Chip-to-board Interconnects for High Performance Computing

Markus B.K. Riester*[,a], Ruth Houbertz-Krauss[b], Sönke Steenhusen[b]
[a] Feldkirchnerstrasse 111, 8055 Seiersberg, Austria; [b]Fraunhofer Institute for Silicate Research ISC, Neunerplatz 2, 97082 Würzburg, Germany

ABSTRACT

Super computing is reaching out to ExaFLOP processing speeds, creating fundamental challenges for the way that computing systems are designed and built. One governing topic is the reduction of power used for operating the system, and eliminating the excess heat generated from the system. Current thinking sees optical interconnects on most interconnect levels to be a feasible solution to many of the challenges, although there are still limitations to the technical solutions, in particular with regard to manufacturability.

This paper explores drivers for enabling optical interconnect technologies to advance into the module and chip level. The introduction of optical links into High Performance Computing (HPC) could be an option to allow scaling the manufacturing technology to large volume manufacturing. This will drive the need for manufacturability of optical interconnects, giving rise to other challenges that add to the realization of this type of interconnection. This paper describes a solution that allows the creation of optical components on module level, integrating optical chips, laser diodes or PIN diodes as components much like the well known SMD components used for electrical components. The paper shows the main challenges and potential solutions to this challenge and proposes a fundamental paradigm shift in the manufacturing of 3-dimensional optical links for the level 1 interconnect (chip package).

Keywords: 3-dimensional optical interconnects, optoelectronic packaging, chip-to-board interconnects, high-performance computing, manufacturability, interconnect hierarchy

1. INTRODUCTION

Today's communication infrastructure is tightly connected to technological advances in computing, data storage, "big data" applications, and optical transport. The amount of data that is being transported has reached unseen and ever increasing volumes. Around the globe large corporations operate data centers to provide distributed computing and redundant storage across continents. These data centers require power at megawatt levels. Often they have their own power supplies, redundant fail-safe power backup, and independent cooling. Figures on the power usage efficiency (PUE) show that most sources attribute 50-60% of power consumption to the server component level (processor, memory, storage), and the remaining 40-50% to support systems like the power supply, power distribution system, UPS system, cooling system and building entrance switchgear and medium voltage transformer. Google states that its PUE is at 12% [1]. A reduction of 1W at the server component level (processor, memory, storage) results in additional savings of 1.84W [2]. Thus, reducing power consumption on the processing side will trigger a cascade effect that further contributes to the reduction of power consumption. All operators of these data centers strive to use technologies that allow data transport, computing and storage being done at lower power consumption, and solutions are sought in all aspects of the data center technology.

Moving forward on the topic of energy-benign technology IBM has introduced a novel hot-water cooling concept in the SuperMUC computer [3] that allows cooling to be done very close to the active component. The excess heat is brought to further use for office heating. The technology allows the reduction of energy used for powering the cooling equipment by 40%. While the technology is a great advancement over legacy cooling systems, the final goal must be to reduce the use of power that is put into the system, so it does not need to be removed from the system. The realization of high performance computing systems thus requires innovation on all aspects, both on the processing systems as well as the infrastructure technology. One of the designated technologies for advancing energy-saving computing systems is moving

*Markus.Riester@maristechcon.com, phone +43 316 228660 www.maristechcon.com

Optoelectronic Interconnects XIII, edited by Alexei L. Glebov, Ray T. Chen, Proc. of SPIE Vol. 8630, 863002 · © 2013 SPIE · CCC code: 0277-786/13/$18 · doi: 10.1117/12.2008817

optical interconnects closer to the processor. In this paper we propose to use the interconnect hierarchy concept known in the electronics manufacturing industry to specify the options for system partitioning. A means is proposed for manufacturing optical interconnects that allow the scaling of optical package manufacturing comprising Chip-to-package optical interconnects to high-volume manufacturing.

2. THE INTERCONNECT HIERARCHY IN ELECTRONICS MANUFACTURING

2.1 The interconnect hierarchy in electrical systems

The evolution of electrical, electronic and optical systems started with simple devices that were composed of individual parts joined together to perform a designated task. When the tasks got more elaborate, the systems got more complex a more structured way of designing and manufacturing systems was required. This need was accompanied by technical evolution that allowed the condensation of functions into smaller, more integrated parts, e.g. integrated circuits. These ICs then were assembled on a substrate, which served as mechanical and electrical interconnection.

The interconnect hierarchy [4] was conceived in the 1960s as a means to describe the connections to be made between the different levels of integration in the system. It allowed to distinguish the requirements on different interconnects Levels. Interconnect Level 0 represents the interconnections on the chip, interconnect level 1 the chip-to-package interconnections to form the IC package, interconnect Level 2 stands for the assembly of components to the printed circuit board (PCB), on interconnect level 3 circuit boards are connected to racks, or multiple PCBs are connected to each other, interconnect Level 4 is the wiring and cabling connection in the cabinet. On interconnect Level 5 cabinets are connected to each other, interconnect level 6 connects the collection of cabinets with the "outside", e.g. as connection out of a data center. Each integration level typically requires different types of interconnects. In this paper focus is on interconnect levels 1 and 2, which typically are similar in nature and hence similarly created.

Today, electrical systems are commonly built in a modular way that requires an interconnect hierarchy that is modular in itself to account for the multitude of variations in packaging technology. The Jisso interconnect hierarchy shown in Figure 1 [5] sets out to formalize this new interconnect hierarchy. It allows a more precise categorization of the level hierarchy and is helpful for system designers to do in-depth Design-for-manufacturing, i.e. optimization for best cost, highest reliability, highest level of miniaturization. It is also useful for the members of the supply chain to identify business opportunities through adjusting their technical capabilities.

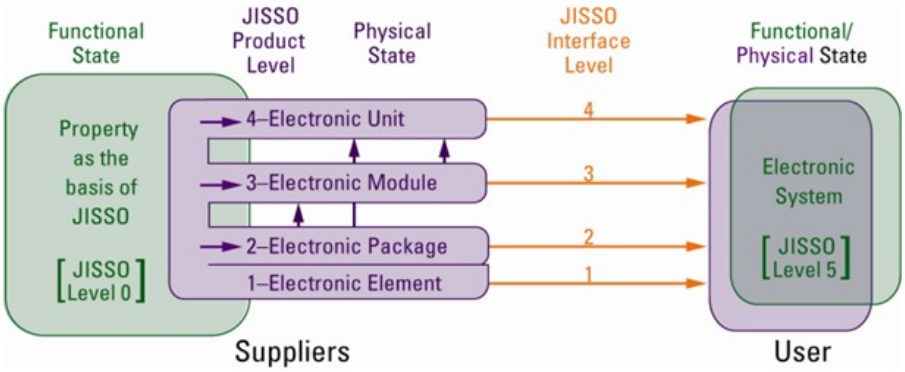

Figure 1: The Jisso interconnect hierarchy provides a systematic modular approach to the interconnect level concept used in electronics manufacturing. It takes into account the various combinations of components and substrates that are conceivable and to a great part used today.

2.2 The use of an interconnect hierarchy in optical interconnections

Complex systems using electrical interconnects nowadays require the more complex and modular Jisso interconnect hierarchy for a precise description of state-of-the-art packages and constructions. Optical systems however, due to the currently observed lower complexity of the interconnect could potentially make use of the original model. This view is supported by the segmentation used for describing the application scenarios for optics in systems. In bandwidth/time vs. link length charts like shown in [9] the need for interconnects between the different levels becomes clear: Wide Area Network Optical interconnects (OIC) have different requirements than OICs on the system level. Advancing further to the module or chip level, the OICs will again have to suit other requirements. Therefore it will be helpful to clearly distinguish between the different application domains when discussing optical interconnect realization strategies. As in electrical interconnects, the technology behind the OIC might be different, or it might be similar, depending on the use case. The advantage of using the interconnect hierarchy is that the thinking about system partitioning can be more dynamic by adjusting the boundaries of the levels to optimize manufacturability, system performance, reliability, or modularization with the goal of reducing overall system cost.

2.3 The physical realization of interconnects

Data and power transport connections between different components in a system can be achieved through physically linking them. These links can be a material connection, or it can be wireless. In the electrical domain, typical representations of material connections are solder joints in pin-in-hole and surface mount technology, galvanic connections, and welded or screwed contacts (in PCB Assembly), connectors (to connect PCBs with larger components or cables) [6] Wireless connections make use of electrical fields for data and/or power transport, are however used only where no direct material contact can be created . Only material connections shall be considered in the following.

A typical example for an interconnect on level 2 is the soldered connection of an electrical component to the printed circuit board (PCB), which today is commonly attached in surface mount technology. Large or heavy components may be screwed to the PCB to enhance mechanical stability; this mechanical connection may also be an electrical connection (e.g. for ground.) The use of conductive adhesives for establishing electrical contact is common practice today. The large variety of types of electrical interconnects results from the fact that the nature of electrons allows low-loss connectivity using all these different forms of establishing electrically conductive interconnections. Standardization of manufacturing processes and materials has helped to advance the technology over the past 50 years to a point where large volumes of interconnects can efficiently be created for even modestly complex products.

In the optical domain, similar representations exist: optical interconnects are usually material connections of light guides for enabling optical data transfer and illumination optics, i.e. for lighting or power transmission), and wireless transfer is represented by free-space optics. In the electrical domain many different realizations of interconnects have been used for designing and building products. The number of optical interconnect types actually in use is lower as optical interconnects are much harder to manufacture efficiently in high volumes. In particular the realization of interconnects on chip, package and board level has been successful only in applications where electric interconnect could not compete for technical reasons, like bandwidth requirements, interference and interconnect density [7]. While means for creating low-loss optical connections on the chip, the package and the board level are known the economically feasible realization has been a more difficult task. The nature of photons requires a more precise way of addressing the interconnect realization, which usually requires more precise alignment, and more complex interconnect designs. As cost needs to be taken into the equation the manufacturability of the optical interconnect will be added to the list of requirements for volume manufacturing.

In the recent past, work on the TERABUS system by IBM has shown the opportunities of optical modules that are surface mount compatible [8, 9, 10]. In its current form it will already allow existing manufacturers use many of their existing process equipment and manufacturing strategies. The technical solution points in the direction of enabling better manufacturability, and may be enhanced by regarding further requirements of volume manufacturing processes. The chip was mounted on a substrate carrier allowing it to be assembled like an electrical SMT device. The interconnection of the "holey" optochip to the waveguides was done using free-space optics with a lens arrangement, requiring a) additional alignment steps to align the lens array with the mirror array reflecting the light out of the plane of the waveguides, b) creation of the mirror array using laser ablation (which is a rather "messy" process", and c) the alignment of the chip to the lens array. By eliminating the lens array, the need for total internal reflection optics on the waveguides and elimination of the alignment of parts a much simpler process seems possible that allows an even more manufacturable process to be implemented.

3. 3D OPTICAL INTERCONNECTS

3.1 Optical Waveguides and Optical Interconnects

The creation of optical interconnects has been addressed by many groups in the past two decades, mainly in the form of optical waveguides (OW). Therefore many different technologies have been explored for manufacturing optical waveguides, like photolithography [11], hot embossing/micro molding [12], laser ablation [13], nano imprint lithography [14], laser direct writing [15] and non-linear optics using two-photon absorption polymerization [16, 17].

These waveguide manufacturing technologies can in principle be used for achieving optical connectivity, but it will be hard for most of them to be scaled to high volume manufacturing. This is true not only for multi-mode waveguides, but in particular for single-mode waveguides. The prime reason is the required precision of the alignment of the optical components to the waveguides. State-of-the-art assembly equipment is capable of both, handling substrates and dies, but will be very slow when sub-micron precision is required. This is critical for scaling to large volumes of product. In addition, the alignment of multiple I/O coming from an optical arrangement, an optical chip or a laser bar will be even harder than for a single I/O, if technical solutions are sought only on the waveguide manufacturing side. Current state-of-the-art in assembly of such optical arrangements is active alignment, where the device needs to be powered up and emit/receive light like in live operation. While active alignment is common practice, passive alignment is the preferred choice as it is faster and simpler to implement.

While some of the waveguide manufacturing technologies have reached a level of sophistication that would allow the transfer to volume manufacturing of optical waveguides, this might still not be the solution for the problem that needs to be solved. As the waveguide only represents one part of the complete optical interconnect. Additional features need to be included in the development to account for the complete system to deliver an added value to the final product. As minimum requirement, the features include the drivers, E/O and O/E conversion, light source, waveguide, detector und amplifier.

The understanding of the requirements of optical interconnects has made progress over the past years, and the main and fundamental challenge of creating optical interconnects (OI) that can be manufactured in larger volumes in an economically viable way has been addressed This was mainly driven by telecom and datacom applications, propelled the development of OI strategies that would allow manufacturing of optical or optoelectronic systems to be done similar to the way that electrical systems are built today [9]. In specific, the challenges for advancing optical technology closer to the PCB and chip package, or into packaging levels one and two are in focus.

To understand the challenge better, a look to the electronics manufacturing industry is helpful. Scaling of manufacturing of products to high volume manufacturing has been demonstrated by this industry. It could also serve as a role model for the optoelectronics industry. There are plenty of examples how innovation in miniaturization and integration enabled new features in mobile consumer products and supported volume manufacturing: a) the advent of surface mount technology enabling faster assembly, b) shrinking components for enabling thinner buildups on a smaller footprint, c) the use of various generations of area array technology (Ball grid arrays, land grid array, …), allowing the reduction of number of assembled components and d) stacking of components to die stacks, yielding highly integrated system-in-packages, to name a few. Electronics manufacturing has come a long way of putting more computing power into silicon, and more integration capability into level 1 and level 2 packaging. Enabling the OI to be efficiently manufactured will allow optical technologies to advance into interconnect levels 0 to 2. Extending the idea, it might be helpful to consider various alternatives of interconnecting components on different interconnect levels.

3.2 Optical interconnects in the interconnect hierarchy

The interconnect hierarchy [4, 5] logically partitions a system into various levels that are interconnected by physical connections. The model had been created and is used for electrical connections, but may well be applied to understand the system partitioning in the optical domain as well. It is helpful to understand the complexity that interconnects may have on the respective levels, or, in other terms, on which level complexity is best aggregated to eliminate the need for complex interconnects. As only a few standards exist today with respect to optical interconnections, let alone system

design, variations of system concepts will persist, and today's solutions will remain specific to individual systems. In the following, two examples are given that outline the concept.

On level 1, optical chips are connected to their package. The chips might come with only one optical I/O, or with many. Thus, either a single interconnect needs to be formed, or multiple. In both cases, the optical output of the chip needs to be precisely aligned to the waveguide or the free-space optical system comprising lenses. A repeatable and standardized manufacturing process will be required to deliver optical packages that can be used in large-scale systems The complexity of the technical solution will inherently influence the scalability of the system. When making the decision for a particular solution, the manufacturability of the optical interconnect is determined by a) how light is delivered (input), b) how light needs to be taken from the input to the output c) how light is delivered at the output of the waveguide (or free-space optics), d) the environmental conditions under which the module needs to operate, e) the reliability (MTTF, MTBF), d) manufacturing yield and e) cost. In Chapter 3.3 a process is introduced that allows the manufacturing of component packages that can be connected using standard fiber connectors, without the need for lens arrays or other optical alignment required at the package level.

On level 2, the module-to-PCB interconnect level, the optical connection between the package and the substrate is formed. On level 2, the reduction of the number of process steps is even more critical, as level 2 is the most cost sensitive interconnect level in the value chain [18]. Various arrangements using mirrors, micro lenses and other means for coupling light in and out of a waveguide can be found in the literature [19 20]. The main challenge is the registration tolerance that needs to be reached during assembly and maintained during operational life. Also, the number of interconnects to be simultaneously formed will be critical, as multiple I/Os will require a precise positioning of interconnects with respect to each other. Thus, the module I/Os are imposing tight design rules for the counterpart on the substrate/PCB side, which needs to be ideally matched. While solutions employing arrayed lens coupling can be used, the registration of the module to the lens array is still critical, and it comprises an additional manufacturing step. In addition to the proposed free-space lens array coupling solutions optical connectors might be used. These again will require an optical interface that needs to be created to interface with the connector, which merely shifts the problem, and adds more complexity to the process.

An alternative and potentially more attractive approach as an intermediate (and potentially final) solution on level 2 is the separation of the electrical and the optical interconnect, which would lead to system concepts where optical connectivity on level 2 is provided through optical cables. A similar concept has been described in [21] for high frequency electrical signal transmission.

Level 0 is not regarded in this paper.

3.3 New manufacturing process flow for optical component surface mount assembly

Existing manufacturing processes for optical waveguides have not been able to meet all requirements for high volume manufacturing of integrated optical modules. A new interconnection technology and manufacturing process flow has been developed that allows the elimination most of the drawbacks and that is scalable to volume manufacturing. In addition to the new way of creating the waveguide structures, a fundamental paradigm shift for the process sequence for the assembly process is proposed. The proof of concept has previously been shown in several arrangements [22], and this new implementation will allow the process to be used for creating optical components that can be used like electrical modules during assembly, while exhibiting the performance of optical modules or optical multi-chip modules (OMCM).

The process flow for creating 3-dimensional optical interconnect structures is outlined in Figure 2a. The advantage of the process flow is that the large number of process steps required for generating waveguide structures on planar substrates is significantly decreased, from more than 40 process steps to as little as 5 steps. The process is significantly shorter, uses less process material, and less energy. In addition the technology enables the creation of 3D waveguide structures not attainable with legacy processes.

The process allows optical components like laser VCSELs, or optical chips to be interconnected with the corresponding receiving devices like PIN diodes. As a first step the optical components are assembled on a substrate, carrier or circuit board. For an individual optical chip the assembly is done into a socket, where also electrical interconnections are created. The socket is surrounded by a frame that is used for creating a defined volume. This volume is filled with one material that later will be the waveguide *and* the cladding material. It may also serve as glob top and/or underfill. The socket containing the optical chip and the waveguide material are subjected to a thermal curing operation where the

material is slightly cured/dried. Subsequently the waveguides are written into the material using a two-photon absorption (TPA) structuring process. This process allows the waveguide structures to be precisely aligned with the designated optical input or output location. The alignment process is separated from the assembly process. This leads to the situation that for the assembly process existing assembly equipment may be used, and no comparably slow and costly equipment is required for creating the optical component-to-waveguide link in parallel to the assembly of the electrical parts. The precision alignment is possible as the geometry of the optical chip is well known, and can be used for precisely aligning the TPA laser with the optical output. At the same time, the socket frame has optical windows at specified locations that can be easily identified in the alignment process. With the location of both end points of the optical I/O defined, the calculation of the waveguide is done in situ by geometrical optimization of the waveguide with respect to insertion loss. Pre-calculated optical waveguide segments may be also used and connected in this process to reduce process time. The TPA structuring process allows multiple waveguides to be created inside the material. The waveguides may be arranged in parallel; they may intersect, and may come in bundles to allow space multiplexing (see Figure 3 for an example). The waveguides may be tapered (in situ) to account for differences in diameter of the optical I/O on either side of the waveguide. After the waveguide and interconnect structuring has been completed, the remaining material is cured to stabilize the waveguides inside the matrix. The completely polymerized package is mechanically stable and carries optical interconnects that connect the optical chip with the package.

As comparison one of the existing process flows for manufacturing an optical waveguide and interconnect is shown in Figure 2b. The number of process steps alone hints towards a very complex manufacturing process chain, with the corresponding price tag. This process sequence and similar process sequences are described in literature (see literature on optical waveguide manufacturing referenced above).

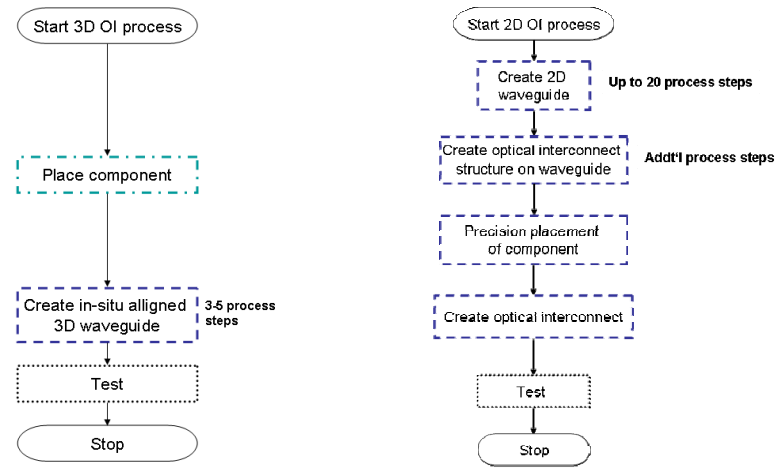

Figure 2: Comparison of legacy process flow with new, more efficient process flow that allows fully flexible manufacturing of optical waveguides in 3-dimensional space. a) shows the new process flow, b) outlines the principle of legacy process flows using subtractive structuring methods.

The number of process steps for attaining 3D-waveguides can be reduced to a minimum of 3 individual steps.

The process flow as outlined for the TPA process in (Figure 2a) can also be used for creating 3-dimensional optical interconnects. This is helpful for multiple reasons. In the application described above, the feature can be used to account

for misalignment in all three axes: x-y and height. The process sequence remains identical, only the programming of the waveguide structures needs to account for the change in depth through which the imaging laser light needs to travel.

This process is carried out separately from the die attach/placement and links the optical output location with the desired destination of the light signal. The distal end of the optical interconnection can be a standard connector, or fiber array connector, or a specified location in the frame, which serves as the base for insertion of such a connector. After the waveguide polymerisation process using TPA structuring, the matrix material is polymerized, thus creating a protective layer on the chip, much like a dam & fill process including, however, the optical interconnect structures from the chip to the connector. The described process has been shown to work in PCB manufacturing [22], with subsequent processing of the optical layer in a lamination press at 200°C. Current materials available maintain a refractive index difference Δn of 0.005, which is large enough for SM waveguides to adequately guide light [16].

The precision of the optical alignment attained in this process is well below 0.5 um, which is sufficient for creating multi-mode and single-mode optical interconnections. The diameter of the waveguides can be designed as needed, to suit the application.

3.4 Requirements and opportunities for real 3-dimensional optical interconnects

3.4.1 Registration

In addition to creating waveguide structures the waveguides can be registered in situ to any optical device, such as laser diodes or optical chips as well as corresponding receiving devices with known geometry, as described in chapter 3.3. The alignment process is decoupled from the assembly process, which fits well with the existing manufacturing industry landscape. The main prerequisite for the precision registration process is to know the exact geometrical data of each position of the optical I/Os in the arrangement and their optical coupling properties, e.g. (numerical aperture, mode field diameter, geometry …). As this data is readily available from design data of the components the end points of the waveguide structures and their geometries can be well described. The registration then is carried out by identifying the component, measuring this geometry, using the resulting data for calculating the location of the I/O using the design data. This defines the starting point and the three-dimensional orientation of the waveguide structure. The same procedure is applied to the other end of the waveguide structure. With all geometrical data known the waveguide structure can be calculated in situ. This allows the process to compensate for any misalignment occurring in the placement process, and precisely aligning the waveguides with the optical I/Os on either side.

3.4.2 Optical Interface design

In addition to simply writing a 3D waveguide structure with the TPA process, the interfaces between the waveguides and the optical I/O may be tailored to maximize the coupling efficiency. While simple I/O structures might merely consider the numerical aperture of the arrangement and allow the waveguide structure to be starting with a standoff, adjusting the waveguide to the mode field diameter by using a taper is significantly more challenging. While the optimization of the optical interconnection using the 3D waveguide structuring is possible, the considerations for optimizing the system design should be done in an earlier development stage. In Figure 3 shows the positioning of perfectly circular waveguide facets. The waveguides were created using two-photon absorption polymerization. The image shows the fluorescence image of the cross section of the waveguide array. The optical power and writing speed for triggering the two-photon absorption polymerization process was varied to identify the potential for dynamically tuning the diameter of waveguides. The diameter was varied between 3…8 micrometers.

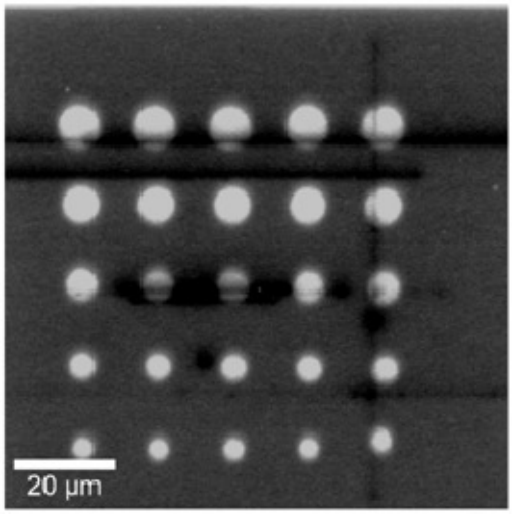

Figure 3: Fluorescence characterization of the facets of a waveguides array created using two-photon absorption. The diameter of the structures can be tuned by varying the optical power and writing speed of the 3D imaging process.

3.4.3 Material requirements

The versatility of the process offers various opportunities for simplifying the creation of optical modules and devices. At the same time the process depends on the usage of designed materials that have very specific properties. The main reason is that the material used for creation of the 3D-structures is applied once and then needs to fulfil multiple purposes: a) It is polymerized in the two-photon polymerization process to yield the desired 3D structures; b) the non-polymerized matrix material is subsequently cured using a different polymerization mechanism so that the optical density, hence the resulting refractive index, is different from the TPA polymerized material; c) the Δn needs to be retained throughout the operational lifetime of the optical device; d) a low material absorption is needed (<0.1 dB/cm); e) the material needs to withstand the process conditions of the optoelectronic substrate and component manufacturing; f) other desirable properties like low water uptake, CTE matching the optical chip material and the surrounding socket, or flexibility.

Materials that suit most requirements have been under development and have been shown to work in the laboratory and demonstration arrangements [22]. These multifunctional materials have shown good performance in the particular arrangement. However, due to their multifunctionality they usually need to be optimized to serve all specific application requirements.

4. SUMMARY

Optical interconnects are required for advancing product performance in high end computing and telecom applications. The prime reasons are performance, interconnect density and energy efficiency. As the requirements towards volume manufacturing are getting more stringent with rising volume, the design of optical interconnects needs to regards manufacturing principles that are known from the electronics manufacturing industry. The experience gathered in that domain over the past decades can be leveraged. One means for aligning the thinking in optical electronics integration with electronics manufacturing is the use of the interconnect hierarchy. It will support understanding system partitioning and creating different options for identifying the most efficient approach for volume manufacturing.

REFERENCES

[1] Google Inc.,"Efficiency: How we do it lt," http://www.google.com/about/datacenters/efficiency/internal/ (28 December 2012)

[2] Emerson Network Power, "Reducing Data Center Energy Consumption by Creating Savings that Cascade Across Systems", Emerson Network Power (2008) http://www.google.com/about/datacenters/efficiency/internal/

[3] "First Commercial IBM Hot-Water Cooled Supercomputer to Consume 40% Less Energy" (18 June 2012); http://www.prnewswire.com/news-releases/first-commercial-ibm-hot-water-cooled-supercomputer-to-consume-40-less-energy-159392045.html; (02 January 2013).

[4] Pinnel M.R., Knausenberger W.H., "Interconnection System Requirements and Modeling," AT&T Technical Journal 66, 4 July/August (1987)

[5] Fjelstad, J., "Rethinking the Hierarchy of Electronic Interconnections," IPC Outlook, 28 December 2012, http://www.ipcoutlook.org/mart/52287F.shtml (28 December 2012)

[6] Groover M.P., [Fundamentals of modern manufacturing materials, processes and systems], Wiley, 830ff (2007).

[7] Pepeljugoski P., et al, "Low power and high density optical interconnects for future supercomputers," Optical Fiber Conference San Diego, 21-25 March 2010, 1-3 (2010)

[8] Doany F. E., "Power-Efficient, High-Bandwidth Optical Interconnects for High Performance Computing," IEEE HOT Interconnects Symposium, Santa Clara, CA, August 21-23 2012 (2012)

[9] Doany F. E., et al., "Dense 24 Tx + 24 Rx Fiber-Coupled Optical Module Based on a Holey CMOS Transceiver IC," ECTC 2010, 247–255 (2010)

[10] Schow C.L. et al., "A 24-Channel 300Gb/s 8.2pJ/bit Full-Duplex Fiber-Coupled Optical Transceiver Module Based on a Single Holey CMOS IC," IEEE JLT, Feb (2011).

[11] J. Kobayashi, T. Matsuura, Y. Hida, S. Sasaki, and T. Maruno, "Fluorinated polyimide waveguides with low polarization-dependent loss and their applications to thermo-optic switches," Journal of lightwave technology 16, 1024-1029 (1998).

[12] E. Griese, "A high-performance hybrid electrical-optical interconnection technology for high-speed electronic systems," IEEE Transaction on advanced packaging 24, 375-383 (2001).

[13] G. Van Steenberge, P. Geerinck, S. Van Put, J. Van Koetsem, H. Ottevaere, D. Morlion, H. Thienpont, and P. Van Daele, "MT-Compatible Laser-Ablated Interconnections for Optical Printed Circuit Boards," Journal of lightwave technology. 22(9), 2083-1090 (2004).

[14] T. Han et al., "Low loss high index contrast nanoimprinted polysiloxane waveguides," Opt. Exp. 17(4), 2623 (2009)

[15] F. Tooley, N. Suyal, F. Bresson, A. Fritze, J. Gourlay, and A. Walker, “Optically written polymers used as optical interconnections and for hybridization,” Optical Materials 17, 235-241 (2001)

[16] Houbertz R., Satzinger V., Schmid V., Leeb W., Langer G, "Optoelectronic printed circuit board: 3D structures written by two-photon absorption," Proc. SPIE 7053, 70530B (2008).

[17] Houbertz R., Wolter H., Schmidt V., Kuna L., Satzinger V., Wächter C., Langer G., "Optical waveguides embedded in PCBs - a real world application of 3D structures written by TPA," Materials Research Society Symposium Fall Meeting 1054-FF01-04 (2007)

[18] 2011 iNEMI Roadmap, iNEMI, Herndon, Virginia, USA

[19] Schröder H., Bauer J., Ebling F., Franke M., Beier A., Demmer P., Süllau W. Kostelnik J., Mödinger J., Pfeiffer K.,. Ostrzinski U., Griese E., "Waveguide and packaging technology for optical backplanes and hybrid electrical-optical circuits," Proc. SPIE, 6124 (2006).

[20] Van Steenberge G., Geerinck P., Van Put S., Van Koetsem J., Ottevaere H., Morlion D., Thienpont H., Van Daele P., "MT-Compatible Laser-Ablated Interconnections for Optical Printed Circuit Board," J. Lightwave Techn. 22, 2083 (2004).

[21] Fjelstad J.C. et. al, "Direct-connect integrated circuit signaling system for bypassing intra-substrate printed circuit signal paths", US Patent US7307293, (11 December 2007).

[22] Langer G., Riester M., "Two-photon absorption for the realization of optical waveguides on printed circuit boards," Proc. SPIE, 6475 (2007).