Electro-Optical PCB Technology for the Future

Istvan Nagy, 2014, BUENOS@OPENCORES.ORG

Introduction

This article is about idea capture, from the point of view of a board hardware designer. I admit it is a lot of speculations, prototypes were not built, but we have to start somewhere. Signalling speeds over long PCB traces (like backplanes or large boards) can be pushed up to about 28Gbps, while on short ones (close chip-to-chip) up to about 50Gbps. This is what the glass-fibre based cheap-enough PCB material allows, due to losses and fibre wave skew. Today most high-end designs use 10.3Gbit signalling, and transitioning to 28gig in 2013-2015. PAM3 and PAM4 modulations can provide increasing the data bandwidth to another 2x, like it is used in the 100GBASE-KP4 backplane Ethernet. Increasing number of lanes and board/connector density can also provide another 2x-4x advantage at higher cost. These ideas can provide one more generation upgrade to current systems. After this, we don't have any ready-to-use technologies for increasing interconnect bandwidth. We will have to "go optical" on the board. Looking back to the past, the on-board interconnect speeds were increasing rapidly. For example PCI-express speeds were doubling every four years, today it is 8Gbit (instead of 10Gbps), and so around 2026 it would be 64Gbit/s. Backplane and on-board Ethernet speeds were increasing similarly, 8 times in 11 years from 3.125Gbps XAUI to 25Gbps 100GBASE-KR4. Another 11 years, around in 2025 it would be 200Gbps on each lane. These would definitely need optical connections. When using electrical interconnects on PCBs the practically usable transmission distance is shrinking with increasing signalling speed. The causes are well known within the domain of signal integrity, like losses, reflections, crosstalk, intra-pair skew, mode transformation and inter symbol interference. If we used optical interconnect on the boards, then most of these problems would either go away or would get reduced significantly. Distance vs speed will not be a problem anymore.

The idea is to develop the technology over 3 generations. In the first generation we would use traditional components, electro-optical PCBs and Electro-optical transceiver chips. At later generations the transceiver chips will disappear, making the main processing silicon electro-optical. To develop the technology, cooperation between silicon chip vendors, PCB fabricators, PCB material manufacturers, soldering material manufacturers, PCB assembly services companies and board/system developers.

Currently proposed solutions by Altera, Reflex Photonics, Avago, Zarlink/Microsemi and others include chip-to-cable connections from transceiver modules with inside-box cabling, while the PCB itself is not containing any printed optical waveguides. Compass-EOS has already built a switch chip with a big opening for optical ports on the bottom center of the chip with unknown connectivity and a router system/chassis with it.

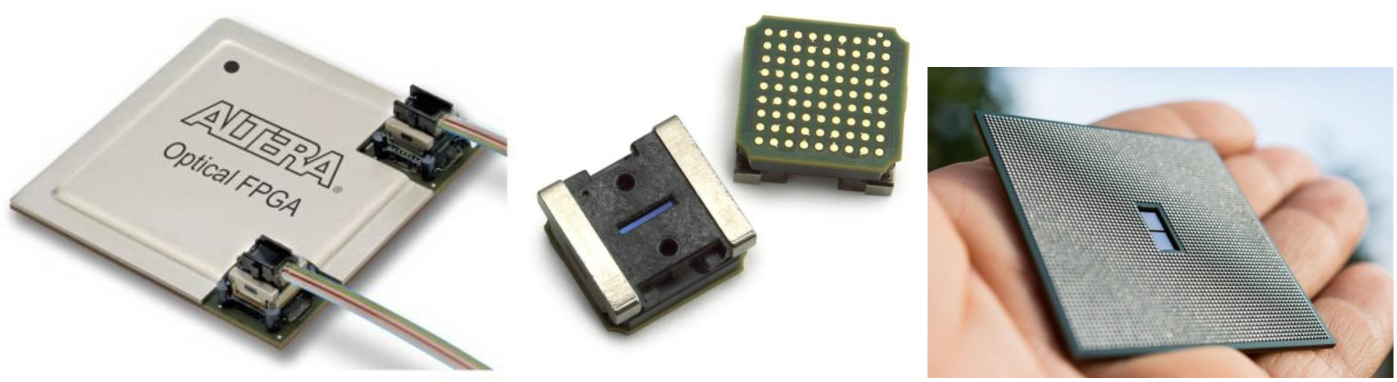

Altera's chip-to-cable solution, the Avago MicroPOD embedded transceiver, and the Compass-EOS switch chip

High-speed connections can be any of 3 types:

1. On-board chip-to-chip.
2. Board-to-board (daughter card or backplane) inside chassis.
3. Chassis to chassis: from pluggable transceiver at the face plate, or from a chip inside through the faceplate.

Some currently proposed solutions (e.g. Intel's external PCIe) only target type-3, by extending the external connection inside the chassis to the interface chip. They only solve one problem which is the chip-to-xFPx PCB trace signal integrity. The way they are constructed (using cables), they will not be suitable for chip-to-chip on the same board or for backplane based connectivity; simply because of the number of connections required cannot be made using cables. In high-complexity non-consumer products it is a much higher number of connections per unit, so high that cable-based solutions would be impractical or physically impossible. For example 75 pieces of quad optical cables inside a box would be a nightmare to connect up, and it would be extremely unreliable, looking not like a bowl of spaghetti but like a large box filled with spaghetti. Not mentioning that the large quantity of cables would fill up the internal volume so much that the air cooling of chips inside the box would be impossible. Possibly the cables would not even fit inside the chassis even with removing all electronics. The other major problem is that the connections required inside a chassis are not only between two points (like the mentioned solutions assume), but between many points, so increasing the bandwidth (speed, or using WDM) will not reduce the number of connections enough to fit in the box. I am not talking about desktop PCs or web servers where you really need to connect one chassis with one other chassis, or where the chassis contains only one main component (e.g. a processor). Unfortunately most currently proposed solutions are aiming for these only, and it is not scalable. Many data center boxes utilize multiple processors, processing ASICs, NIC-ASICs, gearbox chips, FPGAs… and they all need high-bandwidth connectivity with more than one partner.

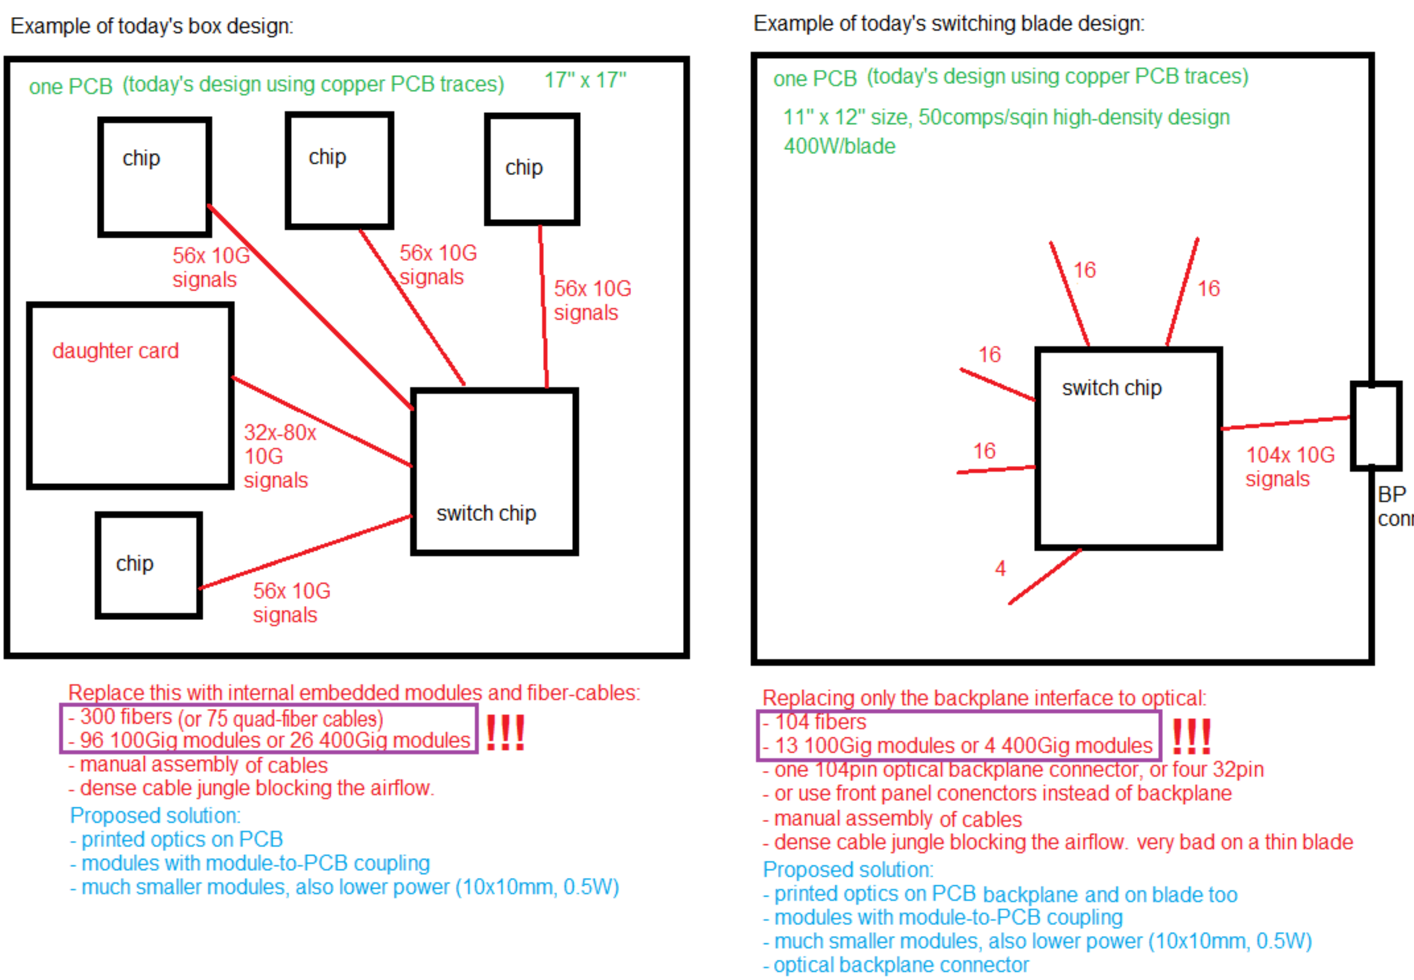

Examples of today's complex designs

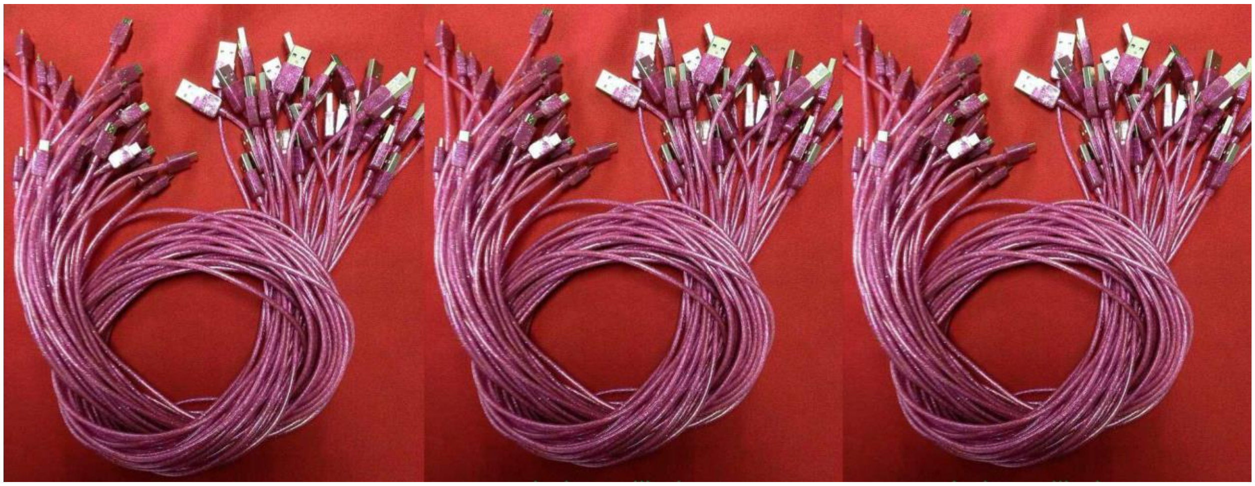

This is how 75 cables look like. Fancy trying to connect all INSIDE a 3U chassis?

It should be clear that intra-chassis cabling (as proposed by Intel, Altera, Avago, Reflex Photonics…) would only work for type-3 connectivity above (simple servers and consumer electronics), but if we want to cover all 3 cases then we need something different, we need printed optics and transceivers with optical coupling (optical pins) to the PCB.
The traditional PCB technology is very flexible and widely available; all types of electronics products can be built with it (low-end to high-end), easily and cheaply. Moving forward we would need a mixed optical-electrical technology that is just as flexible and cheap. This could only be achieved and mass produced with printed optical waveguides and mixed-material soldering on the PCB itself. This article is mainly about this route.

In an article TTI's R&D shown that they implemented printed optics with on-board electro-optical transceivers, but using embedded micro mirrors instead of pins and soldering to couple the light from the component into the trace. The mirrors were made from an inner layer optical trace with carving a V-shape into them forming a 45degree mirror. This requires very high precision and high cost manufacturing. http://www.magazines007.com/pdf/TTM-Immonen08-11.pdf
This might be the way to go, but I think it might not be suitable for cost sensitive and high volume production. Instead, something that resembles today's electrical soldering is proposed in this article.

Components

The electro-optical transceiver chip:

It has to be very small, as we may have to be able to place 128 pieces of them around a 2"x2" Ethernet switch chip. The power dissipation also has to be low enough, so the small package can dissipate it, although we may have to include the transceiver-ring in the main chip's heat sink design. The key would be to use BGA components with both electrical and optical solder balls. The optical BGA balls may be pegs instead of balls, reducing the diameter to reduce dispersion. Soldering a tin-silver peg would be difficult, but a thin polymer peg with high viscosity could retain contact during reflow soldering. A typical 0.5mm BGA ball has a diameter which is half the wavelength of a 200GHz signal (remember, year 2025 for 200Gbps connections). The solder pegs are really straight pins that are soldered or glued on the surface, rather than inside a hole. The chip would probably be a silicon-photonics device integrating laser diode, laser driver block, photodiode, trans-impedance amplifier, electrical transceiver with CDR DFE/FFE and PLL. Basically what we have in an SFP+ module, all that would have to be integrated on a single chip.

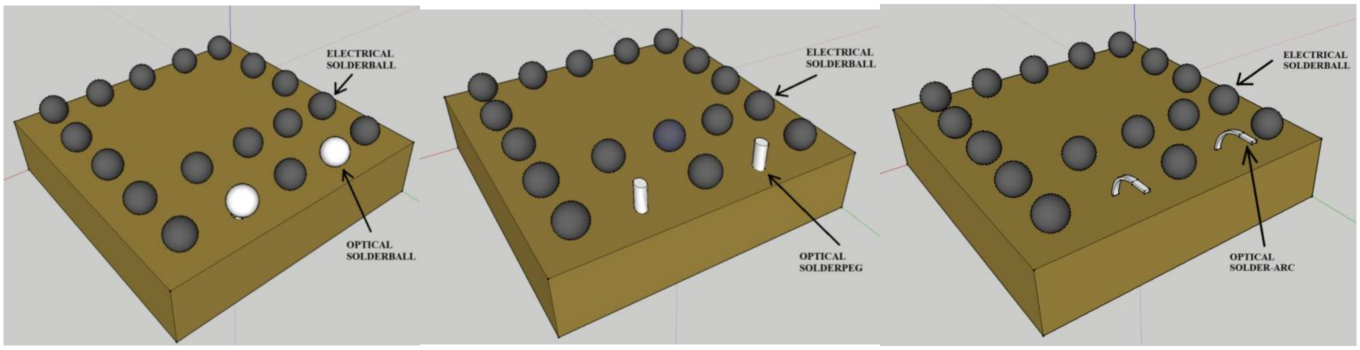

The Electro-optical transceiver chip with different optical connections

An easy initial solution could be based on existing chip-to-cable transceiver technology, for example the Avago MicroPOD optical module or the Reflex Photonics LightABLE optical engine. They would require a custom-made optical plug that attaches to the existing module instead of a cable, and has the optical solder pins/pegs. We attach this to the MicroPOD or LightAble module (instead of the original fibre/cable) and treat the duo as a new component that can be soldered. The chip development can be reduced to an optical plug development, this way the technology can be made available in the very near future. Basically the required solder able mixed-technology transceiver is already 90% existing.

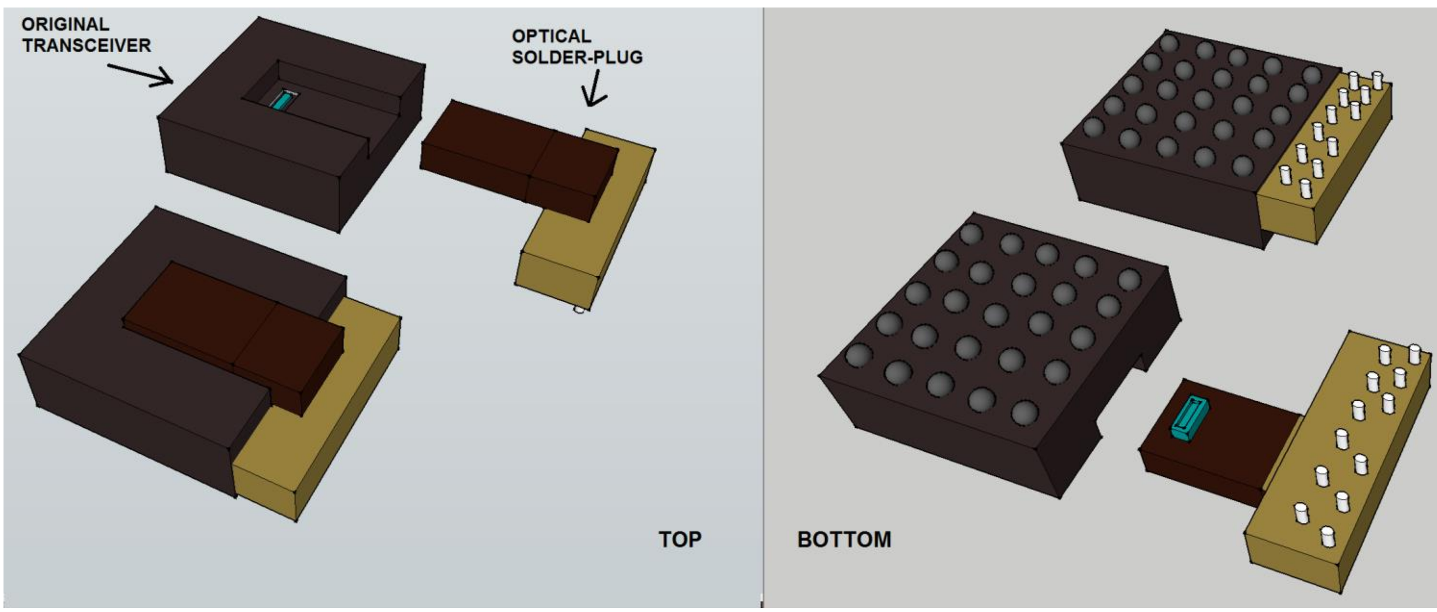

MicroPOD or LightAble module with optical solder plug

High-density optical board-to-board connectors:

This is for daughter cards and for backplane-based blade systems. The connector would be soldered or glued on both the blade and the backplane using optical surface mount pins/pegs. IT has to be surface mount; otherwise the through-hole drilled pins would create a stub that could not even be reduced enough by back drilling. Back drilling leaves a small stub, around 0.2mm, which can be too much at higher data rates. If we use (well, we have to, to reduce stubs) surface mount connectors, then we need to include mechanical reinforcement in the connector designs, for example screws.

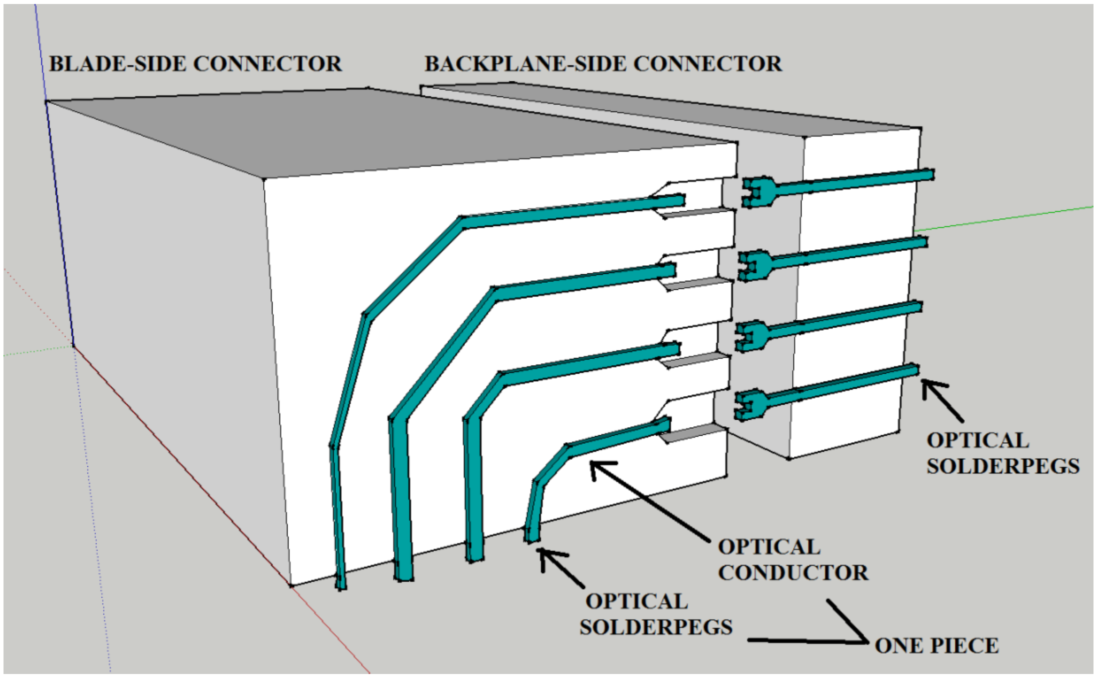

Backplane connector

External interfaces:

The external interfaces would be similar module-based as they are now, but with a few differences. Naturally the size and power dissipation shrinks, and the socket connector to the board interface would be optical not electrical. These sockets could have dual use, either plug in a passive optical-to-optical connector with a fiber cable for short range connections inside the same datacenter, or plug in an actual module with active electronics that has a repeater and signal processing for long range connections.

Boards

Printed circuit board (PCB) technology with optical interconnects is under research. Companies like IBM and Dow Corning Polymer have made optical waveguides on top of PCBs using lithography. See “Polymer Waveguide Silicones” in PCD&F January 2014. The optical trace width and thickness would have to shrink over time as data rate goes up, then dispersion is becoming a problem. We will also need to provide optical isolation to reduce crosstalk. This can be done by using a copper base under the optical trace, and pitch-black solder mask on the top.

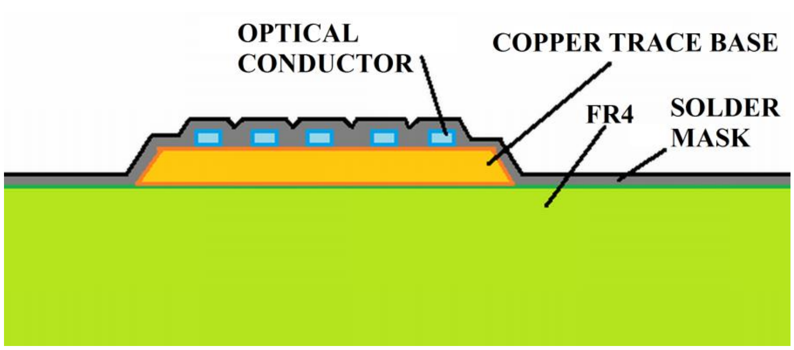

Trace cross section

There are however missing links for building an assembled board. For example reflow soldering of optical pins. The key would be to use BGA components with both electrical and optical solder balls. The electrical balls would be the same as today, Tin-Silver-Copper. For the optical balls we need an optically clear material that melts around the same temperature as the lead free solder, at around 217 ˚C. We would still use the stencil based solder paste application for the electrical pins, but for the optical balls we either use higher viscosity transparent polymer, or we dispense the transparent solder paste one by one using a robot, or we apply it using lithography as the last step of the PCB fabrication process. If we use an optical solder paste, then it would be the same material in liquefied form (using a diluent at room temperature) as the chip's optical solder balls. After soldering, the chips would require "black" under fill, to prevent crosstalk between the nearby pins. It would be practical to put the optical pins on the outer row of the chip, to make the under fill process easier.

It might also be a good idea not to solder the optical interconnect, but after the electrical solder stencilling we dispense high viscosity glue on the optical pads one by one, go through the pick&place machine, cure the glue with UV light or heat, then solder the electrical pins in reflow. This requires the optical balls/pegs to have higher melting point than the reflow temperature. Since the FR4 PCB material conducts light, it could create crosstalk between optical traces. To prevent it, we would need to create a copper basis for the optical interconnects, and use pitch black solder mask on the top.

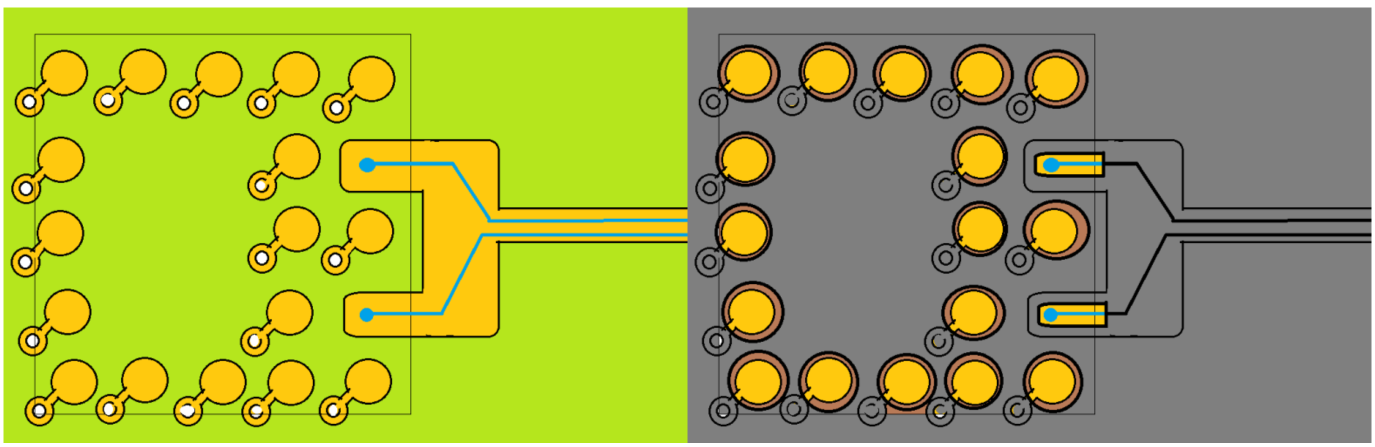

Board surface without and with black solder mask

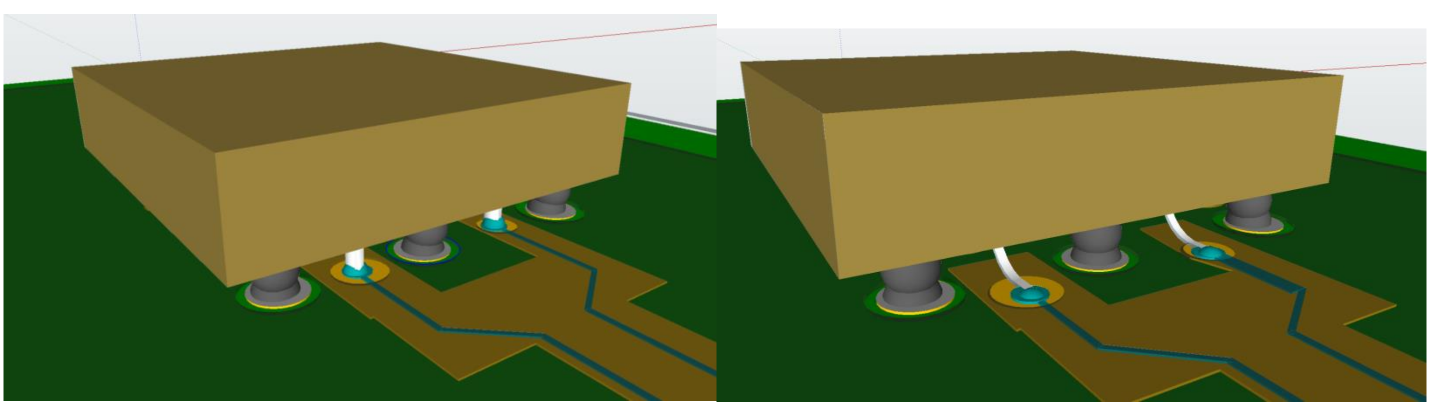

Soldered chip

At the point where the optical solder peg is soldered or glued to the PCB the connection is 90 degrees that might create a problem for guiding the light into the printed optical waveguide. Possible solutions would be either to use a non-clear scattering white glue or solder with pitch-black under fill, or use clear solder/glue with under fill that looks like chrome and creates mirror walls on the solder peg and on the glue/solder. If we use scattering glue then every soldering point or component pin would create some amount of loss on the signal that might or might not be acceptable.

Non-BGA (gullwing) approach:

For the optical pins we could use gullwing pins like those on the traditional QFP and SOP packages. This could allow more flexibility for the manufacturing process, and it would also better guide the light beams into direction at the point of contact to the PCB. It still needs reflective or black under fill, like the under-the-package approaches mentioned earlier (ball, peg, and arc). The electrical pins could be BGA balls as it is used today for very complex and high-performance devices. Instead of vertically connecting to (traditional) pads as attachment point, we could use an edge-interface to attach the gullwing pin to the optical trace. This means gradually widening the trace to the attachment point, and gluing it to the pin side by side. The pin has to be thin at the tip, nearly as thin as the optical trace. This further improves the directional attachment efficiency. With this method we could even make the optical connections with glue after the electrical soldering.

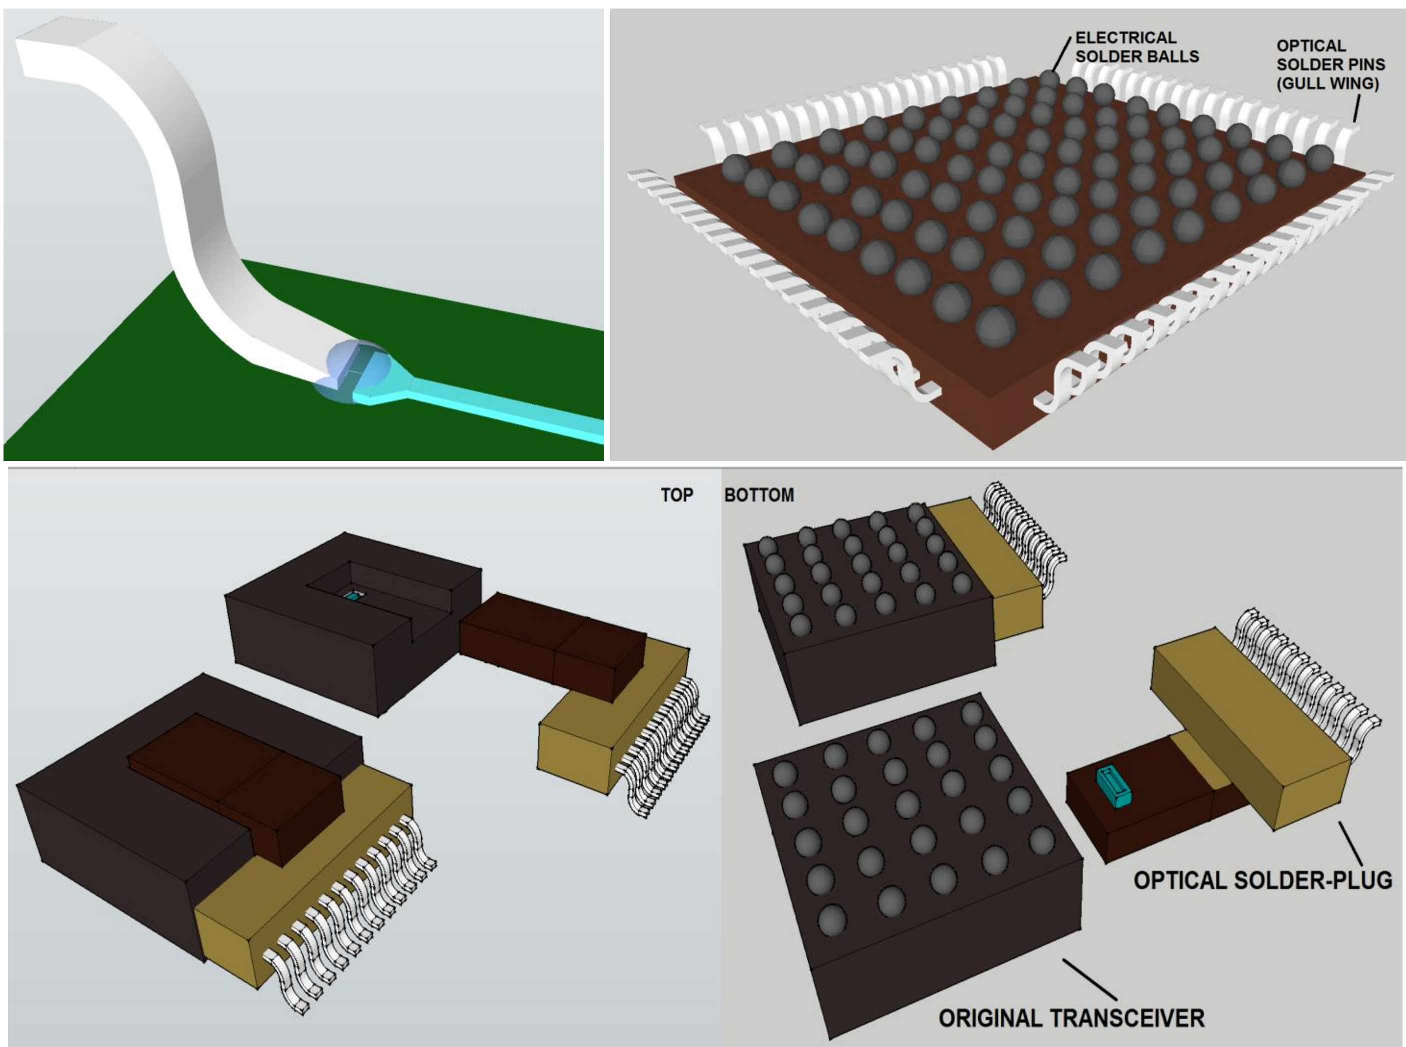

Gullwing options for the optical pins only

The gullwing pin would be much thicker (200um) than the optical trace (25um-50um), so the light would have to be funnelled in between, by the glue and the under fill. The idea is that the amount of glue and the glue-trace wetting has to be such that the glue blob takes up a certain shape, with low angle on the sides and flattening out on the edges. Important to note that the optical pin, the optical trace and the glue would together form an optical conductor. Since the trace is so thin, we might or might not create a "pad" area on the optical layer. After applying the glue they will all fuse into one piece of optical connection anyway.

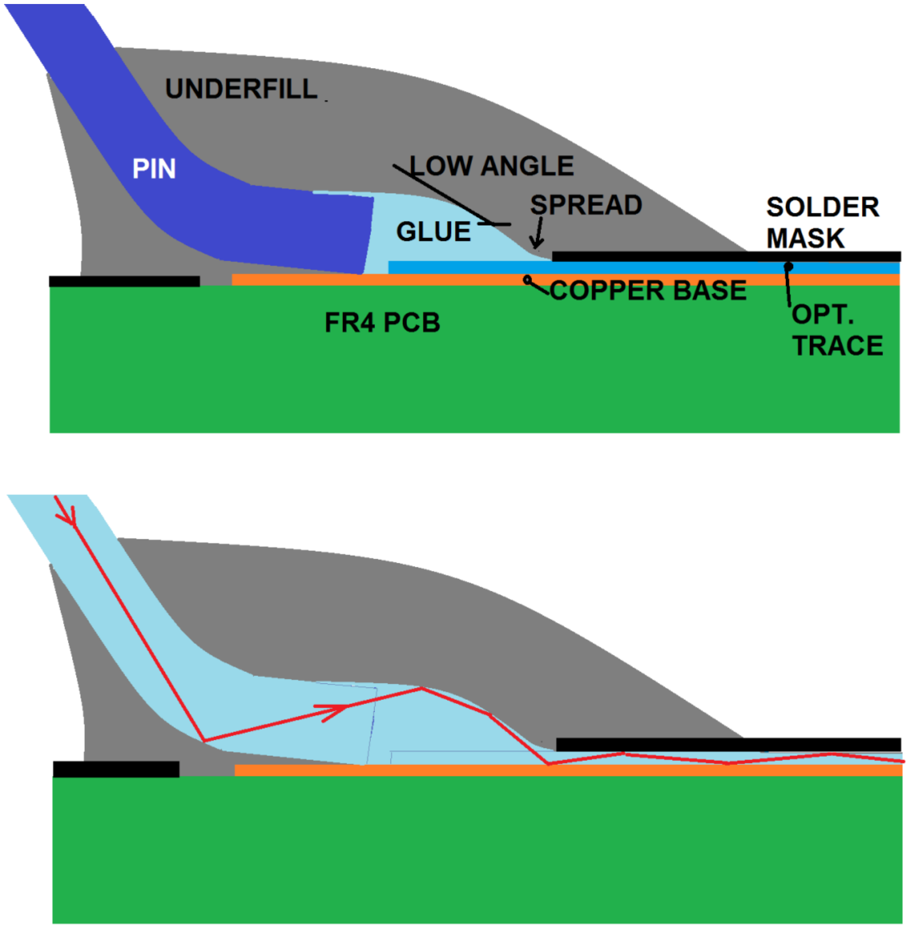

Gullwing to trace funnelling

Component to PCB attachment (coupling) methods summary:

It is still to be determined whether we need pitch-black or normal solder mask for crosstalk isolation and for providing reflective side walls for the optical trace/waveguide. The same goes for the black or reflective under fill that would cover the pin/trace attachment area. Also to be determined whether we need to use clear glue, or milk-like white glue that scatters the light for directional coupling at the 90 degrees angled pin to PCB attachment. It should be a lot easier and cheaper (time spent per board and equipment investment) to dispense glue on the PCB surface, rather than manufacturing high-precision micro mirrors on the inner layers.

Right angle coupling methods		
Method	Pro	Con
Bent pin (Gullwing or solder-arc) with clear glue	Easy/cheap to manufacture. Existing PCB design software could be used.	The glue blob shape (wetting and viscosity) and size have to be controlled.
Straight pin (ball, peg) with White glue (pin attach and via-fill)	Easy/cheap to manufacture. Existing PCB design software could be used.	Might have too much loss

Embedded micro-mirror (under chip, inner layer)	exists on prototype	Super-expensive to manufacture. No PCB CAD support.

Chip pin types		
Method	Possible problems	Advantages
Direct fibre attachment	Low number of connections possible, high cost manufacturing.	Exists today
Optical solder ball	Too much dispersion above 100-200Gbps. Directional coupling might be problematic.	Almost the same technology as exists today.
Optical solder peg	Might be fragile. Directional coupling might be problematic.	Good above 200Gbps, reasonably simple manufacturing.
Gullwing pin	Glue blob dimension control might or might not be difficult, requires research. Probably not suitable for backplanes.	Probably the safest and easiest way of making low accuracy positioning with good directional low loss attachment.
Trace-end micromirror	Very high tolerance required for manufacturing the mirror walls and for the component position.	Test vehicle has been made and was functional.

Electro-optical PCB-design CAD software support:

Either trick the current cad software by assigning fake internal layers for optical related (interconnect, optical pad, optical under-fill, optical solder paste or glue) layers, or implement support for them both in symbol editor and in PCB editor.

Backplanes:

Backplanes would contain much more optical interconnects than a normal PCB. This way they will need multiple optical layers, optical vias and pitch black substrates. A blade design could be made using one optical layer, normal FR4 material (not black), and optical trace on copper trace for crosstalk insulation. A backplane would need pitch black substrate anyway, as inner layer insulation cannot be done otherwise, so the copper trace base will not be required. The optical vias are drilled mechanically or with laser, but then instead of copper plating them we fill them with optical material. Mechanical drilling would always have long stubs, even with back drilling, so laser drilled micro vias will be necessary, and therefore the stackup would be a build-up type. Back drilling leaves a small stub, around 0.2mm, which can be too much at higher data rates. The backplane could be optical-only, not containing any copper conductors or planes. The power delivery and management would require a separate backplane, like in ATCA systems.

In backplanes it will be necessary to have 90 degrees optical connections, so the gullwing or solder-arc methods could not be used to provide optical coupling due to mechanical arrangement. These right angle connections would be the optical solder pins and the optical vias. For these we need to find a way to direct the light in the right angle. One possible method would be the use of white light-scattering glue for both the pin-to-PCB connection and for the via-fill, together with a black or reflective underfill or coating material. It would be expected to have a few dB loss at these 90 degree connections, so the number of vias per signal will have a hard constraint. We will have to build a test vehicle with 2 perpendicular light pipes glued together with white glue to measure how much light gets through (or how much loss incurs), comparing to a straight light pipe. The test vehicle can be a low speed simple setup, as the loss is not high-frequency loss like in high speed electrical interfaces, but directly reducing the total amplitude.

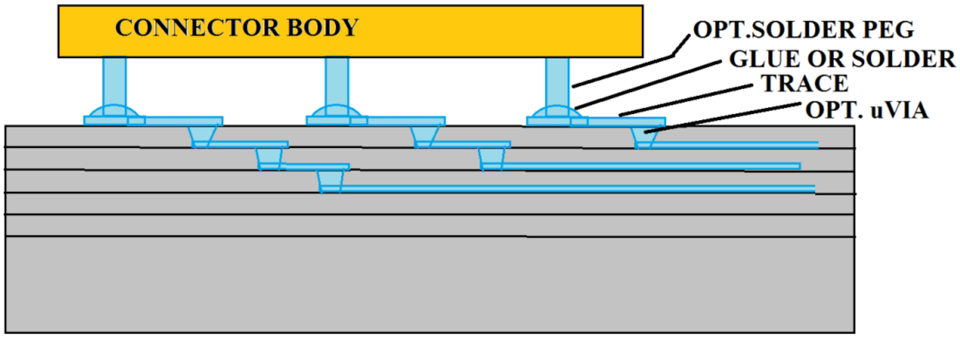

Layer Stack up of an optical-only backplane

Another possible way of making an optical backplane is by using a flex PCB cut into a comb shape and twist the teeth. This way the connector could be gullwing style (that can provide good directional coupling), and it would not require optical vias neither multiple layers as the optical traces will not have to cross connectors and each other. This is the case for single-star and chain topologies. For mesh and dual star we would still need multiple optical layers, or multiple flex backplanes (one for each switch slot). For a dual star 14 slot chassis with 4-lane ports (4 TX and 4 RX) we would use two flex backplanes, instead of 192 optical fibre-cables (that would be impossible to wire up correctly).

The connectors will have to be screwed onto the back wall of the chassis, or into a solid mechanical backplane that has rectangular openings on it, to allow for blades to be plugged in/out of the chassis without having to also handle the backplane or the back side in the process.

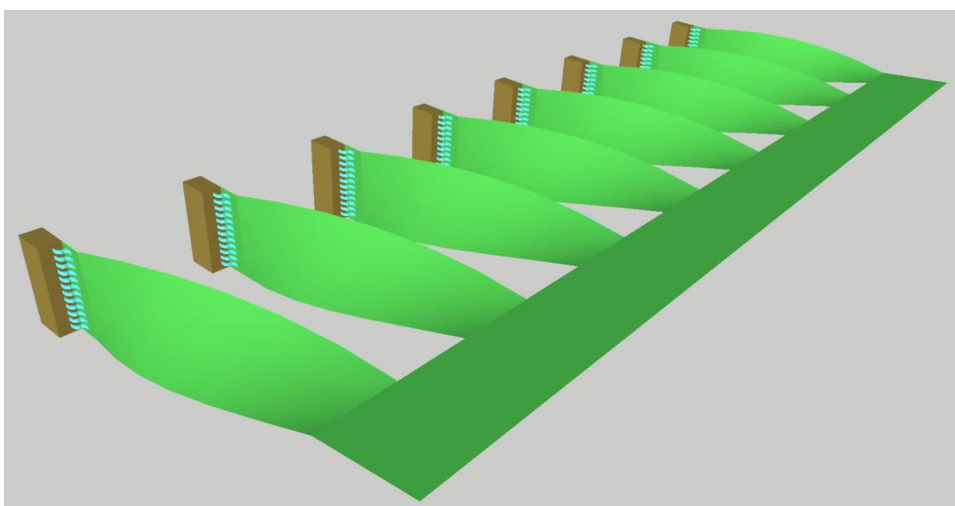

Flexible optical backplane

The backplane does not have to be flex, when we use traditional vertical connectors. However it will have to be extra wide (route signals around the connectors, not under) to avoid signal-signal and signal-connector crossings that is required for manufacturing with a single optical layer without optical vias. Optical vias are the difficulty here. This would be similar to the Xyratex optical backplane concept.

Systems

The printed interconnect is an important element of PCB-optics, but a working product would require additional developments, like silicon photonics transceiver chips and mixed material soldering using electrical and optical BGA balls on the same chip. Only the high-bandwidth data-place interconnects would be needed to be optical, the rest of the "circuit" would remain multi-layer electrical. The optical interconnects would be routed on the surface, that is normally filled up by components, so the new technology would only allow for less high-density designs. We will have to leave “routing channels” between the components. The high bandwidth optical data plane would probably only occupy one layer on the top, the electrical will remain multi-layer. Trace width would normally be much narrower than typical copper traces, so higher density of the optical portion would be a norm.

As I see it, the technology jump could be implemented in three phases:

1. Gen 1: Use many pieces of electro-optical transceiver chips around large electrical-only silicon devices like switch chips, FPGAs, ASICs, processors. The trick is to use both copper and optical transmission in a chain. Most (90%) of the route we use optical, because at higher speeds it is very difficult or impossible transmit signals over copper. We implement part of the route as copper, as this way we can use already existing high-complexity silicon devices from various vendors. Then we convert electrical to optical using small size transceiver chips. The chip has to be very small, as in a typical system we might have 100+ lanes coming out of a switch chip, and 4-50 lanes from one of today’s FPGA devices. For example the Altera Stratix-10 device (announced) will have many lanes of 50Gbps transceivers that can only be used for close chip-to-chip connections, where the link partner is placed very close on the board. For real systems, the board size and chip distances might be a lot larger. To make it work, we have to extend the range of those 50Gbps transceivers by using optical on most part of the path. The electrical to optical conversion takes place the closest to the chips, not at the backplane/daughter card connectors, to reduce the length of the electrical segment. The length of the optical segment can even be kilometres without problems, while the electrical portion struggles with inches. One might try to avoid having any optical features on the blade design, then build the transceivers into the backplane connector, but then the technology would go obsolete in 1-2 generations, as the chip-to-connector distance (even 1-3 inches) will become too long at certain speeds. So the main data plane routings will have to be optical on the blades, on the backplanes and throughout the connectors.

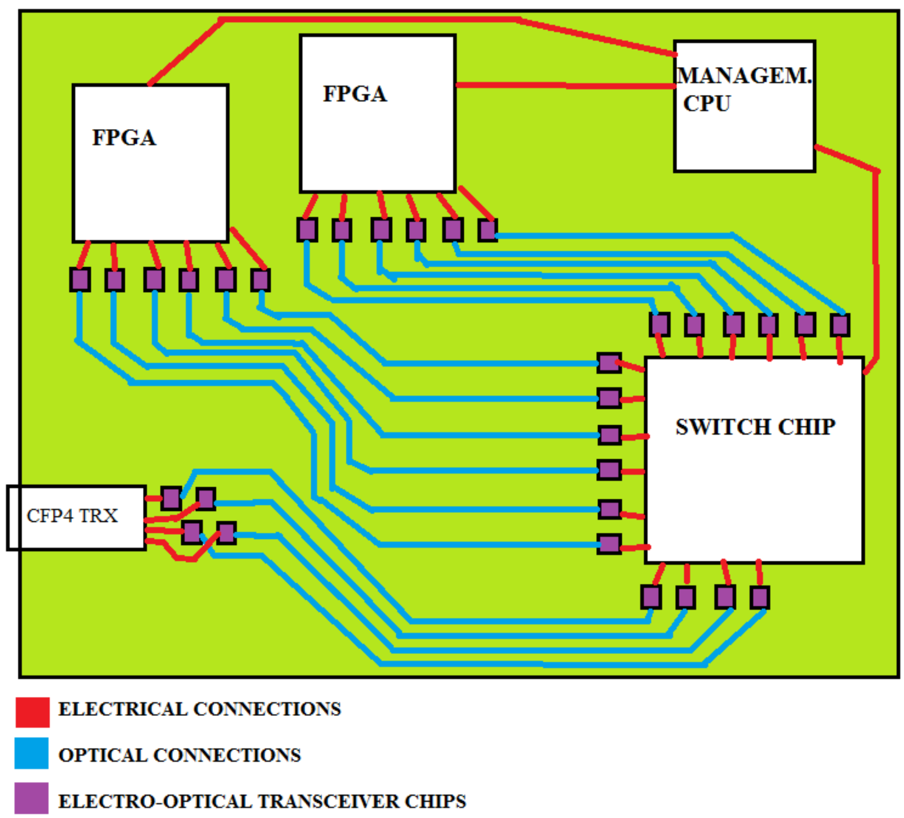

Gen1 system

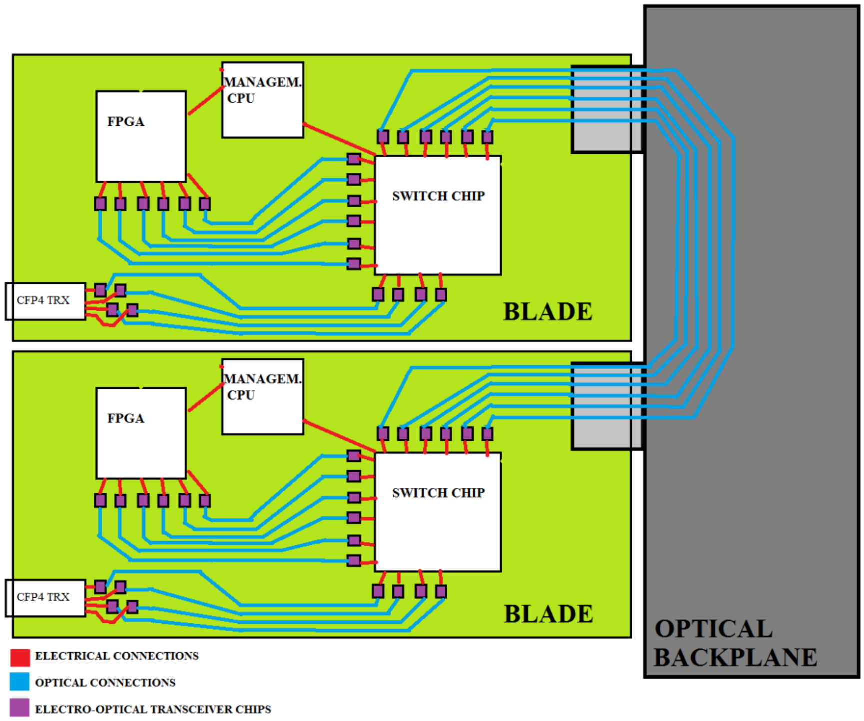

Gen1 system with backplane

2. Gen 2: Integrate the transceivers on the BGA package of large silicon devices. Taking the gen1 technology one step further, we can reduce the copper portion of the signal path to allow for even higher line rates. The gen2 transceivers have to be even smaller than gen1, to be able to fit 10-128 pieces of them on a typical 2”x2” BGA package substrate. The external interface with a pluggable module would not be an electro-optical transceiver anymore, but an optical-to-optical repeater. The small on-board transceivers could not drive longer (km) distances, hence the need for a high-power repeater. For a few meter connections within the same data centre we might not need repeaters, so the pluggable module will really be just an optical fibre connector.

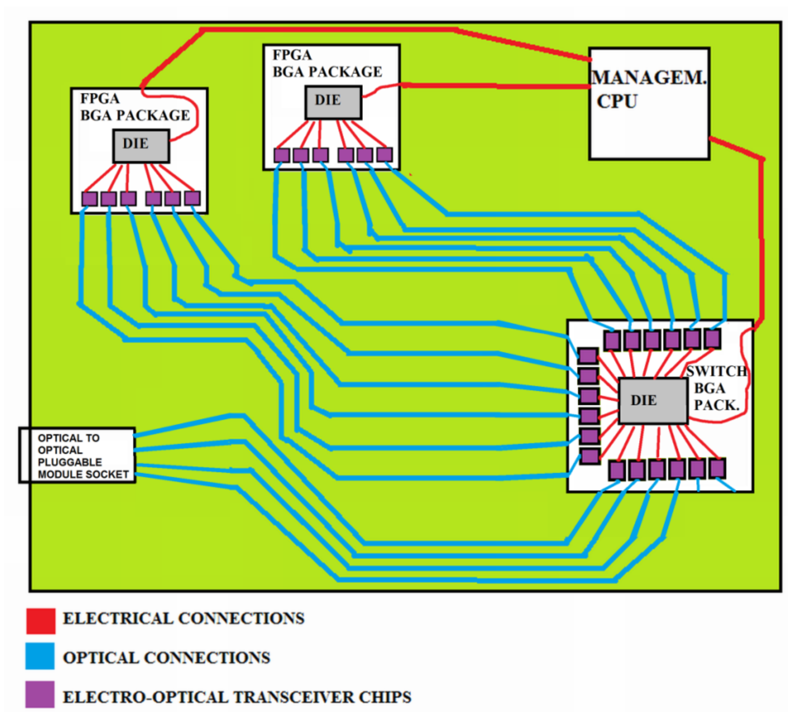

Gen2 system

3. Gen 3: Integrate the electro-optical transceivers on the main silicon device die, like switch chips, FPGAs, ASICs, processors. The copper interconnect is completely removed from the data plane signal path, this way the speed limit is removed too. This would allow for much higher speeds, and likely be used for many generations of the systems and devices. We may need 3 generations to get here, as the technology investment can be made one step at a time, reducing risk and cost of initial systems. We can mix gen1, gen2 and gen3 devices on the same board, allowing for great flexibility to build systems. Compass-EOS, a start-up already has a similar device, where they have optical from the die, although it's connection to the outside world might be different than the one suggested by this article.

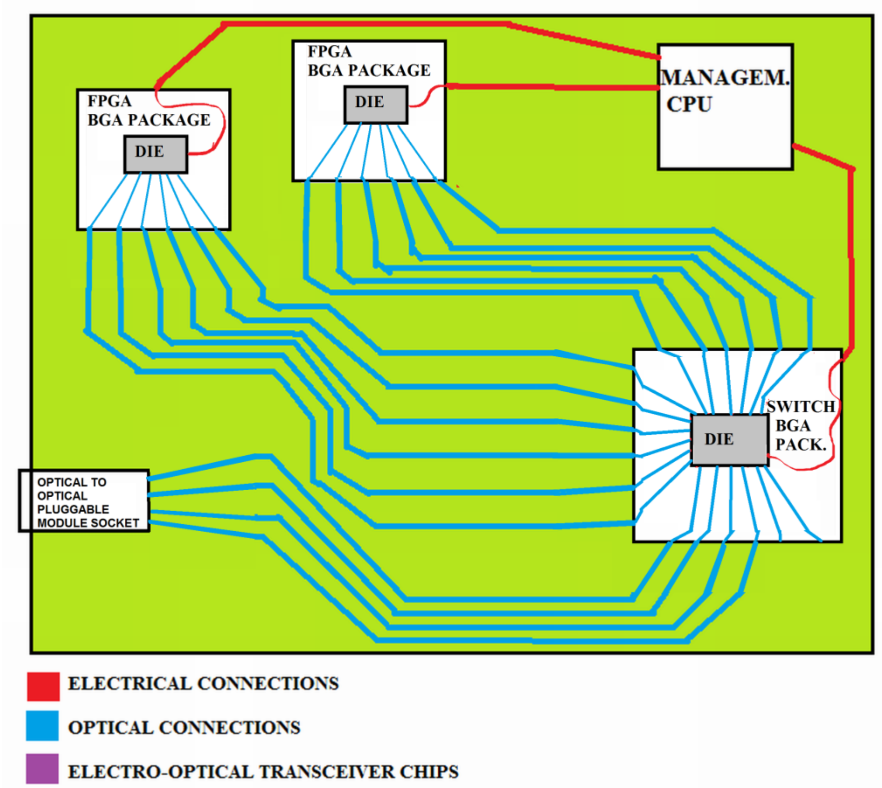

Gen3 system

Summary

We need to develop several new technologies:

- Electro-optical PCB manufacturing. Already in development.
- Multi-layer optical-only PCB manufacturing for backplanes.
- Mixed-material reflow soldering, new glue and under fill materials for optical attachment.
- Electro-optical PCB design CAD software support.
- Electro-optical transceiver chips using silicon photonics or other semiconductor technology, or an optical solder-plug for existing embedded transceivers.
- High-density optical board-to-board connectors for backplane and daughter-card.

Once we define new board/blade form factor standards, those should be designed in a way that we can upgrade them for higher speed later for several generation without redefining connectors and assembly technology, like we have today with the electrical-only systems, like ATCA, XTX, VPX…