic while supporting ML training in parallel, which conveniently brings me to wrap this up.

AI is the next wave of business innovation. The advantages it brings from operational cost savings, additional revenue streams, simplified customer interaction, and much more efficient, data-driven ways of working are just too attractive – not just to your CFO and shareholders but also your customers.

We have seen how the major tech companies have embraced AI and how deploying parallel single-mode has helped them to achieve greater capital and operational costs over traditional duplex methods, which promise lower costs from day one. But operating a data center starts at day two and continues to evolve as our habits and ways of interacting personally and professionally continue to change, increase in speed, and add further complexity. Installing the right cabling infrastructure solution now will enable your business to reap greater financial benefits from the outset, retain and attract more customers, and give your facility the flexibility to flourish no matter what demands are placed on it.

The logo for Corning, featuring the word "CORNING" in a white, serif, all-caps font centered on a dark blue rectangular background.

Tony Robinson is a global marketing applications manager for data center segment with Corning Optical Communications (www.corning.com/opcomm).

Corning Optical Communications LLC • PO Box 489 • Hickory, NC 28603-0489 USA
800-743-2675 • FAX: 828-325-5060 • International: +1-828-901-5000 • www.corning.com/opcomm

Corning Optical Communications reserves the right to improve, enhance, and modify the features and specifications of Corning Optical Communications products without prior notification. A complete listing of the trademarks of Corning Optical Communications is available at www.corning.com/opcomm/ trademarks. All other trademarks are the properties of their respective owners. Corning Optical Communications is ISO 9001 certified.
© 2019 Corning Optical Communications. All rights reserved. LAN-2549-AEN / August 2019