How Fiber Optics Painted Itself into a Corner

and How Electro-Optic Polymers Can Get It Out

Fiber optic communications is the lifeblood of our information economy. As the amount of data that the world relies on continues to grow inexorably, the speed of fiber connections has reached astounding rates. Back in the late 1990s, only the big data pipes of the "information superhighway" national infrastructure relied on fiber optics that ran at 1 or 2 gigabits per second. Today, the consumer at the edge of the network is being offered gigabit broadband service; tomorrow, the next big cellular upgrade to 5G promises "Gigabit (per second) Internet" to your phone. Meanwhile, the big pipes that aggregate entire cities' worth of traffic to send to the giant datacenters for processing have grown to 100's of gigabits per second -- and are still expected to grow significantly over the next decade. Fiber optics, in many forms, permeate the network from residential broadband to intra-datacenter to national and international backbones.

But can we count on speeds to continue increasing into the future?

There are no guarantees. The end of electronics' Moore's Law is widely discussed. Similarly, leading fiber optic technologists are giving papers about approaching fundamental limits. It does not mean progress will stop, but it does mean continued progress will require more ingenuity (and complexity) than before. In this white paper, we will look at a particularly rapidly developing segment of fiber optics, the type that carry traffic in and between giant datacenters. We will see how it has not been straightforward to get to where we are today and why a new direction is needed now. We will also discuss why polymer EO modulators could be key to our continued access to more data.

Both electronics and fiber optics have already been forced to rely on some clever innovations to get us to where we are today. For electronics, multi-core processors, FinFET transistor structures, and specialty processors are some of the ways that the industry has circumvented the difficulty of continuing to shrink transistors. In the past two decades, optoelectronic devices have increased their operating speeds dramatically—but not as much as data rates have increased!

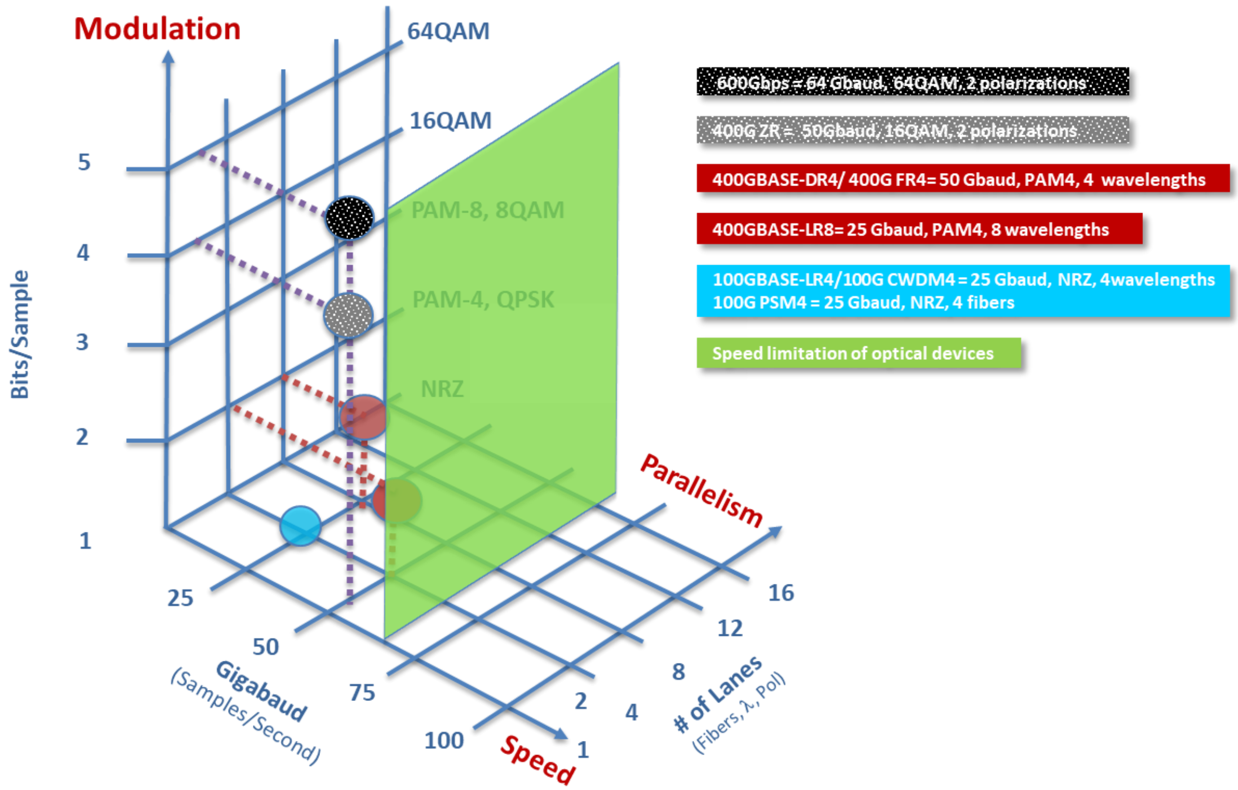

Figure 1 Limited speed has forced the use of complex multi-dimensional transmission schemes

Several popular current and new transmission formats are shown in Figure 1. Notice how complicated and constrained the optics industry has become. This is a result of being constrained by optical device speed to keep to the left side of the green plane.

Figure 2 expands on how and why the industry arrived at the condition shown in Figure 1. When the industry went from 1 Gbps to 10 Gbps, around 1999-2000, progress was a relatively simple matter of developing faster transmitter and receiver devices. It was like pushing down on the accelerator in a car and going faster automatically. Device engineers made lasers go faster, and fiber optic links sped up from 1 Gbps to 10Gbps. Figure 2 shows where the industry started at the yellow circle representing 1 Gbps, then moved along the orange arrow to the orange circle representing 10 Gbps.

Optoelectronics only managed to keep pace so far by using adding complexity

---

But to go from 10 Gbps to 40 Gbps--after an abortive attempt at a simple speed increase in the early 2000's--the industry was forced to use four parallel lanes each at 10 Gbps. This new direction is represented by the purple arrow in Figure 2 leading to the purple circle representing 40 Gbps of aggregated data. In some cases, these lanes are separate fibers. In other cases, the lanes are wavelengths carried in the same physical fiber but staying separate from each other (wavelength division multiplexing or WDM).

By 2015, millions of units of 40 Gbps optics in the form of 4 x 10 Gbps transceivers populated the biggest datacenters. This bought time for the device speeds to progress to 25 Gbps, enabling 100 Gbps still using 4 parallel lanes. Figure 2 shows the resulting progress as the blue arrow that leads to the blue circle which is 4 lanes x 25 Gbps = 100 Gbps. 100 Gbps using these parallel lanes is where the industry is today with formats such as 100GBASE-LR4 standardized by IEEE, and industry published multi-source agreement specifications for 100G PSM4 and 100G CWDM4.

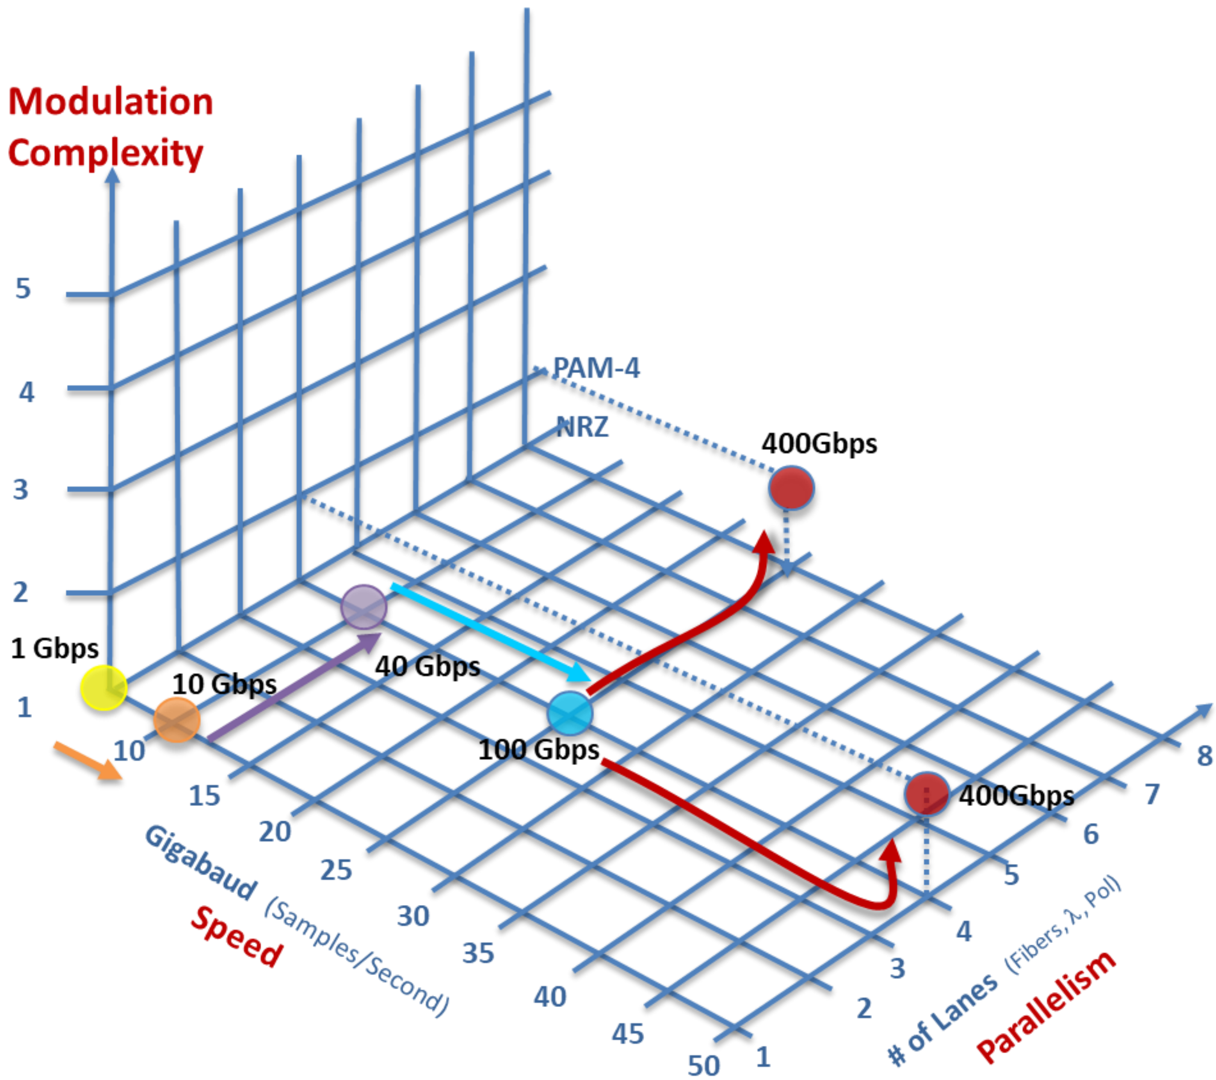

Figure 2 The long and winding roadmap to achieve aggregated 100 and 400 Gbps

(If you are really paying attention, you may have noticed that the speed axis is labelled Gigabaud rather than Gigabits per second. Don't worry for now about the difference between Gigabaud and Gigabits, we'll get to that below.)

Since 2015, appetite for more data is still increasing, and this has challenged the industry to consider an aggregated speed of 400 Gbps. Currently the industry is struggling with yet another directional change as it advances to 400 Gbps. Two popular routes shown in Figure 2 as red arrows both required adding a vertical dimension to the graph. One route utilizes 8 lanes, and the other route utilizes 4 lanes of data. Simple math says that 8 lanes of 25 Gbps only does not get to 400 Gbps: in fact, 4 x 25 = 200 Gbps. So, to double up the speed, there has been another innovation: the vertical axis shows the use of advanced signaling or modulation which we explain next.

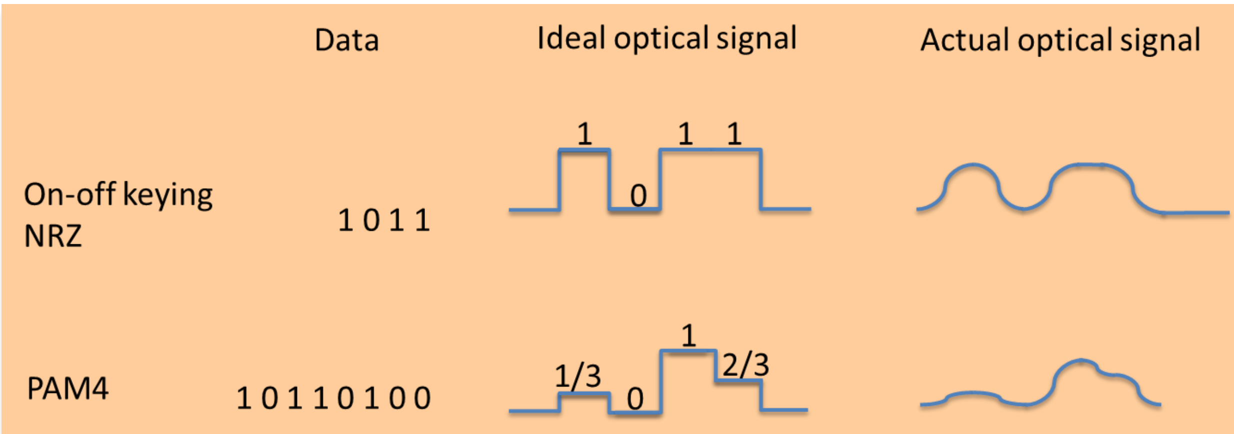

(1) NRZ (non-return to zero) or sometimes referred to as OOK (on-off-keying) is simply 2 levels: 1 and 0
(2) PAM4 (Pulse Amplitude Modulation) uses 4 levels and has the effect of doubling of the data rate
Figure 3 Comparing data modulation schemes

The top of Figure 3 shows a data stream of 1's and 0's sent simply as "light" and "no light"—as scheme called NRZ (non-return to zero) or OOK (on-off-keying). It also shows how a transmitter with limited speed blurs the data. But if data is encoded to be sent as multiple levels, more than one bit can be sent in one symbol period, or a faster bitrate (bits/second) than baudrate (symbols/sec).

The lower part of Figure 3 shows twice as many bits sent using advanced modulation with the same transmitter speed as before. The particular coding illustrated is PAM4 (pulse amplitude modulation, 4 levels) where each pairs of bits are converted to a single symbol. ('10' is mapped to 1/3 height signal, '11' is mapped to 0 height signal, etc.) This innovation comes at the cost of quite a lot of complexity in the form of signal processing electronics.

Back to Figure 2, both ways shown for 400 Gbps rely on this factor of 2 multiplication in bit rate from using PAM4. The two choices are 4 lanes of 50 Gbaud PAM4 (4 lanes x 50 Gbaud x 2 bits/symbol is used for IEEE standard 400GBASE-DR4 and 100G Lambda MSA's specification for 400G-FR4) or 8 lanes of 25 Gbaud with PAM4 (8 lanes x 25 Gbaud x 2 bits/symbol is used in IEEE standard 400GBASE-LR8).

So far, we have been discussing the history of short, typically Ethernet carrying, fiber links such as those found inside enterprise networks and data centers. Figure 2 maps the history of

Ethernet optics from 1 Gbps to 400 Gbps. Bigger pipes which aggregate multiple streams of traffic for transmission over longer distances such as between datacenters, have been forced to use even more complex schemes not discussed here. The gray and black textured bubbles in Figure 1 illustrate a few of these additional higher order vertical axis modulation schemes used for these longer distances as having a multiplying factor of up to 5 bits/symbol!

Cost, size and power consumption are pushing back on additional increases in parallelism and modulation complexity

---

These directional changes have not been as straightforward as represented. Why didn't the industry simply continue adding more parallel channels? Additional real-world requirements on cost, size, error rate and electrical power consumption must also be met. The industry struggled to accept the complexity and cost due to multiple lasers and receivers but ultimately the appetite for data outreached technological progress in faster optoelectronic devices. These considerations are already limiting the degree that parallelism and modulation complexity can be used. It gets even harder the further one goes along each axis!

Recall that two different solutions for 400G are shown in Figure 2. One reason that multiple solutions exist is discomfort with the use of 8 lanes. Parallelism beyond 4 lanes has not proved economical or small enough previously. But the industry is even more leery of increasing the modulation complexity. Complex modulation schemes are more sensitive to noise and are more error-prone. Notice that in Figure 3 the difference in PAM4 levels is only 1/3 has much as for NRZ. As a result, the optoelectronics on both transmit and receive ends must distinguish between these smaller differences in signal level. They require power-hungry signal processing electronics not only to encode and decode the data but to offset these errors. The industry already has experience with very complex modulations schemes (e.g. the gray and black circles in Figure 1) and deems them too large, too power-hungry, and way too expensive—in other words, practical only for national backbone networks. Moreover, they are approaching fundamental limits.

The industry has painted itself into a corner…literally! What is beyond 50Gbaud?

The whole industry today (epitaxial vendors, foundries, various chip suppliers, standards organization, test equipment vendors, package houses) are gearing up for 50Gbaud with some players are stretching their frontier of performance to 64 Gbaud. The industry is wondering where to go next. The customers - both giant data center companies and the telecommunication companies are expecting data rates to continue to move quickly beyond 400Gbps and on to 800Gbps and even 1600Gbps.

The obvious, yet most difficult next move, is to revisit increasing the optoelectronic device speed. At the same time, these new optoelectronic devices must be very small, and operate with very low voltage to keep power consumption low. Today's roadmaps are expressing desire but at the same time doubts about getting beyond 50 or 64 Gbaud based on incumbent technologies.

Electro-optic polymers are positioned to extend the roadmap past 100 Gbaud

Unlike conventional modulator materials, the polymer material system is naturally fast. Lightwave Logic, Inc. is designing commercial high speed optical modulators made from electro-optic polymers that will be capable of 100 Gbaud. Technical data showing 130 GHz (corresponding to 150 Gbaud) indicates even higher speeds should be possible in future.[i] The polymer material system also has the potential of extremely low voltage levels on the order of 1V as well as small device size.

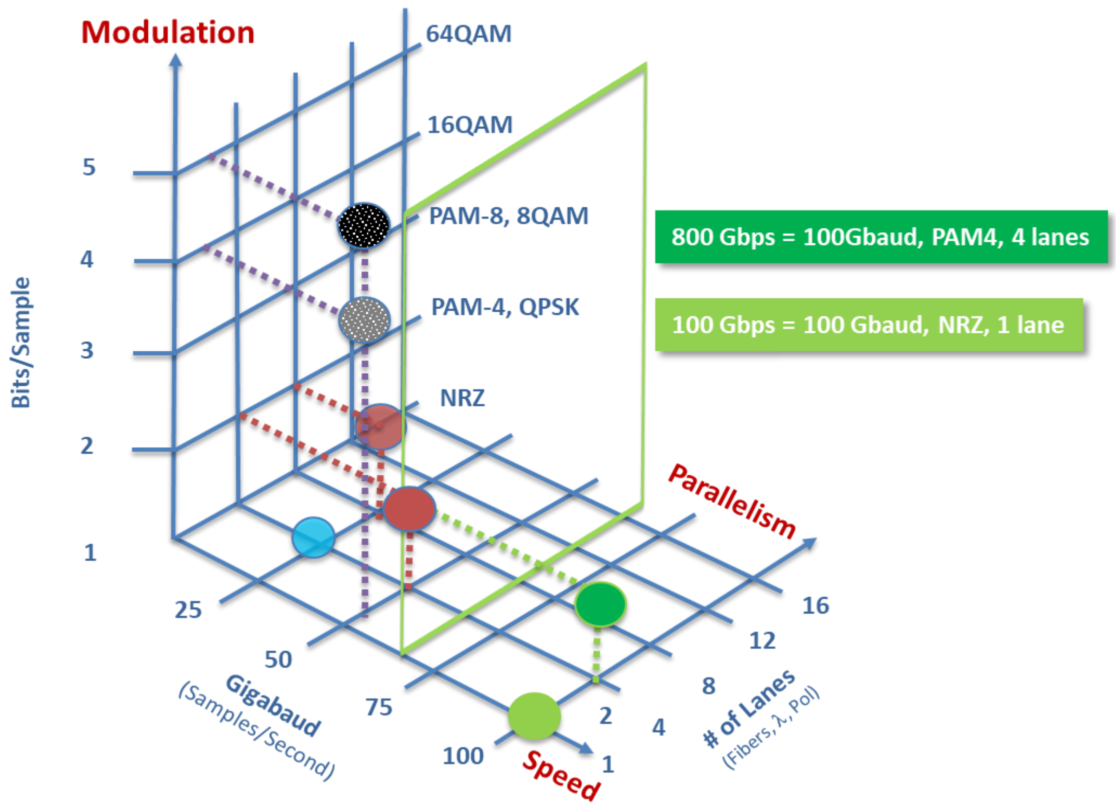

Figure 4 Devices at 100 Gbaud restore access to the speed dimension

The polymer material breaks through the technology barrier and opens up new possibilities. Two examples of solutions that could inhabit the new design space are shown in green in Figure 4. We believe it is ideal to extend the industry roadmap to 100Gbaud and beyond.

[i] M. Lebby. [Online]. Available: http://lightwavelogic.com/external.asp?b=2252&from=dl&ID=175125