Richard Pitwon
Pluggable Optical Connector Interfaces for Electro-Optical Circuit Boards

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without the author's prior consent. Unless otherwise specified, all content in this thesis, including images, are the copyright of the author.

# PLUGGABLE OPTICAL CONNECTOR INTERFACES FOR ELECTRO-OPTICAL CIRCUIT BOARDS

by

**RICHARD CHARLES ALEXANDER PITWON**

A thesis submitted to Plymouth University

in partial fulfilment for the degree of

**DOCTOR OF PHILOSOPHY**

Faculty of Science and Technology

**September 2016**

# ACKNOWLEDGEMENTS

I would like to sincerely thank Dr Jasper Graham-Jones and Dr John Summerscales for their incredible support and guidance throughout the period of my PhD.

I would like to thank my colleague and friend Dr Kai Wang, as well as Dr David Selviah and Dr Ioannis Papakonstantinou from University College London, with whom I started this journey a long time ago and have collaborated extensively over the years.

I would also like to acknowledge my collaborating partners on the European projects SEPIANet, PhoxTrot and Nephele, whose collective expertise in these areas has deeply enriched my study and complemented my research activities. It has been a great privilege to know them and learn from them.

Finally my deepest and most heartfelt thanks go to my partner Anna for her unceasing support, encouragement and patience over the years, without which this work would not have been possible.

I dedicate this work faithfully to my mother.

# AUTHOR'S DECLARATION

At no time during the registration for the degree of Doctor of Philosophy has the author been registered for any other University award without prior agreement of the Graduate Sub-Committee.

Work submitted for this research degree at the Plymouth University has not formed part of any other degree either at Plymouth University or at another establishment.

This study was financed with the aid of a studentship from Xyratex Technology Ltd and carried out in collaboration with Xyratex Technology Ltd.

A programme of advanced study was undertaken.

Relevant scientific seminars and conferences were regularly attended at which work was often presented; external institutions were visited for consultation purposes and several papers prepared for publication.

Word count of main body of thesis: 73,746

Signed ……………………………………………………………

Date ………………………………………………………………

# ABSTRACT

A study is hereby presented on system embedded photonic interconnect technologies, which would address the communications bottleneck in modern exascale data centre systems driven by exponentially rising consumption of digital information and the associated complexity of intra-data centre network management along with dwindling data storage capacities. It is proposed that this bottleneck be addressed by adopting within the system electro-optical printed circuit boards (OPCBs), on which conventional electrical layers provide power distribution and static or low speed signaling, but high speed signals are conveyed by optical channels on separate embedded optical layers. One crucial prerequisite towards adopting OPCBs in modern data storage and switch systems is a reliable method of optically connecting peripheral cards and devices within the system to an OPCB backplane or motherboard in a pluggable manner. However the large mechanical misalignment tolerances between connecting cards and devices inherent to such systems are contrasted by the small sizes of optical waveguides required to support optical communication at the speeds defined by prevailing communication protocols. An innovative approach is therefore required to decouple the contrasting mechanical tolerances in the electrical and optical domains in the system in order to enable reliable pluggable optical connectivity.

This thesis presents the design, development and characterisation of a suite of new optical waveguide connector interface solutions for electro-optical printed circuit boards (OPCBs) based on embedded planar polymer waveguides and planar glass waveguides. The technologies described include waveguide receptacles allowing parallel fibre connectors to be connected directly to OPCB embedded planar waveguides and board-to-board connectors with embedded parallel optical transceivers allowing daughtercards to be orthogonally connected to an OPCB backplane.

For OPCBs based on embedded planar polymer waveguides and embedded planar glass waveguides, a complete demonstration platform was designed and developed to evaluate the connector interfaces and the associated embedded optical interconnect.

Furthermore a large portfolio of intellectual property comprising 19 patents and patent applications was generated during the course of this study, spanning the field of OPCBs, optical waveguides, optical connectors, optical assembly and system embedded optical interconnects.

# LIST OF CONTENTS

1 Introduction ... 1
1.1 Background ... 1
1.2 Motivation and aim ... 3
1.2.1 Chapter 1: Introduction ... 4
1.2.2 Chapter 2: System embedded optical interconnect technology survey ... 4
1.2.3 Chapter 3: First generation pluggable active optical circuit board connector for polymer waveguide based optical circuit boards ... 5
1.2.4 Chapter 4: Second generation pluggable active optical circuit board connector for polymer waveguide based electro-optical circuit boards ... 7
1.2.5 Chapter 5: Pluggable passive optical circuit board connector for planar glass waveguide based optical circuit boards ... 7
1.2.6 Chapter 6: Intellectual Property ... 10
1.2.7 Chapter 7: Conclusion ... 10
1.3 Optical interconnect migration in data centres ... 11
1.3.1 Data centre overview ... 12
1.3.2 Data storage system interconnect topologies ... 13
1.3.3 Data centre switching architectures ... 14
1.3.4 Resource disaggregation ... 17
1.3.5 Optical interconnect migration ... 19
1.4 Data centre network communication protocols ... 20
1.4.1 Serial Attached SCSI ... 20
1.4.2 Infiniband ... 21
1.4.3 PCI Express ... 22
1.4.4 Fibre Channel ... 23
1.4.5 Ethernet ... 23
1.4.6 Network Convergence ... 25
1.5 International research projects targeting system embedded optical data communications ... 26
1.5.1 SEPIANet project ... 26
1.5.2 PhoxTroT project ... 27
1.5.3 Nephele project ... 29
1.5.4 HDPuG Optoelectronics project ... 30
1.6 Summary ... 31
2 System embedded optical interconnect technology survey ... 32
2.1 Introduction ... 33
2.2 System embedded optical interconnect technologies ... 35
2.2.1 Host side pluggable transceivers ... 35
2.2.2 Passive optical connector interface ... 36
2.2.3 Board-mountable optical transceivers ... 39
2.3 Performance limiters on high frequency electronic channels ... 43

2.3.1 Overview ....43
2.3.2 Crosstalk ....44
2.3.3 Impedance mismatch ....45
2.3.4 Signal attenuation ....46
2.4 Board-level optical interconnect ....47
2.4.2 Milestones in polymer OPCB research and development ....52
2.5 Design attributes of embedded optical channels ....56
2.5.1 Density comparison between printed copper and polymer waveguides ....56
2.5.2 Bend radius routing constrictions ....60
2.5.3 In-plane vs out-of-plane optical coupling to OPCBs ....62
2.6 Summary ....62
3 First generation pluggable active optical circuit board connector for polymer waveguide based optical circuit boards ....64
3.1 Introduction ....64
3.1.1 Objective ....65
3.1.2 Methodology ....65
3.2 OPCB connectivity in data storage enclosures ....67
3.2.1 Electro-optical midplane architecture ....67
3.2.2 Midplane interconnect requirements ....68
3.2.3 Electro-optical PCB connection scheme ....69
3.2.4 Electro-optical PCB connection concept ....70
3.3 Electro-optical PCB demonstration platforms ....72
3.4 Storlite – first generation passive optical backplane system with active pluggable connectors ....73
3.4.1 Storlite project summary ....73
3.4.2 System design overview ....74
3.4.3 StorConn2 - Storlite quad 10 Gb/s parallel optical transceiver ....75
3.4.4 Storlite connector mechanism ....93
3.4.5 StorConnOpt2 - Storlite optical waveguide backplane ....98
3.4.6 Storlite optical waveguide receptacle ....102
3.4.7 Passive optical assembly process ....104
3.4.8 StorConnTest2 - Storlite 10 Gb/s test daughtercard ....116
3.4.9 Storlite platform firmware design ....125
3.4.10 StorConnPwr2 – passive electrical backplane ....131
3.4.11 Storlite demonstration platform ....133
3.4.12 Test and measurement ....135
3.5 Summary ....151
4 Second generation pluggable active optical circuit board connector for polymer waveguide based electro-optical circuit boards ....152
4.1 Introduction ....152
4.1.1 FirstLight project summary ....152
4.1.2 FirstLight system design ....153
4.2 FirstLight pluggable active electro-optical backplane connector ....155
4.2.1 StorConn3 – FirstLight quad parallel optical transceiver ....157

4.2.2 StorConn3 circuit hardware design ....161
4.2.3 Dual lens coupling solution ....163
4.2.4 StorConn3 component layout ....166
4.2.5 StorConn3 card assembly process ....169
4.2.6 Pluggable connector mechanism ....174
4.3 StorConnOpt3 - electro-optical backplane with embedded polymeric waveguides ....177
4.3.1 Polymer optical waveguide layer ....177
4.3.2 Optical interconnect design ....178
4.3.3 Electro-optical PCB fabrication process ....180
4.3.4 Insertion loss measurements ....185
4.3.5 Optical waveguide signal integrity characterisation ....189
4.4 Passive alignment and assembly method ....191
4.4.1 Improvements over Storlite design ....191
4.4.2 Fabrication of passive alignment features ....192
4.4.3 Lensed waveguide receptacle design ....196
4.4.4 Assembly of microlens array into receptacle ....199
4.4.5 Assembly of waveguide receptacle onto OPCB ....199
4.5 StorConnTest3 – FirstLight 10 Gb/s test daughtercard ....204
4.5.1 Functional overview ....205
4.5.2 StorConnTest3 component listing ....207
4.5.3 FirstLight demonstration platform ....211
4.5.4 FirstLight platform FPGA firmware design ....212
4.5.5 FirstLight Graphical User Interface (GUI) ....213
4.5.6 FirstLight test and measurement results ....216
4.6 Summary ....222
5 Pluggable passive optical circuit board connector for planar glass waveguide based optical circuit boards 224
5.1 Introduction ....224
5.1.1 Glass waveguide based optical circuit boards ....225
5.1.2 SEPIANet project ....225
5.1.3 SEPIANet technology overview ....225
5.1.4 Objectives ....226
5.2 Optical waveguide layout design ....227
5.3 Glass waveguide fabrication ....231
5.3.1 Glass waveguide fabrication process flow ....231
5.3.2 Waveguide characterisation by Fraunhofer IZM ....237
5.4 Electro-optical circuit board fabrication using low temperature lamination processes ....241
5.5 Pluggable optical connector design and assembly ....244
5.5.1 MT ferrule compliant waveguide receptacle ....245
5.5.2 Board-to-board connector plug and backplane receptacle design ....247
5.5.3 Board assembly of fibre ferrule receptacles ....250
5.5.4 Assembly of connector plugs and receptacles ....252
5.5.5 SEPIANet platform design ....254

5.6 Optical interconnect test and measurement ....259
5.6.1 Bidirectional insertion loss measurements ....259
5.6.2 Insertion loss measurements through both daisy chained demonstration platforms ....262
5.6.3 Correlation between contamination and non-reciprocal insertion loss profiles ....266
5.6.4 Bidirectional signal integrity measurements at 10.3 Gb/s ....268
5.6.5 Bit Error Rate characterisation of on-board 850 nm transceiver subassemblies ....269
5.6.6 Bit Error Rate characterisation up to 32 Gb/s ....270
5.7 Summary ....272
6 Intellectual Property ....273
6.1 Introduction ....273
6.2 Optical intellectual property summary ....275
6.3 Areas of invention ....278
6.4 Optical intellectual property portfolio ....279
6.4.1 Patent 1: Optical transceiver connector module [170] ....279
6.4.2 Patent 2: High-precision passive component alignment to PCB [171] ....282
6.4.3 Patent 3: Daughtercard optical waveguide interconnects [172] ....287
6.4.4 Patent 4: Reconfigurable optical printed circuit board [173] ....289
6.4.5 Patent 5: Crosstalk suppression on optical waveguides [174] ....292
6.4.6 Patent 6: Passive optical waveguide integrated CWDM device [175] ....294
6.4.7 Patent 7: Method of manufacturing multimode CWDM multiplexer / demultiplexer structures [176] 296
6.4.8 Patent 8: Optical adapter module [177] ....297
6.4.9 Patent 9: Fabricating a hybrid electro-optical PCB with optical surface layers [178] ....299
6.4.10 Patent 10: Lossless tapered waveguides [179] ....302
6.4.11 Patent 11: Directly pluggable orthogonal in-plane optical PCB connector [180] ....304
6.4.12 Patent 12: Optical capstan to simplify routing of optical waveguides [181] ....306
6.4.13 Patent 13: Directly pluggable orthogonal in-plane optical PCB connector [182] ....308
6.4.14 Patent 14: Modular interconnectable electro-optical backplane [183] ....310
6.4.15 Patent 15: Optical amplification devices for polymer optical printed circuit boards [184] ....312
6.4.16 Patent 16: Design and manufacture of polymer optical waveguide amplification structures [185] 314
6.5 Summary ....315
7 Conclusion ....316
7.1 Summary ....316
7.1.1 Introduction to data centres and system embedded photonic eco-system ....316
7.1.2 Pluggable active optical connectors for polymer waveguide interfaces ....317
7.1.3 Pluggable passive optical connectors for glass interfaces ....317
7.1.4 Intellectual property ....318
7.2 Challenges and future work ....318
7.2.1 Higher density active connectors ....318
7.2.2 Standardisation of polymer waveguide connectorisation ....319
7.2.3 Connectorisation of single mode polymer waveguides ....320
7.2.4 Method of futher improving robustness of passive alignment technique ....321

7.2.5 Multilayer glass waveguide couplers ....... 323
7.2.6 High yield alignment and assembly processes on glass OPCBs ....... 323
7.2.7 Wavelength dependent performance ....... 323
7.3 Publications ....... 324
7.3.1 Journal papers ....... 324
7.3.2 Conference papers ....... 325
7.3.3 Articles ....... 326
7.4 Invited talks and presentations ....... 327
8 References ....... 331
9 Appendix - FirstLight platform firmware coding ....... 349
9.1 PCI command structure ....... 349
9.1.1 General static control lines ....... 349
9.1.2 XFP $I^2C$ interface (XFPIO) ....... 353
9.1.3 Crosspoint switch interface (CPSIO) ....... 354
9.1.4 StorConn3 transceiver interface (SC3IO) ....... 359
9.1.5 LED control interface (LEDIO) ....... 365

# LIST OF FIGURES

Figure 1-1: Optical circuit board varieties and associated demonstration platforms ..... 4
Figure 1-2: SEPIANet glass waveguide optical circuit board connector platform..... 9
Figure 1-3: Optical interconnect migration in data centre environments ..... 11
Figure 1-4: Floor plan of a single storey data centre (Source: NTT Japan) ..... 12
Figure 1-5: a) Generic data storage array system, b) Data storage integrated application platform, c) 42U data storage rack ..... 13
Figure 1-6: Example of a Dual Star interconnect topology whereby 2 controller nodes each have a link to 16 storage devices, thus eliminating the controller module as a single point of failure ..... 14
Figure 1-7: Schematic view of traditional fat tree with links becoming thicker as we move up the hierarchy ..... 15
Figure 1-8: Fat tree network where all three levels are built out of uniform rate port commodity TOR switches in a (folded) Clos topology (Source: Nephele project [11]) ..... 16
Figure 1-9: a) Building blocks of modern data centres, b) Schematic of optical links connecting geographically dispersed subsystems within a data centre architecture ..... 18
Figure 1-10: Serial Attached SCSI Roadmap (Source SCSI Trade Association – March 2014) ..... 21
Figure 1-11: Infiniband™ Roadmap (Source: InfiniBand™ Trade Association) ..... 22
Figure 1-12: Ethernet roadmap [29] ..... 24
Figure 1-13: SEPIANet project technology overview ..... 27
Figure 1-14: PhoxTrot vision: On-chip, chip-to-chip, board-to-board and rack-to-rack optical interconnect solutions deployed in PhoxTrot demonstration platform ..... 28
Figure 1-15: Multimode optical interconnect deployed in systems, chips and OPCBs ..... 29
Figure 1-16: Nephele hybrid electro-optical scale out data centre architecture ..... 30
Figure 1-17: HDPuG demonstrators a) 1.4 metre spiral polymer waveguide sample (source: Dow Corning), b) MT ferrule terminated optical waveguides with visible pink and green light illuminating 2 waveguides (Source: Optical Interlinks), c) OPCB with straight waveguide illuminated with green light (source: TTM Technologies), d) Two MT ferrules assembled onto the edge of a polymer waveguide based OPCB (Source: Optical Interlinks) ..... 31
Figure 2-1: Generic data storage enclosure with passive and active optical connectors ..... 35
Figure 2-2: MT parallel optical interface showing the positions of six rows of 12 small circular fibre channels relative to two large circular mechanical alignment slots on the left and right hand side ..... 37
Figure 2-3: a) MT ferrule, b) PrizmMT™ ferrule, c) MXC connector ..... 38
Figure 2-4: PCB stack-up showing electrical signal propagation between a copper trace on the top layer and a copper trace on an intermediate layer, passing along part of a via. A) fully filled via causes the exposed unused section to create localised impedance mismatch, which gives rise to a partial reflection, b) back-drilling of the via to minimise discontinuity and therefore partial reflection ..... 45
Figure 2-5: Schematic view of electro-optical midplane architecture comprising copper layers to accommodate power, static and low speed signal distribution and at least one optical transmission layer to convey high speed signals along optical waveguides ..... 47
Figure 2-6: Electro-optical PCB interconnect technologies a) Fibre-optic flexible circuits, b) PCB embedded polymer waveguides, c) PCB embedded planar glass waveguides (Source: Fraunhofer IZM) ..... 48
Figure 2-7: Scanning Electron Microscope image of exposed 50µm x 50µm polymer waveguide core without upper cladding (left) and array of exposed parallel polymer waveguide cores (right) ..... 50
Figure 2-8: Male MT adapter passively aligned and assembled onto an OPCB (Source: IBM Research [107]) ..... 52
Figure 2-9: Optical transceiver connectivity to OPCBs a) active optical transceiver interface in a fixed butt-coupled arrangement to embedded polymer waveguide interfaces in OPCBs (Source: IBM Research [108]), b) active optical transceiver and pluggable connector mechanism developed by author and described in Chapter 4, allowing optical transceiver interface to be selectively butt-coupled to embedded polymer waveguide interfaces on OPCB during engagement of a daughter card ..... 53
Figure 2-10: Channel density comparison between 24 Gb/s copper differential trace and polymeric optical channels ..... 57
Figure 2-11: Waveguide crossovers a) Schematic view of waveguide crossover, b) Photo of 50µm x 50µm polymer waveguides patterned photolithographically with multiple 90° crossovers (Source: Exxelis Ltd), c) Photo of 50µm x 50µm polymer waveguides patterned photolithographically with multiple 60° crossovers (Source: Exxelis Ltd) ..... 59
Figure 2-12: Waveguide splitters a) Schematic view of 1 x 2 waveguide splitter, b) Photo of 50µm x 50µm polymer waveguides with 1 x 2 symmetric Y-branch splitter (Source: IBM Research), c) Photo of 50µm x 50µm polymer waveguides with 1 x 2 splitter (Source: Exxelis Ltd) ..... 60
Figure 2-13: Photo of polymer waveguide 90° arc bend illuminated with 650 nm visible light ..... 62

Figure 3-1: Schematic view of electro-optical midplane architecture comprising copper layers to accommodate power, static and low speed signal distribution and at least one optical transmission layer to convey high speed signals along optical waveguides ....67
Figure 3-2: Data storage array with 24 drives, two controller modules (one shown) and two power supply modules (one shown) ....68
Figure 3-3: Schematic view of line cards with optical edge connectors plugging in orthogonally to an electro-optical backplane with embedded optical interconnect ....69
Figure 3-4: Butt-coupled optical connection between surface emitting and receiving devices on the line card and planar optical waveguide interface on electro-optic backplane ....70
Figure 3-5: a) Storlite Demonstrator - First generation optical backplane connector platform, b) FirstLight Demonstrator - second generation optical backplane connector platform ....72
Figure 3-6: Storlite optical backplane connection scheme ....74
Figure 3-7: Storlite parallel optical transceiver circuit ....76
Figure 3-8: Storlite parallel optical transceiver side-view: a) MT pin holder, b) Opto-mechanical interface, c) Heat sink, d) Flexible bridge, e) GigArray® high speed connector ....77
Figure 3-9: Functional diagram of quad transceiver circuit ....78
Figure 3-10: StorConn2 raw card outline dimensions ....79
Figure 3-11: StorConn2 PCB layer stack-up ....80
Figure 3-12: Optical interface a) Quad VCSEL and PIN photodiode die array aligned and assembled onto PCB, b) GRIN micro-lens array in ceramic holder aligned over active areas of VCSEL and PIN photodiodes ....81
Figure 3-13: Physical drawings of Gradient Index Micro-lens array and custom ceramic holder ....83
Figure 3-14: 6x12 MT ferrule interface. The lowest (or highest) row of optical channels is offset from the datum of the MT pins by 625 µm ....83
Figure 3-15: ZEMAX simulation provided by GRINTech of optimum GRIN lens length (ZL) required to enable: a) imaging of the 75 µm x 75 µm waveguide output into the 75 µm circular aperture of the PIN photodiode and b) imaging of the VCSEL output into the waveguide ....85
Figure 3-16: StorConn2 component layout: a) top view, b) side view ....88
Figure 3-17: Gig array receptacle (*Source: FCI Interconnect - GS-20-016 Rev B)* ....89
Figure 3-18: Equaliser Chip – 10.3 Gb/s ....90
Figure 3-19: MT pin block ....91
Figure 3-20: Photo of VCSEL array and photodiode assembly prior to wire bonding (photo through microscope) ....92
Figure 3-21: GRIN lens assembly (photo through microscope) ....93
Figure 3-22: Storlite connector CAD model - side view a) optical interface retracted, b) optical interface elevated ....94
Figure 3-23: Photo of Storlite connector - side view a) photonic interface retracted, b) photonic interface elevated ....95
Figure 3-24: a) Storlite connector with retracted optical interface, b) Storlite connector bottom view showing bifurcated pivoted lever mechanism and electronic high speed GigArray® connector ....96
Figure 3-25: Storlite connector exploded view ....96
Figure 3-26: Storlite connector attached to line-card - front view ....97
Figure 3-27: Storlite optical waveguide PCB: a) four groups of 12 parallel straight waveguides. MT compliant receptacle with MT pins slots offset from waveguide array by 625 µm (centre to centre), b) Single waveguide illuminated with 635 nm visible light to demonstrate effectiveness of passive alignment method and devices. ...99
Figure 3-28: Polymeric optical PCB interface ....100
Figure 3-29: Fabrication process for Storlite optical waveguide PCB ....101
Figure 3-30: Mechanical drawings of Storlite waveguide receptacle ....103
Figure 3-31: Diagrams from patent "Optical printed circuit board and manufacturing method" [144]: a) optical waveguide receptacle socket showing exposed registration stubs, b) optical waveguide receptacle assembled into optical waveguide receptacle socket ....105
Figure 3-32: Passive self-alignment assembly procedure: a) Storlite waveguide receptacle and OPCB waveguide interface with mechanical recesses and exposed registration structures, b) Waveguide receptacle assembled onto OPCB (view facing waveguide interface), c) Waveguide receptacle assembled onto OPCB (side view showing front face of unit flush with waveguide interface) ....107
Figure 3-33: Shadowgraph image of an optical waveguide group with a receptacle passively aligned over it using the mechanical registration features fabricated in the peripheries of the 12 communication waveguides ....108
Figure 3-34: First stage alignment verification using visible light: a) Class 2 635 nm laser source is connected to one fibre of a lowest row populated 6x12 MT to 12 LC fibre jumper, b) view of lowest row populated 6x12 MT interface with channel five illuminated (fifth channel from left), c) MT ferrule about to engage with optical waveguide receptacle, d) MT ferrule fully engaged with optical waveguide receptacle, e) optical waveguide receptacle passively aligned into waveguide registration recess with one waveguide channel illuminated ....110

Figure 3-35: Contour maps of relative insertion loss compared to the optimum coupling position for a) VCSEL misalignment at z = 0, b) PD misalignment [2] ....111
Figure 3-36: SEM images of waveguide array interface which has been diced and polished: a) Full view of array of 12 communications waveguides and peripheral recesses in upper cladding layer to enable mechanical access to mechanical registration features, b) Close view of roughness and surface degradation over waveguide interface....114
Figure 3-37: optical backplane assembly: a) Exploded view of optical backplane assembly and parts including the receptacle holder, b) optical waveguide interface ....115
Figure 3-38: Storlite test daughtercard: a) GigArray® connector receptacle, b) microcontroller, c) four commercial XFP transceivers, d) RS232 connector and transceiver....116
Figure 3-39: StorConnTest2 functional diagram....118
Figure 3-40: StorConnTest2 card profile ....119
Figure 3-41: StorConnTest2 raw card stack-up ....120
Figure 3-42: a) Schematic showing StorConnTest2 component layout, b) photo of StorConnTest2 ....124
Figure 3-43: StorConnTest2 microcontroller flow diagram....125
Figure 3-44: Main screen / comm control tab ....129
Figure 3-45: Optical control tab....130
Figure 3-46: XFP control tab ....131
Figure 3-47: StorConnPwr2 physical layout ....132
Figure 3-48: StorConnPwr2 PCB layup....132
Figure 3-49: Assembly process for Storlite demonstration enclosure: a) unpopulated enclosure, b) Enclosure populated with only StorConnPwr2, c) enclosure populated with StorConnPwr2 and StorConnTest2, d) fully populated enclosure....133
Figure 3-50: Storlite demonstration platform a) Passive optical waveguide backplane (copper clad side shown), b) Pluggable optical backplane connectors, c) Electrical backplane, d) two test daughtercards ....134
Figure 3-51: a) View of Storlite demonstration system from the back, b) Single waveguide illuminated with 635 nm visible light ....135
Figure 3-52: Screen shots from Xyratex proprietary 10GbE LAN test traffic generator and error counter ....136
Figure 3-53: a) Storlite transceiver opto-mechanical interface with all VCSELs activated, b) Storlite transceiver opto-mechanical interface with MT patchcord attached for stand-alone testing....138
Figure 3-54: Storlite test line card (powered) with transceiver attached (without connector housing) for stand-alone evaluation of the transceiver with an MT patchcord ....139
Figure 3-55: Eye diagrams of all four VCSEL channels driven at 10.3 Gb/s with a bias current of 11.91 mA and a modulation current of 9.8 mA. ....141
Figure 3-56: Eye diagrams of 10.3 Gb/s optical signals after propagation through 12 waveguides under test. a) Eye diagrams captured without index damping fluid applied on waveguide end facets, b) Eye diagrams captured with index damping fluid applied on both ingress and egress waveguide end facets ....143
Figure 3-57: a) Total jitter of 10.3 Gb/s signals measured across 12 waveguides without index damping fluid (blue graph) and with index damping fluid (red graph), b) Comparative average optical power measured at the communications signal analyser across 12 waveguides without index damping fluid (blue graph) and with index damping fluid (red graph) ....144
Figure 3-58: View of exposed optical backplane waveguide egress interface with Storlite connector attached to ingress interface. a) VCSEL channel one (furthest left) activated and 850 nm light exiting single egress waveguide, b) VCSEL channels one and four activated and 850 nm light exiting two egress waveguides 750 µm apart ....146
Figure 3-59: Characterisation setup for Storlite optical interface connection to optical backplane ....147
Figure 3-60: Relative loss / dB and total jitter / UI for each use case listed in Table 3-9. ....148
Figure 3-61: Eye diagrams for four use cases: a) Storlite VCSEL is connected directly to communications signal analyser via MT fibre-optic patchcord (Reference), b) VCSEL connected to waveguide with isopropanol applied to ingress and egress end facets, c) VCSEL connected to waveguide with end facets diced and polished, d) VCSEL connected to waveguide with end facets only diced ....149
Figure 3-62: Storlite demonstration platform populated with two test line cards with proprietary pluggable connectors plugged into a passive optical backplane and a passive electrical backplane. ....150
Figure 4-1: a) FirstLight demonstration enclosure fully populated, b) FirstLight demonstration enclosure powered....155
Figure 4-2: FirstLight pluggable optical PCB connector concept. Pluggable active connector modules are inserted into the top of an optical backplane and engage with an optical layer on the bottom side of the backplane: (a) Electro-optical backplane connection scheme; active pluggable connectors housed on the edge of peripheral line cards engage with the embedded optical layer in the backplane PCB. ....156
Figure 4-3: Functional diagram of StorConn3 ....158

Figure 4-4: Parallel optical transceiver circuit and optical interface a) Front ramped plug, b) Optical platform guide pins, c) Optical interface, d) Flexible bridge, e) Rigid base section, f) Microcontroller, g) MT compatible alignment pins, h) Micro-lens array, i) Quad VCSEL array, j) Quad PIN photodiode array, k) Quad VCSEL Driver array, l) Quad transimpedance amplifier/limiting amplifier TIA/LA array ....159
Figure 4-5: StorConn3 circuit board form factor ....161
Figure 4-6: StorConn3 PCB layer stack-up....162
Figure 4-7: a) Model and simulation results of optimum dual lens arrangement to image VCSEL to multimode ingress waveguide end facet predicting 0.72 dB loss in this transition, b) Model and simulation results for optimum dual lens arrangement to image egress waveguide end facet onto PD active area predicting 1.11 dB in this direction. (Source: Omron)....164
Figure 4-8: Dual lens coupling interface between transceiver and OPCB ....165
Figure 4-9: StorConn3 component layout: a) MT compatible alignment pins, b) lead frame holding the VCSEL array and PIN photodiode array, c) quad VCSEL driver array IC, d) transimpedance amplifier/limiting amplifier TIA/LA array IC, e) microcontroller, f) rectangular slot in PCB for copper slug, g) adaptive signal equalisers, h) electrical array connector ....168
Figure 4-10: MT pin holder and ramped plug part....170
Figure 4-11: Custom heatsink: a) heatsink design, b) optical head section prior to heatsink attachment, c) optical head section with heatsink attached ....170
Figure 4-12: Zarlink lead frame holding VCSEL array and PIN photodiode array ....171
Figure 4-13: Lead frame alignment: a) optical head with two MT pins protruding through the PCB, b) lead frame slots aligned over the MT pins, c) lead frame fully descended into place. During actual assembly the VCSEL driver and PIN photodiode array ICs will already have been secured by this time....172
Figure 4-14: Dimensioned drawings of Omron P1L lens array (Source: Omron) ....173
Figure 4-15: Optical interface section with Omron microlens assembled ....173
Figure 4-16: a) Optical transceiver circuit mounted on flexi-rigid substrate, b) Connector module housing, c) Grooves to enable required movement of optical interface during mating process. ....175
Figure 4-17: a) Optical connector during first stage of coarse engagement with OPCB connector receptacle, b) Optical connector fully plugged into OPCB connector receptacle....176
Figure 4-18: Schematic illustrating waveguide embedded in lower and upper claddings (Source: IBM Research – Zürich) ....177
Figure 4-19: a) Waveguide routing pattern with four optical engagement apertures for the FirstLight demonstrator, b) Negative cascading waveguide, c) Positive cascading waveguide ....179
Figure 4-20: Top view micrographs showing two waveguides intersecting with (a) a crossing angle of 90$^o$ and (b) a crossing angle of 145$^o$ (Source: IBM Research)....180
Figure 4-21: a) Schematic of optical surface scan performed on an OPCB containing up to ten laminated copper layers. b) Experimental results of optical surface scan showing height variations up to ±50 µm. (Source: IBM Research – Zürich)....181
Figure 4-22: Schematic depiction of polymer layer deposition using doctor blade method. (Source: IBM Research – Zürich)....181
Figure 4-23: Separate flexible polyimide substrate with optical polymer layers deposited and patterned. (Source: IBM Research – Zürich) ....182
Figure 4-24: a) Waveguides measured by microscopic top view fulfil width specification of 70 µm within ± 2 µm. b) Waveguide heights experimentally determined by optical surface scanning satisfy target height of 70 µm within ± 2 µm. (Source: IBM Research – Zürich) ....183
Figure 4-25: StorConnOpt3 electro-optical backplane prior to slots being milled out ....184
Figure 4-26: StorConnOpt3 electro-optical backplane with optical connector slots milled out and electronic CompactPCI connectors populated. A close-up view of a connector aperture is shown with a single curved waveguide and its egress point illuminated with 635 nm visible light....184
Figure 4-27: Far field pattern of launch fibre with and without mode filtering (Source: UCL)....185
Figure 4-28: Predicted and measured values for insertion loss on a group of nine waveguides on the OPCB. The dashed line shows the 9.62 dB threshold ....186
Figure 4-29: Use of large area thin silicon photodetector to fit into the mid board waveguide connector interface slots in order to directly measure optical output from waveguides ....187
Figure 4-30: Process to deposit smooth thin film over waveguide end facets in order to reduce coupling loss [155]....188
Figure 4-31: Top view of OPCB backplane with one waveguide illuminated with 650 nm light ....189
Figure 4-32: Requirements specification for FirstLight waveguide receptacle: a) waveguide receptacle with 1x12 MT compliant interface, b) waveguide array with self-alignment features and PCB machined recesses, c) waveguide receptacle passively aligned over waveguides using self-alignment features....192
Figure 4-33: Fabrication process for the passive alignment features on the optical layer....193

Figure 4-34: View of the PCB layout of the FirstLight waveguide interface using the Cadence Allegro PCB Layout software tool. ....194
Figure 4-35: Photographs of upper cladding opening showing waveguides for signal transmission as well as connector alignment features (dummy waveguides). (Source: IBM Research – Zürich) ....195
Figure 4-36: (a) Clearance areas in upper cladding layer to provide mechanical access for optical connector receptacle and waveguide end facet ....195
Figure 4-37: Lensed waveguide receptacle front views: a) trimmed receptacle without lens array holder section, b) complete receptacle, c) complete receptacle with Omron microlens array integrated and receptacle assembled onto OPCB ....197
Figure 4-38: Dimensioned drawing of waveguide receptacle ....198
Figure 4-39: Photo of complete micro-lens receptacle assembly next to British ten pence coin ....198
Figure 4-40: Mechanical lens assembly jig A designed to align the Omron microlens array into the custom receptacle a) open position, b) engaged position ....199
Figure 4-41: a) OPCB waveguide interface without receptacle, b) OPCB waveguide interface with receptacle passively aligned, but not glued in ....201
Figure 4-42: Receptacle assembly jig B required to register lensed receptacle with waveguide alignment features and hold in place during adhesive curing process. ....202
Figure 4-43: a) Assembly of primary receptacle and secondary receptacle in optical engagement slots, b) Pluggable optical connector prior to engagement with OPCB. Parallel MT compliant optical interface is visible as connector starts to engage with connector receptacle ....203
Figure 4-44: Demonstration assembly with one test card inserted. a) Electro-optical backplane, b) FirstLight connector plugged into backplane receptacle, c) PCI bridge chip next to CompactPCI connector connected to electro-optical backplane providing electrical power and low speed electronic signal interface, d) FPGA to regulate test card, e) crosspoint switch, f) commercial XFPs ....205
Figure 4-45: StorConnTest3 functional diagram ....206
Figure 4-46: StorConnTest3 board dimensions and component layout: a) CompactPCI connector plug, b) high speed electronic array connector, c) PCI bridge chip, d) FPGA to regulate test card, e) crosspoint switch, f) four commercial XFP cages, g) linear voltage regulators to provide different voltage supplies to card components, h) four banks of indicator LEDs for each XFP ....208
Figure 4-47: StorConntest3 card with StorConn3 transceiver circuit (without connector module): a) CompactPCI connector plug, b) StorConn3 transceiver circuit, c) PCI bridge chip, d) FPGA to regulate test card, e) crosspoint switch, f) four commercial XFP cages, g) linear voltage regulators to provide different voltage supplies to card components, h) four banks of indicator LEDs for each XFP ....209
Figure 4-48: StorConnTest3 raw card stack-up ....210
Figure 4-49: FirstLight demonstration platform fully populated with all four test line cards and powered up....211
Figure 4-50: Compact PCI slot population by singleboard computer motherboard and four StorConnTest3 daughtercards. Slots C, E and G are unpopulated as the daughtercards require a separation of at least 2 slots ...212
Figure 4-51: FirstLight GUI overview ....213
Figure 4-52: Characterisation of optical jitter from VCSEL channel one operating 10.3 Gb/s with modulation current ....216
Figure 4-53: FirstLight demonstration platform with optical test cables attached to XFPs ....217
Figure 4-54: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel one ....218
Figure 4-55: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel two ....218
Figure 4-56: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel three ....219
Figure 4-57: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel four....219
Figure 4-58: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel five ....220
Figure 4-59: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel six....220
Figure 4-60: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel seven ....221
Figure 4-61: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel eight ....221
Figure 5-1: a) SEPPLANE 1 glass panel dimensions, b) SEPPLANE2 glass panel and gap dimensions ....227
Figure 5-2: Schematic layouts and photos of fully populated OPCB backplane variants: a) SEPPLANE1 with single glass waveguide panel and connector layout, b) iterated waveguide crossover section, c) top view of

embedded glass panel showing single waveguide illuminated, d) SEPPLANE2 with dual glass waveguide panel and connector layout ....228
Figure 5-3: SEPPLANE1 with waveguide groups G1 to G8, test card slots (vertical) 1 – 5 and connector positions A to D. Individual waveguides in each group are always numbered 1 – 12 going from top to bottom in the horizontal component section of the waveguide group ....229
Figure 5-4: SEPPLANE2 with waveguide groups G9 to G16, test card slots (vertical) 1 – 5 and connector positions A to D. Individual waveguides in each group are always numbered 1 – 12 going from top to bottom in the horizontal component section of the waveguide group ....230
Figure 5-5: Change of optical interface dimensions for non-orthogonal waveguide - interface intersections ....231
Figure 5-6: Glass waveguide fabrication process flow (Source: Fraunhofer IZM) ....232
Figure 5-7: a) Cross-section of glass waveguide sample with thickness of 500 µm and centre-to-centre waveguide pitch of 250 µm. b) Refractive index profile at measurement wavelength of 678 nm as measured by a refractive near field (RNF) scan (Source: Fraunhofer IZM) ....234
Figure 5-8: Layout diffusion mask on glass containing waveguide layout, partner logos and alignment marks (Source: Fraunhofer IZM) ....235
Figure 5-9: a) Dip-coating, b) positioning mount and c) mask exposure with LDI (Source: Fraunhofer IZM) ...236
Figure 5-10: Aluminium diffusion mask layer on glass panel (297 x 210 mm²) (Source: Fraunhofer IZM) ....236
Figure 5-11: a) Detail of mask layer on glass, b) protected partner logos and fiducial marks after removing aluminium mask layer (Source: Fraunhofer IZM) ....237
Figure 5-12: Propagation loss and coupling loss measurements based on best fitting line of cut-back measurements (Source: Fraunhofer IZM) ....239
Figure 5-13: Glass foil laminated into PCB stack-up with conventional hot lamination technique shows bulging (Source: ILFA GmbH) ....242
Figure 5-14: a) SEPPLANE1 stack-up with one single panel optical and two electrical packages with solder mask, b) SEPPLANE2 stack-up with one dual panel optical and two electrical packages, c) Photo of OPCB cross-section, d) photo of unpopulated OPCB backplane circuit board ....243
Figure 5-15: Complete pluggable connection system comprising connector plug on connecting edge of daughtercard, backplane receptacle and fibre ferrule receptacles ....244
Figure 5-16: a) MT ferrule receptacle made of polycarbonate material, b) exposed dual glass waveguide array interfaces, c) front view of waveguide array interface with MT compliant receptacle port assembled, d) vertical offset between MT pins datum and waveguide array, e) dual ferrule mounts with MT fibre jumper plugged into one, f) four adjacent waveguide ferrule receptacles made of Ultem™ material ....246
Figure 5-17: a) Pluggable connector assembled on card edge (front view), b) rear view of pluggable connector showing fibre ribbon from internal MT ferrule, c) pluggable connector and receptacle head, d) connector plugs assembled on edge of test daughtercard, e) 4 variants of full backplane receptacle with receptacle head (black) in different orientations relative to receptacle body (clear) ....248
Figure 5-18: OPCB backplane type SEPPLANE 1 with ferrule receptacles and jumpers populated, a) schematic view of ferrule receptacle populated with MT terminated fibre ribbon, b) schematic view of dual edge board fibre jumper (connecting to 2 adjacent interfaces), c) schematic view of dual mid-board fibre jumper (connecting to 2 opposite interfaces) ....250
Figure 5-19: Ferrule receptacle assembly process steps: a) backplane suspended in ficonTEC assembly rig, b) active launch fibre alignment with 3-axis translation stage, c) active detection fibre and ferrule receptacle alignment with 5-axis translation stage, d) MT ferrule receptacle glued in place, e) fibre patch cord plugged into assembled MT multi-fibre ferrule ....252
Figure 5-20: a) SEPPLANE1 after assembly of all MT ferrule receptacles, b) SEPPLANE1 after assembly of all backplane receptacles ....253
Figure 5-21: Two test and measurement platforms SEPDEM1 and SEPDEM2 ....255
Figure 5-22: SEPIANet test cards: a) SEPTEST1 test card with vacant slots for mezzanine cards and populated with 1 edge connector plug, b) SEPTEST1C mezzanine card populated with alignment brace and MT patchcord attached, c) SEPTEST2 populated with 2 edge connector plugs, d) SEPTEST3 fully populated with 4 edge connector plugs ....257
Figure 5-23: Schematic view of complete optical connector and waveguide link under test ....259
Figure 5-24: Bidirectional insertion loss on SEPPLANE1 - waveguide group 4 in both directions as recorded directly using a large area photodetector and CSA: a) at 1310 nm, b) at 850 nm. Propagation direction >>> signifies propagation from left to right along a given waveguide as viewed from the top of the backplane, while <<< indicates right to left propagation ....260
Figure 5-25: Schematic view of complete daisy chained link under test through both SEPDEM1b and SEPDEM2a ....262
Figure 5-26: a) Bidirectional insertion loss and b) jitter measurements at 1310 nm as measured through daisy-chained link over both demonstration platforms, through SEPPLANE1b G2 in propagation direction 2A -1B and through SEPPLANE2a G12 in propagation direction 3D-1D ....264

Figure 5-27: Eye diagrams measured through SEPPLANE1b G2 in propagation direction 2A-1B and through SEPPLANE2a G12 in propagation direction 3D-1D ....265
Figure 5-28: Correlation between glue contamination pattern on waveguide interfaces and bidirectional insertion loss profile before and after removal of glue contamination ....267
Figure 5-29: Eye diagrams for 10.3 Gb/s with PRBS $2^{31}$-1 test signal conveyed at 1310 nm over all 12 waveguides of group 2 of SEPPLANE1b. A B2B eye diagram is included at the bottom ....268
Figure 5-30: BER bathtub curves and eye diagrams for 850 nm optical test signals conveyed through full optical connector and waveguide link and corresponding B2B profiles. The increased degradation observed on the 32 Gb/s profile is partly due to a reduced power budget and insufficient light being received to resolve a clean eye ....271
Figure 6-1: (a) Electro-optical midplane connection scheme; active pluggable connectors housed on the edge of peripheral line cards engage with the embedded optical layer in the backplane PCB. ....280
Figure 6-2: a) Optical transceiver circuit mounted on flexi-rigid substrate, b) Connector module housing, c) Grooves to enable required movement of optical interface during mating process, d) Electro-optical midplane in FirstLight demonstration platform fully populated with all 4 test line cards and powered up, e) Connector module housed on line card and plugged into electro-optical backplane, f) 4 test line cards ....281
Figure 6-3: Fabrication process for the passive alignment features on the optical layer ....283
Figure 6-4: Diagram from Pitwon precision alignment patent [144] ....284
Figure 6-5: Photographs of upper cladding opening on FirstLight electro-optical midplane showing waveguides for signal transmission as well as connector alignment features (dummy waveguides). ....285
Figure 6-6: a) Storlite waveguide receptacle, b) FirstLight waveguide receptacles ....286
Figure 6-7: Electro-optic daughtercard with embedded flexible optical layers ....288
Figure 6-8: Various passive and active optical components which are passively aligned into optical sockets and their associated channels through mechanical alignment features ....290
Figure 6-9: Electro-optical printed circuit board with optical sockets housing optical transceiver devices and out-of-plane deflection structures ....291
Figure 6-10: Boundary effect on leaked light with air trench ....292
Figure 6-11: Graded particulate concentration to reduce back-reflection ....293
Figure 6-12: Embedded prism style WDM multiplexer / demultiplexer structure ....295
Figure 6-13: Photolithographic mask to create WDM structures in photo-cureable polymer ....296
Figure 6-14: Optical drive adapter module ....297
Figure 6-15: Electro-optical PCB with partial optical layer on surface ....300
Figure 6-16: FirstLight electro-optical midplane with optical connector slots milled out and electronic CompactPCI connectors populated. A close-up view of a connector aperture is shown with a single curved waveguide and its egress point illuminated with 635 nm visible light. ....301
Figure 6-17: Tapered waveguide with core index variation ....303
Figure 6-18: Orthogonal optical connector with alignment rails and grooves. A free space lens coupling system allows physical contact between optical interfaces to be omitted ....305
Figure 6-19: Rotary optical brace with angle adjustable optical interface ....307
Figure 6-20: Drive dongle with optical engine incorporated ....309
Figure 6-21: Rack scale data centres with modular interconnectable electro-optical backplanes ....311
Figure 6-22: Optical amplification module arranged over an embedded polymer waveguide layer ....313
Figure 6-23: a) Optical amplification node in which optical waveguides are brought together for collective pumping, b) Method of fabricating doped section of polymer waveguides ....315
Figure 7-1: IEC polymer MT (PMT) ferrule standard for the termination of optical waveguides a) PMT connector components, b) PMT connector structure after assembly, c) photo of commercial PMT assembled onto a polymer waveguide strip ....320
Figure 7-2: Firstlight precision alignment concept deployed in singlemode polymer waveguide connectorisation – a) precision alignment concept diagram from Pitwon patent [171], b) shadowgraph photo of precision alignment of an MT compliant Storlite receptacle passively aligned to an array of multimode waveguides, c) concept diagram for silicon V-groove to singlemode polymer waveguide [145], d) photo of silicon V-groove passively aligned to singlemode polymer waveguide ....321
Figure 7-3: Excerpts from published passive alignment patent showing slanted registration guides a) straight passive alignment features will require clearance, b) slanted alignment features will require clearance however will allow MT mount to accurately self-centre by squeezing the registration guides ....322

# LIST OF TABLES

Table 1-1: Serial Attached SCSI protocol generations....20
Table 1-2: PCI Express protocol generations....22
Table 1-3: Fibre Channel protocol generations....23
Table 2-1: Strategic purchases by optical connector companies of mid-board transceiver start-ups....40
Table 2-2: Survey of commercial midboard optical engine technologies....41
Table 3-1: Storlite card designations....75
Table 3-2: Component layout and description....86
Table 3-3: Optical waveguide testbed attributes....100
Table 3-4: StorConnTest2 component listing....121
Table 3-5: StorConnTest2 User Interface ASCII Command Format....127
Table 3-6: StorConnTest2 Microcontroller ASCII Command Exchange....128
Table 3-7: StorConnPwr2 component listing....132
Table 3-8: Test conditions to estimate BER on a 10 Gb/s signal from error counts with a 99% confidence level....137
Table 3-9: Waveguide end facet characterisation....148
Table 4-1: Storlite card designations....154
Table 4-2: StorConn3 component list and description....166
Table 4-3: StorConnTest3 key component listing....207
Table 5-1: Insertion, propagation and bend losses at 1310 nm of waveguide groups in SEPPLANE1 panel dependent on bend radii and launching condition (Source: Fraunhofer IZM)....240
Table 5-2: Summary of insertion loss as measured on all backplanes with 1310 nm and 850 nm optical test signals modulated at 10.3 Gb/s. The standard deviation of each measurement set is also shown to provide an indication of the measurement spread and variability of results due to contamination....261
Table 5-3: Summary of peak to peak jitter as measured on all backplanes with 1310 nm and 850 nm optical test signals modulated at 10.3 Gb/s. The standard deviation of each measurement set is also shown to provide an indication of the measurement spread and variability of results due to contamination....269
Table 6-1: Summary of filed patents....275
Table 9-1: PCI command map - PCI write and read command set for general static control status lines....349
Table 9-2: PCI command sequence example - general IO PCI commands....351
Table 9-3: PCI command sequence - execute read of XFP I²C memory device....353
Table 9-4: PCI command sequence - read back contents of previous read instruction....353
Table 9-5: PCI command sequence - execute write to XFP I²C memory device....354
Table 9-6: PCI command sequence - CPS write sequence to configure switch....354
Table 9-7: PCI command sequence - CPS read sequence to read switch configuration....355
Table 9-8: PCI command sequence - CPS Output pre-emphasis configuration....356
Table 9-9: PCI command sequence - CPS global output level configuration....356
Table 9-10: PCI command sequence - CPS individual output level configuration....357
Table 9-11: PCI command sequence - CPS output boost mode configuration....358
Table 9-12: PCI command sequence - CPS input equalisation configuration....359
Table 9-13: PCI command sequence - execute I²C write of VCSEL Driver to turn channel 4 on....360
Table 9-14: PCI command sequence - execute I²C write of VCSEL Driver to turn all VCSELs off....361
Table 9-15: PCI command sequence - execute I²C write of VCSEL driver to turn all VCSELs on....361
Table 9-16: PCI command sequence - execute I²C read of VCSEL Driver....361
Table 9-17: PCI command sequence - execute I²C write of optical driver (TIA) to turn all photodiode channels on....362
Table 9-18: PCI command sequence - execute I²C write of optical driver (TIA) to turn all photodiode channels off....363
Table 9-19: PCI command sequence - execute I²C read of optical receiver....363
Table 9-20: PCI command sequence - Execute I²C read of VCSEL driver temperature....363
Table 9-21: PCI command sequence - execute I²C read of VCSEL driver monitor current....364
Table 9-22: PCI command map - LED control interface....365

# ABBREVIATIONS AND ACRONYMS

| | |
|---|---|
| **AOC** | Active Optical Cable |
| **API** | Application Programming Interface |
| **AWG** | Arrayed Waveguide Grating |
| **BER** | Bit Error Rate |
| **BERT** | Bit Error Rate Tester |
| **BGA** | Ball Grid Array |
| **CDR** | Clock Data Recovery |
| **CMOS** | Complementary Metal-Oxide-Semiconductor |
| **CTE** | Thermal Expansion Coefficient |
| **DEMUX** | Demultiplexer |
| **DFB** | Distributed Feedback Laser |
| **DWDM** | Dense Wavelength Division Multiplexing |
| **HPC** | High Performance Computing |
| **I/O** | Input / Output |
| **IC** | Integrated Circuit |
| **LA** | Limiting Amplifier |
| **MM** | Multimode |
| **MT** | Mechanical Transfer |
| **MUX** | Multiplexer |
| **NA** | Numerical Aperture |
| **NRZ** | Non Return to Zero |
| **OPCB** | Optical Printed Circuit Board or (Electro-optical Circuit Board) |
| **PAM4** | Pulse-Amplitude Modulation with 4 signal levels |
| **PCI** | Peripheral Computer Interface |
| **PCB** | Printed Circuit Board |
| **PD** | Photodiode |
| **PIN** | Positive intrinsic negative (photodiode) |
| **PRBS** | Pseudo Random Binary Sequence |
| **QSFP** | Quad Small Form-Factor Pluggable |
| **SAS** | Serial Attached SCSI |
| **SATA** | Serial ATA |
| **SCSI** | Small Computer System Interface |
| **SFP** | Small Form Factor Pluggable |
| **TIA** | Trans-Impedance Amplifier |
| **TCO** | Total Cost of Ownership |
| **UI** | Unit Interval |
| **VCSEL** | Vertical Cavity Surface Emitting Laser |
| **WDM** | Wavelength Division Multiplexing |

# 1 INTRODUCTION

## 1.1 Background

Over the last 10 years, there has been a surge in the amount of digital information being captured, processed, stored and conveyed from one place to another, precipitated for the most part by the proliferation of mobile data devices and media communications. This exponential increase in data demand is pushing modern information and communications systems beyond their design limits and towards a crippling “data cliff”, but electro-optical printed circuit boards (OPCBs) offer a way past this bottleneck.

Printed circuit boards (PCBs) form the basis of all modern information and communication systems. The increase in data communication speeds will incur a toll on electronic systems when higher frequency electronic signals are conveyed along the copper channels used in conventional PCBs today. Dielectric absorption and resistive loss mechanisms will more strongly attenuate higher frequency signals conveyed along a copper channel, while reflections, signal skew and interference from other electronic channels will distort the data.

Furthermore the environmental effects of system operation, such as temperature and humidity, will cause changes in the circuit board substrate, thus altering the carefully balanced characteristics of the electronic channels. Many of these constraints can be mitigated to some degree, however, at an ever mounting cost to the overall system design and with an increasing power penalty.

OPCBs are PCBs, in which optical channels have been embedded or to which they have been directly attached. In future, OPCBs will almost always still require electronic layers for electrical power distribution and static control or low-speed signals, however higher frequency signals can be confined to dedicated optical communication layers, which benefit from the substantially higher data bandwidth density that optical channels enable. Therefore it is envisaged that in future OPCBs will almost always comprise both optical and electrical layers.

Embedded optical interconnect technologies, whether deployed at the cable level, circuit board level or chip level offer significant performance and power advantages over conventional electronic interconnect including higher data rates, reduced electromagnetic interference, reduction in power consumption, higher channel density and corresponding reduction in the amount of cable or PCB materials used. Therefore in order to cope with the exponential increase in capacity, processing power and bandwidth density inside information communication systems, there has been a trend over the past decade to migrate optical channels down from the higher communication tier optical fibre networks into the data communication system enclosure itself. The introduction of OPCB technology is a critical part of this migration as they provide a medium, which can accommodate hundreds of times the volume of data compared to conventional PCBs and are therefore seen as a key enabling technology for future, high bandwidth data communication systems.

One area in which this migration is particularly apparent is in modern data centres. A major consequence of the widespread adoption of smaller mobile data devices (smart phones and tablets) over fixed larger computer terminals (PCs and laptops) is that a dramatic shift is now occurring in where customers need to store their information. While it has, until recently been sufficient to store data locally (such as on the user's local laptop or desktop computer hard drive), the average size of data objects generated on the fly now, such as high definition pictures or short videos, has grown to the extent that the storage available on mobile devices is rapidly becoming insufficient for long term accumulation and retention of data. So called "Cloud" services are therefore emerging to meet a burgeoning customer demand to store data remotely and securely. Data centres provide the dedicated compute, storage and server equipment required to meet the remote data processing and storage requirements of these emerging Cloud environments, but in order to cope with rapidly changing customer demand, the

architectures underlying the data centres themselves need to evolve and a critical part of that evolution is the deployment of optical connections at all levels of the data centre environment.

## *1.2 Motivation and aim*

Widespread adoption of OPCBs will herald substantial performance, cost and environmental benefits for the data communications industry, however there are still a number of technical barriers that need to be overcome before OPCBs can become commercially viable, the most challenging of which is the provision of reliable connector interface technologies allowing pluggable (interfaces can be repeatably attached and detached) optical connectivity to OPCB embedded optical waveguides.

The purpose of this thesis is to investigate and advance the deployment of OPCB technology into data centre systems, by designing and developing new connector and interface solutions to enable pluggable optical connectivity directly to planar waveguides (waveguide structures of uniform thickness subtending a plane) embedded in OPCBs. The thesis focusses on research advances in polymer waveguide based OPCB connector interfaces [1], [2] and low cost high precision optical assembly methods [3], however also explores the application of the same connector interface principles to planar glass waveguide based OPCBs [4]. The author has invented, designed and developed a variety of optical waveguide connector interfaces, optical assembly and coupling mechanisms. In order to evaluate these technologies, the author has also designed and developed a suite of complete optical backplane connection platforms (Figure 1-1).

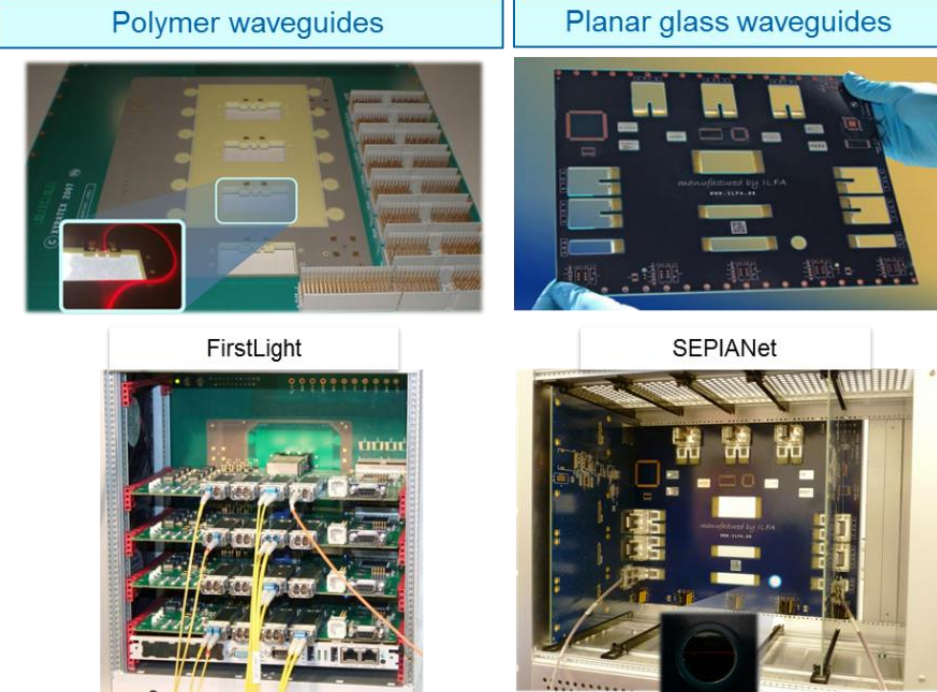


**Figure 1-1: Optical circuit board varieties and associated demonstration platforms**

### 1.2.1 Chapter 1: Introduction

Chapter 1 provides an introduction to the motivation and aim of the thesis and a description of the thesis structure.

This chapter will introduce the trends in modern data centres, which are fuelling the migration of optical interconnect including the adoption of disaggregated architectures and the increase in network protocol speeds across all tiers of the data centre. It will also describe international collaborative research projects, which are advancing system embedded optical interconnect technologies and with which the author has been involved to varying degrees over the course of the thesis. The broad aim of these projects is to develop a complete suite of dense, high-performance and low-energy optical interconnect solutions spanning the system connectivity hierarchies of data centres and high performance computers including rack-to-rack, board-to-board, chip-to-chip and on-chip optical links [5].

### 1.2.2 Chapter 2: System embedded optical interconnect technology survey

Chapter 2 provides an overview of the state of the art in commercial and prototype technologies, which enable optical connectivity within the system enclosure itself. This suite of technologies includes midboard optical

transceivers, passive optical connectors for front fascia and backplane connections and the different media for optical channels within the system.

Chapter 2 will also consider in more detail the design constraints on high frequency copper channels on the PCB and the benefits of replacing these with optical channels as well as the the relative merits of different OPCB interconnect technologies.

#### *1.2.2.1. Methodology*

Throughout the project, comprehensive commercial and academic literature searches will be carried out to determine the state of the art in waveguide, OPCB, connector and transceiver interconnect technology. As this is a new field of research, publications will be continuously surveyed and where possible appropriate conferences attended to establish research and supplier networks.

### 1.2.3 Chapter 3: First generation pluggable active optical circuit board connector for polymer waveguide based optical circuit boards

In order for OPCB technology to become commercially viable, it is crucial that the embedded waveguides can be terminated and that connector technologies are developed that allow either other PCBs or external optical cables to connect directly to the optical waveguides in the OPCB.

This will require both novel optical connector technologies and the ability to reliably and accurately align and assemble components onto OPCBs with respect to the optical waveguides therein. In the past such assembly could only be achieved through active alignment, however this would be unsuitable for high volume OPCB assembly due to the impact on equipment cost, assembly, assembly time and board yield.

#### *1.2.3.1. Objective*

- Develop a novel method of connecting peripheral devices orthogonally to a polymer waveguide OPCB (i.e. with the principal axis of the peripheral device connection at 90° to the plane of the OPCB), which precludes the need for embedded mirrors
- Develop active transceiver based connector system

- Develop novel, low-cost method of assembling components (optical or mechanical) to an OPCB to allow suitably accurate alignment to multimode optical waveguides embedded in the PCB substrate
- Develop bespoke mechanical coupling elements to form part of the waveguide connector receptacle
- Develop test platforms to characterise these novel technologies

#### *1.2.3.2. Methodology*

An in-plane optical connector interface concept will be invented and a number of iterative prototypes designed and developed to prove the concept and evaluate its suitability for commercial deployment. The design will incorporate a high density parallel optical interface and would need to accommodate high speed serial modulation rates of over 10 Gb/s per channel. To this end a parallel optical transceiver circuit incorporating such an interface will also be developed to form part of an active connector mechanism. This mechanism will reside on the edge of pluggable peripheral devices to allow those devices to be optically plugged and unplugged to and from an optical printed circuit board. A suitable programming interface will be developed to allow user configuration of critical transceiver control parameters such as channel enable, laser modulation current, laser bias current and receiver squelch.

Peripheral test cards will also be designed to accommodate the connector prototypes. These test cards will serve as a conduit for external serial test data to the transceiver channels in the connector. Finally a complete proof of concept demonstration platform will be constructed, comprising a test chassis, single board computer, test cards, prototype connectors and an OPCB to allow comprehensive optical and mechanical characterisation of the connectors.

A crucial part of this OPCB connection system will be the optical waveguide receptacle for the pluggable active in-plane connector, which is fixed to the board and meets the tolerance requirements to enable the optical connector interface to align with the embedded waveguide interface.

To this end, a novel low-cost method will be invented and developed of assembling components (optical or mechanical) to an OPCB such as to allow suitably accurate alignment to multimode optical waveguides embedded in the PCB substrate. A suite of receptacles will be designed and developed, which incorporates features to allow passive accurate alignment of the receptacle onto the board waveguides. In addition, compliant

features will be designed in the board itself to accommodate the accurate mechanical registration of the receptacle.

The final iteration of the prototype will allow a lens array to be accurately fixed to the OPCB and form part of a dual lens expanded beam coupling solution.

### 1.2.4 Chapter 4: Second generation pluggable active optical circuit board connector for polymer waveguide based electro-optical circuit boards

This chapter is an extension of chapter 3 and describes the second improved iteration of pluggable optical connector for OPCBs and corresponding demonstration platform.

### 1.2.5 Chapter 5: Pluggable passive optical circuit board connector for planar glass waveguide based optical circuit boards

While polymer waveguides are well suited to applications requiring 850 nm optical signals, such as low cost optical links based on commodity Vertical Cavity Surface Emitting Lasers (VCSELs), planar glass waveguides would be preferable in applications requiring longer wavelengths such as 1310 nm or 1550 nm due to their superior transmissivity at these wavelengths. The emergence of affordable longer wavelength transceiver solutions based on photonic integrated circuits such as silicon photonics makes glass waveguides an attractive OPCB technology.

In this chapter, the author describes how some of the connector principles developed for polymer waveguide interfaces can be successfully deployed to allow connector termination of planar glass waveguides.

#### *1.2.5.1. Objective*

- Co-develop a waveguide receptacle based on those already developed for polymer waveguides to be assembled onto planar glass waveguide based OPCBs allowing direct fibre-to-waveguide coupling
- Lead the design and development of passive board-to-board optical connector system, which incorporates the waveguide receptacle, but also makes use of the commercial parallel optical ferrule jumpers

- Design and develop an optical connector and backplane test platform, which incorporates a glass waveguide based OPCB backplane and peripheral test cards, which can be plugged directly into the OPCB backplane using the optical connectors developed.

#### *1.2.5.2. Methodology*

The author will work with our partner organisations in the consortium of the EU Piano+ SEPIANet project to provide the key elements of the demonstration platform:

The planar multimode glass waveguides will be fabricated within thin glass foils based on a two-step thermal ion exchange process by Fraunhofer IZM. Novel lamination techniques will be developed by ILFA GmbH to allow glass waveguide panels to be reliably integrated into a conventional electronic multi-layer printed circuit board.

The author will lead and co-design a complete suite of optical connector technologies to enable both direct fibre-to-board and board-to-board connectivity.

Both on-card and externally generated 850 nm and 1310 nm optical test data will be conveyed through the connector and waveguide system and characterised for in-system and system-to-system optical connectivity at data rates up to 32 Gb/s per channel.

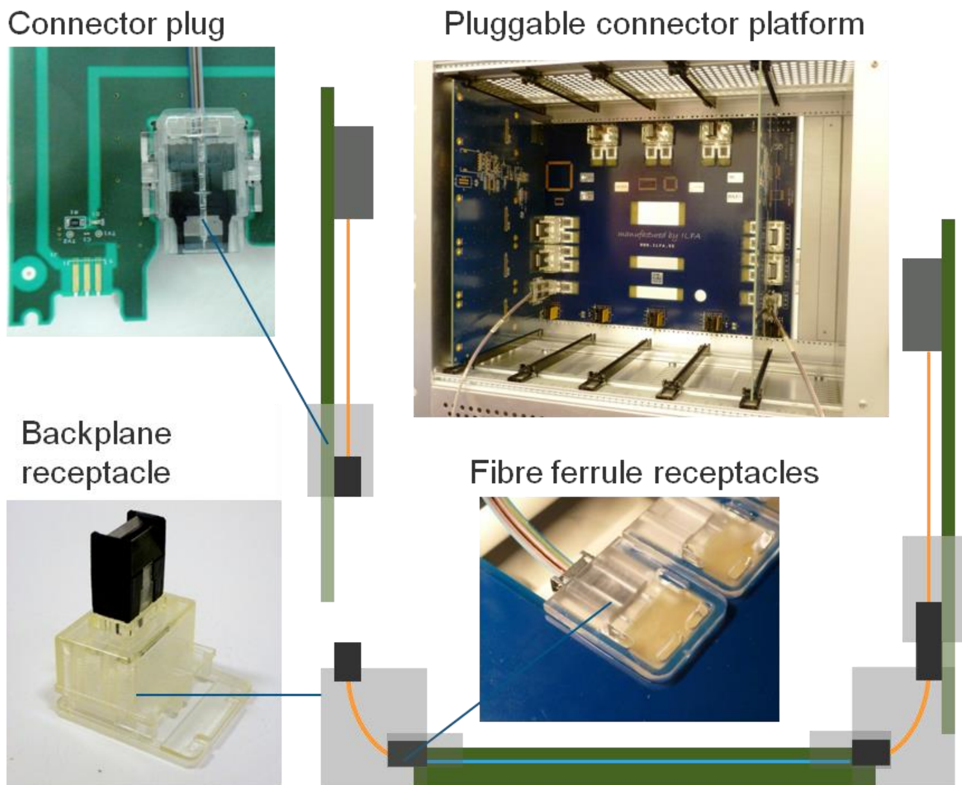


**Figure 1-2: SEPIANet glass waveguide optical circuit board connector platform**

### 1.2.6 Chapter 6: Intellectual Property

Chapter 6 will provide an overview of the intellectual property portfolio developed during the project.

#### *1.2.6.1. Objective*

- Develop Intellectual Property to address problems encountered and capture innovative solutions, preferably in the form of a patent portfolio broadly covering the field of OPCB interconnect including passive and active optical connectors, precision alignment and assembly techniques and OPCB waveguide structures.

#### *1.2.6.2. Methodology*

The author will identify problem areas and develop innovative solutions, which he will capture as invention disclosures using the Xyratex Intellectual Property (IP) capture and submission system. The author will submit invention disclosures for review by the Xyratex IP review team, which will determine whether inventions warrant being pursued as patent filings. He will assist patent lawyers in converting the invention disclosure into proper patent applications to be filed.

As of November 2015, the intellectual property portfolio developed by the author consisted of 17 patent families divided into three broad categories:

1. Connectors, which includes a comprehensive variety of optical connector solutions
2. Waveguides and fabrication, which includes advanced optical waveguide structures, fabrication techniques
3. System and assembly, which includes solutions for rack scale optical interconnect and optical assembly

### 1.2.7 Chapter 7: Conclusion

Chapter 7 will summarise the achievements of the project and describe the future work that will be carried out in this field.

## *1.3 Optical interconnect migration in data centres*

Our growth as an information affluent society will be gated by our ability to consolidate and control storage, processing and switching of digital information on a massive scale. By 2020, it is predicted that ~44 ZB of data will be created, of which ~13ZB will need to be stored, however the amount of data that installed capacity will be able to hold will only be ~6.5ZB [6]. A consequence of the continuing proliferation of mobile devices coupled with the exponential demand in data capacity and performance is now pushing us towards an architectural bottleneck, where traditional data centre infrastructures will not support the required level of information migration from local client side operations to exascale cloud environments.

There is a critical need for data centre architectures to evolve to provide the flexible control to eliminate underutilisation of compute and storage resources, minimise latency for "east-west" communication and support linear scalability of equipment. The prevailing wisdom advances network management solutions coupled with disaggregated modular architectures, in which data centre nodes are optically interconnected across different hierarchical tiers.

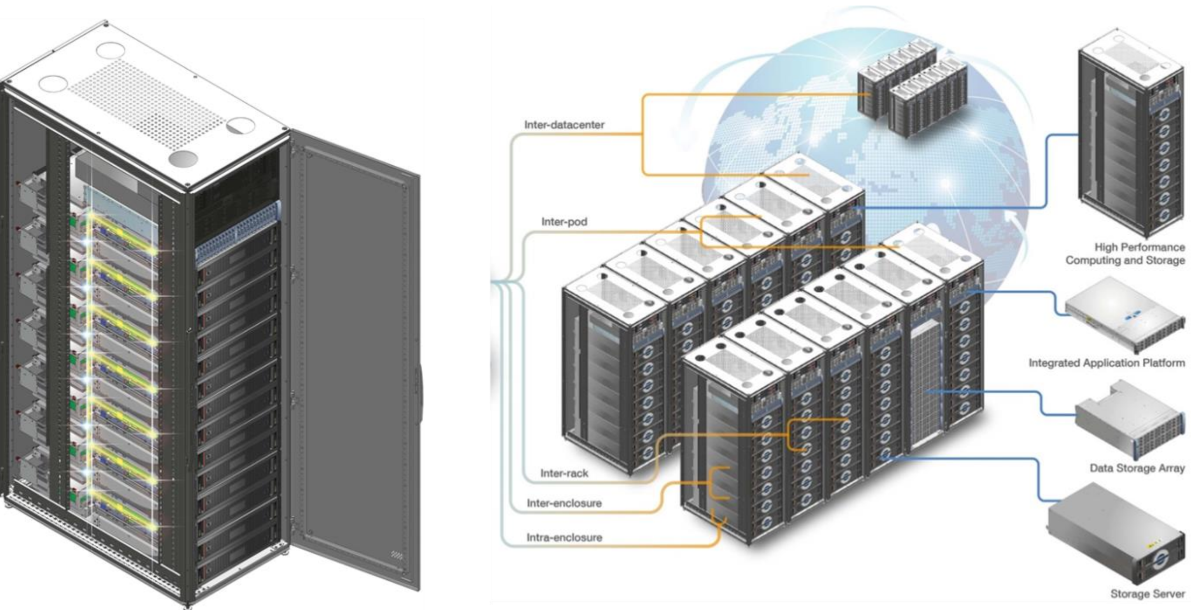


**Figure 1-3: Optical interconnect migration in data centre environments**

### 1.3.1 Data centre overview

Data centres provide dedicated storage and server equipment designed to meet the data processing and storage requirements of an organisation, where such requirements can vary strongly from organisation to organisation. Modern large data centres can contain hundreds of thousands of servers and data storage systems, which are typically placed in racks, with about 20-40 servers in a given rack. Often server and storage racks are grouped into pods and connected either physically in the same enclosed area or logically networked to a common pod switch. Figure 1-4 shows the floor plan of a small data centre facility.

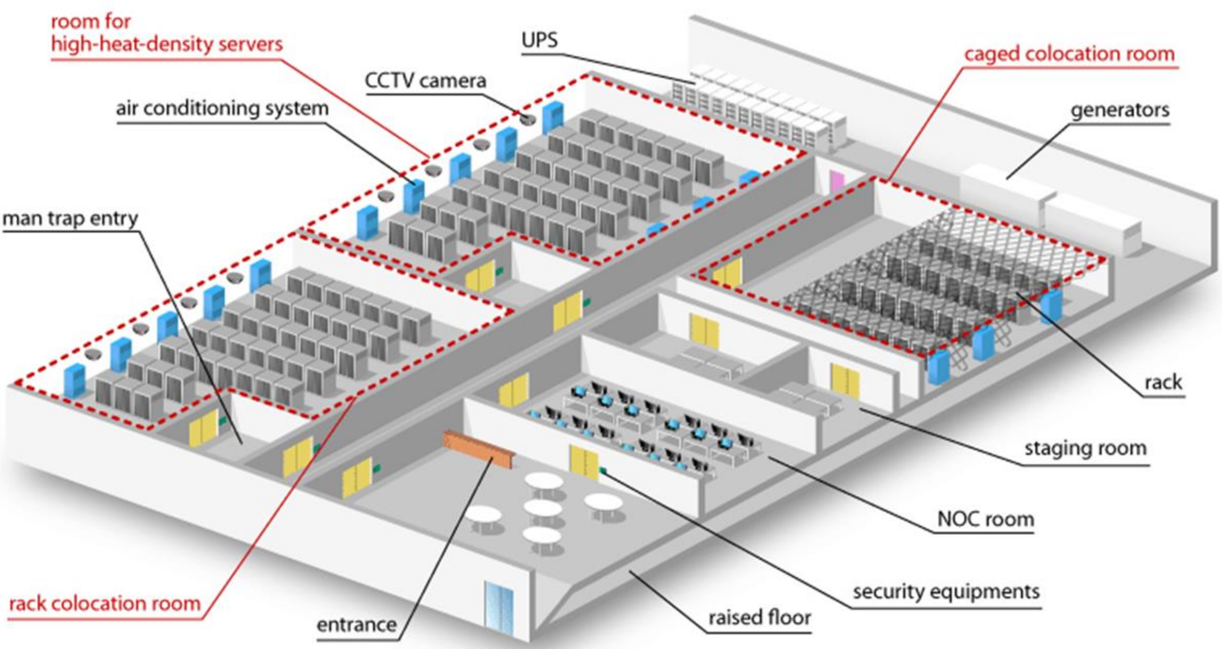


**Figure 1-4: Floor plan of a single storey data centre (Source: NTT Japan)**

#### *1.3.1.1. Data storage systems*

Data storage system technologies, as shown in Figure 1-5, form the crucial building blocks of modern data centres, wherein data storage arrays (typically incorporating data storage devices based on magnetic or solid state media) are connected within systems of varied complexity and size ranging from simple high capacity storage racks to high performance computing data storage systems.

A generic data storage array system (Figure 1-5a) may comprise an array of hard disk drives or solid state drives connected to one side of a passive midplane while controller modules and power supplies are connected to the

other side. The controller modules are designed to fit into a canister of a specific standard form factor defined in the Storage Bridge Bay (SBB) Specification [7].

As shown, the midplane and its peripherals are connected in a mutually orthogonal geometry, which, as will be described later, is advantageous for in-plane optical connections.

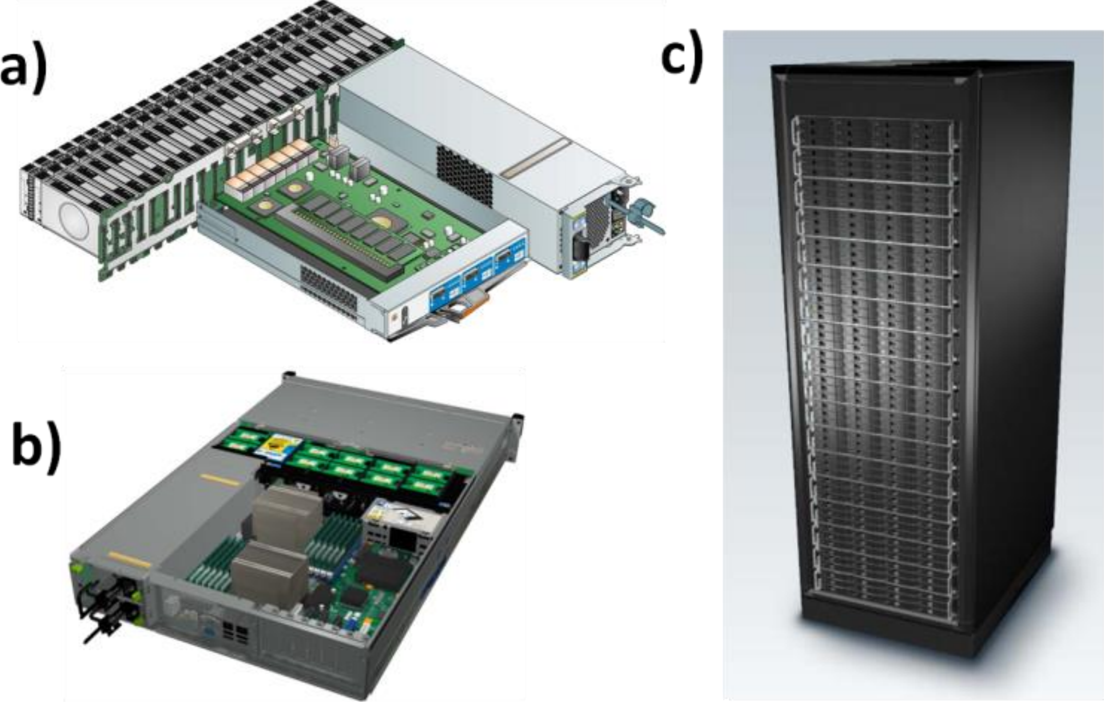


**Figure 1-5: a) Generic data storage array system, b) Data storage integrated application platform, c) 42U data storage rack**

Figure 1-5b shows a data storage integrated platform, which incorporates greater functionality such as server hardware into a data storage system combining computational and storage capabilities. Figure 1-5c shows a data storage rack in which assorted data storage subsystems are incorporated according the required capacity and processing power of the rack and data centre application.

### 1.3.2 Data storage system interconnect topologies

The interconnect topology in high availability enterprise class data storage systems is typically defined by a dual star configuration (Figure 1-6), whereby each data storage device supports two duplex data links on the midplane, one to each of at least two separate controller modules. As a current example the Storage Bridge Bay

(SBB) Specification [7], which defines mechanical, electrical and low-level enclosure management requirements for an enclosure controller slot, allows for a maximum of 48 disk drives in a given data storage system enclosure, each supporting two duplex links to each of two controller modules, with a further 17 duplex links directly connecting the controllers to each other. The midplane of a 48 drive storage array would therefore have to accommodate 113 duplex links or 226 high speed transmission lines. Furthermore there is an enhancement of the drive interface standard in development [8], which allows up to four independent duplex links per drive scaling the number of high speed links on the data storage midplane accordingly.

The level of fault tolerance and scalability offered by these topologies forces increased complexity and cost into the midplane, particularly when interconnect protocols define serial data rates beyond 24 Gb/s.

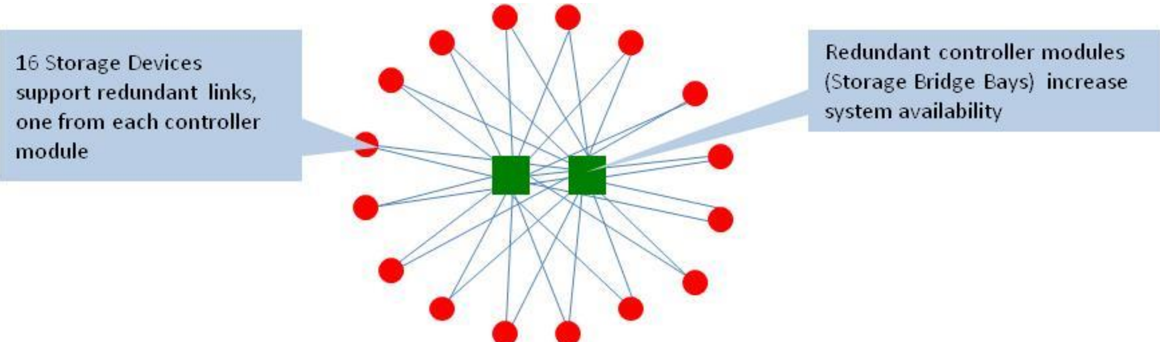


**Figure 1-6: Example of a Dual Star interconnect topology whereby 2 controller nodes each have a link to 16 storage devices, thus eliminating the controller module as a single point of failure**

### 1.3.3 Data centre switching architectures

The term "fat tree" is used to describe the network topology connecting large numbers of different end-hosts and is likened to an upside down tree with the fat trunk and root at the top, representing the highest switching tier. As one moves down from the root to the lower switching tiers with more numbers switching nodes across "branches", which become progressively thinner until one reaches the eventual "twigs" holding the "leaves" of the fat tree. The twigs, branches and trunk of the fat tree represent the increasing bandwidth of data links required as one moves from one tier up to the next.

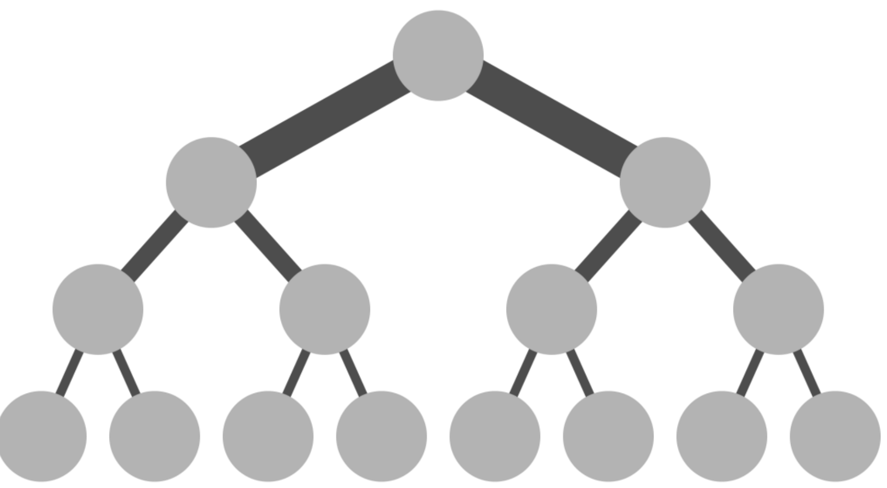

**Figure 1-7: Schematic view of traditional fat tree with links becoming thicker as we move up the hierarchy**

Fat-tree networks are built out of Ethernet or Infiniband packet switches in several tiers or levels. Typically, end-hosts (servers) are installed in racks, and the servers in a rack are connected to a top-of-rack (TOR) switch, that provides southbound interfaces to the rack's servers and northbound interfaces to the higher levels of the data centre switching network. The TOR switches of all racks form the "leaves" of the tree network.

There are three basic fat tree topologies:

**Fat-link** (traditional fat tree) method where the rate of links (size of branch) increases as we move closer to the root (the rate of the port of a switch facing upwards is higher than the rate of ports facing downwards). This method would be impractical for bigger data centres, as increasing the number of end-hosts would require higher bandwidth ports at the switches in the higher tiers of the fat tree. This would mean switches would become prohibitively expensive as the number of end-hosts was increased and the data centre could not be scaled beyond the limit of achievable link bandwidths.

**Folded Clos uniform rate** method where all links have the same rate and the required bandwidth is provided by having multiple paths between end-hosts [9]. This method allows use of low cost, commodity switching equipment at the higher interconnection levels as well as multiple paths between the end-hosts (servers) [10]. Figure 1-8 shows a folded Clos fat tree architecture based on the deployment of commodity TOR switches at all interconnection levels solves the link, port and switch failure problems.

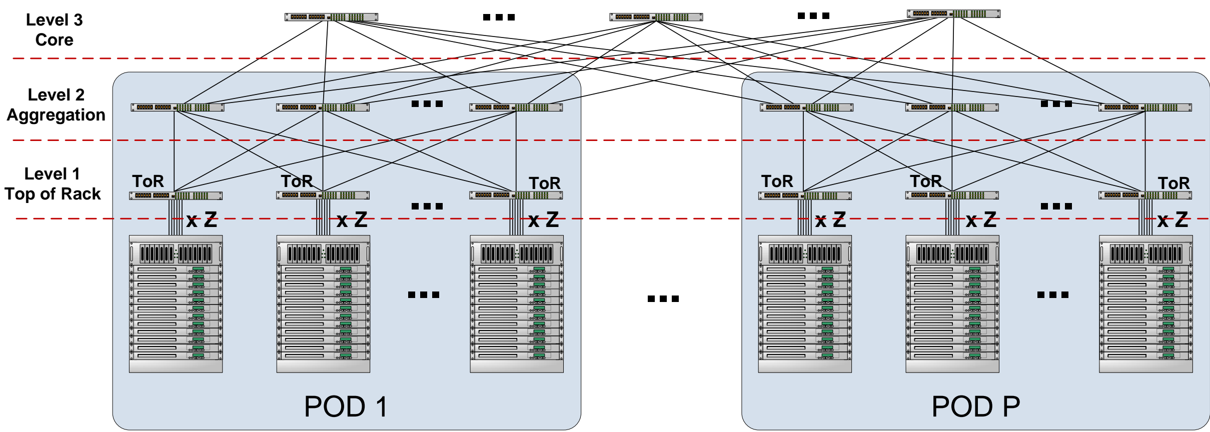


**Figure 1-8: Fat tree network where all three levels are built out of uniform rate port commodity TOR switches in a (folded) Clos topology (Source: Nephele project [11])**

**Folded Clos non-uniform rate method** combines attributes of the aforementioned Clos uniform rate and Fat link topologies. This method would leverage a multiple path Clos topology while also providing increasing link bandwidths at higher levels in the tree, but within the constraints of commercially available transceiver interconnect. For example, 10 Gbps bandwidths can be deployed at server and TOR switches ports, while emerging 40 Gbps and 100 Gbps optical link technologies can be deployed at higher levels.

**Modern data centres are typically designed with a fat-tree or oversubscribed fat-tree data network topology, employing folded Clos uniform or non-uniform rate topologies.**

A (full) fat-tree network provides **full bisection bandwidth**, that is at any given instance each end-host can talk at full rate to another end-host and the network can support simultaneous communication of all end-host pairs. Most data centres have tens to hundreds of thousands of servers and require 3 to 4 fat-tree levels to achieve full bisection bandwidth. However the use of folded Clos topologies mean that the data centre scales super-linearly, that is, as the number of end-hosts increases by a given factor, the number of required switching nodes to support full bisection increases by a larger factor. Moreover fat tree networks are under-utilised most of the time with even oversubscribed trees reporting less than 20% average utilization [12], hence full bisection bandwidth is not needed.

Sometimes however data traffic congestion, in the form of hotspots can be created whereby data packets are lost, due to the rigid allocation of the available (reduced) bandwidth. Even when lossless operation is guaranteed, congestion adds queuing delays leading to latencies. This can be mitigated to some degree through the use of

protocols and algorithms that allow traffic loads to be properly balanced across the data centre network e.g. the equal-cost multi-path routing (ECMP) protocol [13].

### 1.3.4 Resource disaggregation

According to the conventional data centre model, each server system consists of a fixed ratio of computing, memory, storage and communication resources that are all "aggregated" into a single enclosure. ICT requirements on a data centre can differ vastly depending on the user and application. Different ICT requirements can be satisfied by different ratios of the compute, storage and memory subsystems, which form the building blocks of modern data centres (**Figure 1-9**a), but as requirements change, so too will the necessary ratio of compute, storage and memory utilisation. These building blocks include, but are not limited to servers, data storage arrays, switches and high performance storage and computer subsystems.

Data centre customers will typically lease a fixed cluster of equipment and to ensure that they can always accommodate demand including transient peak demand, they will tend to strongly over-provision resources, that is, on average, their equipment utilisation will be quite low, as evidenced by the less than 20% average utilisation reported by Benson [12].

Resource disaggregation is an emerging paradigm, which allows resources to be shared across the data centre in an on-demand fashion by taking advantage of the modular nature of data centre systems and subsystems. This physical decoupling of resources allows for more fine-grained resource provisioning, and the ability to multiplex available resources according to need will give rise to higher utilisation [14].

If these resource subsystems could be arranged to be truly modular and work in independence of their location within the data centre, and if interconnect length and bandwidth constraints between these subsystems could be neglected, then this would allow a disaggregated architecture as shown in **Figure 1-9**b). In this case the combination of subsystems required to satisfy a given set of ICT requirements needn't be constrained to the same rack or cluster of racks, they could be physically dispersed across the data centre.

The Disaggregated Rack-Scale Server (DRS) architecture proposed by Facebook and Intel as part of the Open Compute project [15] proposes the separation of computing, storage and communication hardware components within the rack and the interconnection between them with distributed switching functions [7]. The Open Compute Project anticipates a 24% reduction of costs and an efficiency increase of about 38% with this new disaggregated rack paradigm [16].

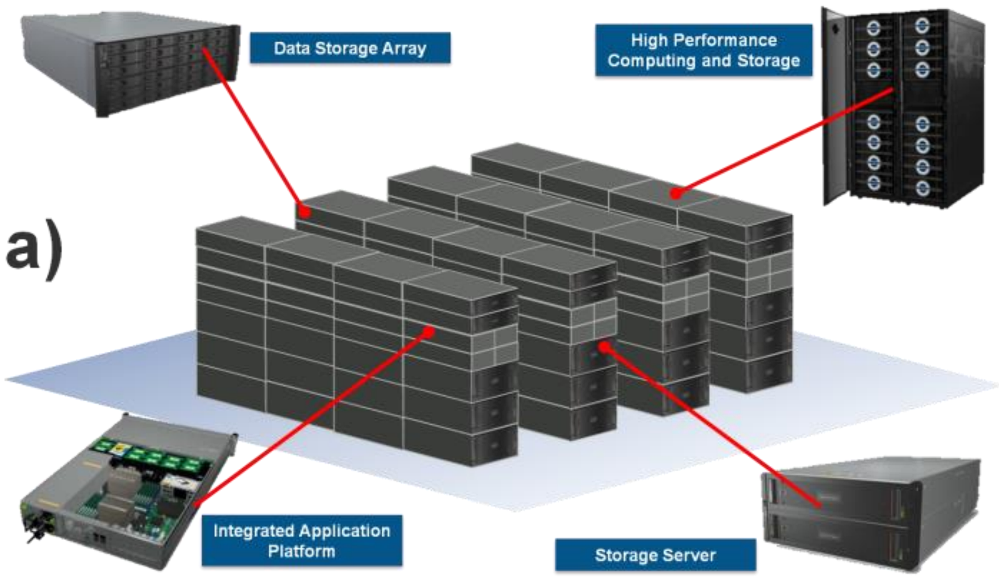


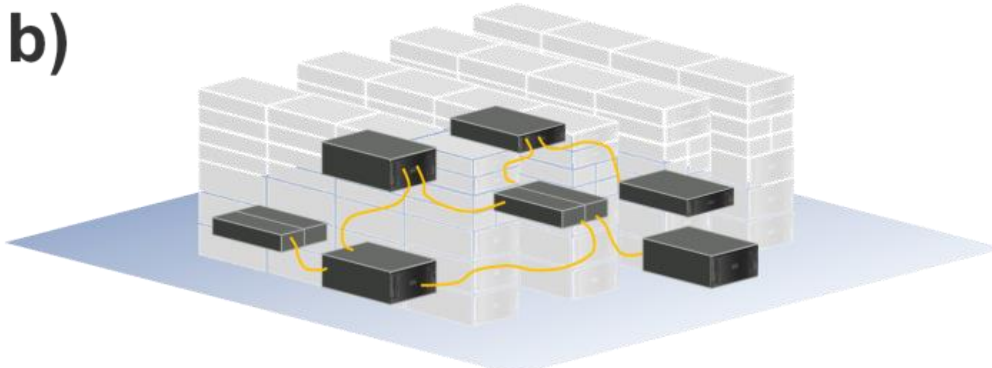


**Figure 1-9: a) Building blocks of modern data centres, b) Schematic of optical links connecting geographically dispersed subsystems within a data centre architecture**

In a similar vein, the drive to increasing virtualisation of the data centre through Software Defined Network (SDN) architectures also promises to provide significantly greater user control, Quality of Service and flexibility while optimising resource use. Ideally the user can be provided a virtual data centre solution with the optimum ratio and amount of compute, memory and storage, which varies dynamically depending on the user's transient needs, even though the actual corresponding hardware allocated could be dispersed. In order to satisfy these requirements without over-provisioning of hardware resources, one must have the capability to convey high bandwidth data over far longer distances then is typical or possible today between subsystems, and this can only

be satisfied by low-cost, high-bandwidth optical links. Indeed fibre-based commercial optical modules are now common in data centres for rack-to-rack connectivity [17].

### 1.3.5 Optical interconnect migration

However, while, as system bandwidths increase, the provision of ubiquitous optical links would remove the interconnect bottleneck between racks or subsystems within a rack, new bottlenecks will emerge or existing ones will become more exposed deeper in the system enclosure itself. Thus the need for commercially viable, dense interconnect solutions will continue inevitably to migrate down through the data communication tiers of the system from board-to-board, chip-to-chip and ultimately to the chip itself.

The migration of optical connectivity within data centres is already underway, with hybrid electro-optical infrastructures proposed and numerous proof-of-concept technologies developed [18]–[20]. Notable examples include the reported deployment by IBM of optical interconnect for POWER7-IH [18] systems with 100,000s of high-performance CPU cores by leveraging dense optical transceiver and connector technologies to construct chip module optical IOs. Fujitsu Laboratories proposed dense optical interconnect architectures for next-generation blade servers [21], with a demonstration of an electro-optical midplane with 1920 embedded optical fibres to meet the projected bandwidth requirements [22].

HP developed an optical backplane with broadcast and micro electro-mechanical systems (MEMS) based optical tapping capabilities along an embedded plastic waveguide, suitable for non point-to-point interconnect topologies [23], which was demonstrated within a proof-of-concept network switch chassis.

Commercial adoption of system embedded photonic solutions will be gated by the priority requirements relevant to the application space or market in question. Applications that prioritise performance and bandwidth density, would be amongst the first adopters. For example, in the IBM Blue Waters Supercomputer, optical links are deployed at both the inter-rack and intra-rack levels [18]. In other application spaces, such as high volume ICT equipment, internal optical interconnect technologies will most likely only be adopted once it becomes cost competitive with traditional copper interconnect solutions or once traditional interconnect can no longer meet the evolving system bandwidth requirements.

Xyratex, Finisar, Vario-optics and Huber+Suhner demonstrated an optically enabled data storage platform, in which 12 Gb/s SAS traffic was conveyed optically between two internal controller cards along 24 PCB embedded polymer optical waveguide channels, thereby showing, for the first time, how in-system optical channels could be successfully deployed within a 12G SAS architecture [24].

## *1.4 Data centre network communication protocols*

Data centres can play host to a variety of data communication protocols including Ethernet, Serial Attached SCSI, Infiniband and PCI Express. The increase in data bandwidth requirements in modern data centres is strongly reflected in the roadmaps of these protocols, which are briefly introduced below.

### 1.4.1 Serial Attached SCSI

Serial Attached SCSI (SAS) is a point-to-point serial bus protocol providing connectivity to storage devices including Hard Disk Drives (HDD), Solid State Drives (SSD) and Solid State Hybrid Disk Drives (SSHD) and, as such, defines the speed with which data is conveyed between the peripheral devices in a storage system environment.

The T10 technical committee of the International Committee for Information Technology Standards (INCITS) develops and maintains the SAS protocol. The SAS roadmap [25] maintained by the T10 committee [26] currently defines a serial data rate of 12 Gb/s and is set to double to 24 Gb/s by 2017 (Figure 1-10).

**Table 1-1: Serial Attached SCSI protocol generations**

| Name | Raw bit rate / Gigabaud | Encoding | Availability |
|---|---|---|---|
| 3G SAS | 3 | 8b/10b | 2002 |
| 6G SAS | 6 | 8b/10b | 2006 |
| 12G SAS | 12 | 8b/10b | 2010 |
| 24G SAS | 24 | 8b/10b | 2016 |

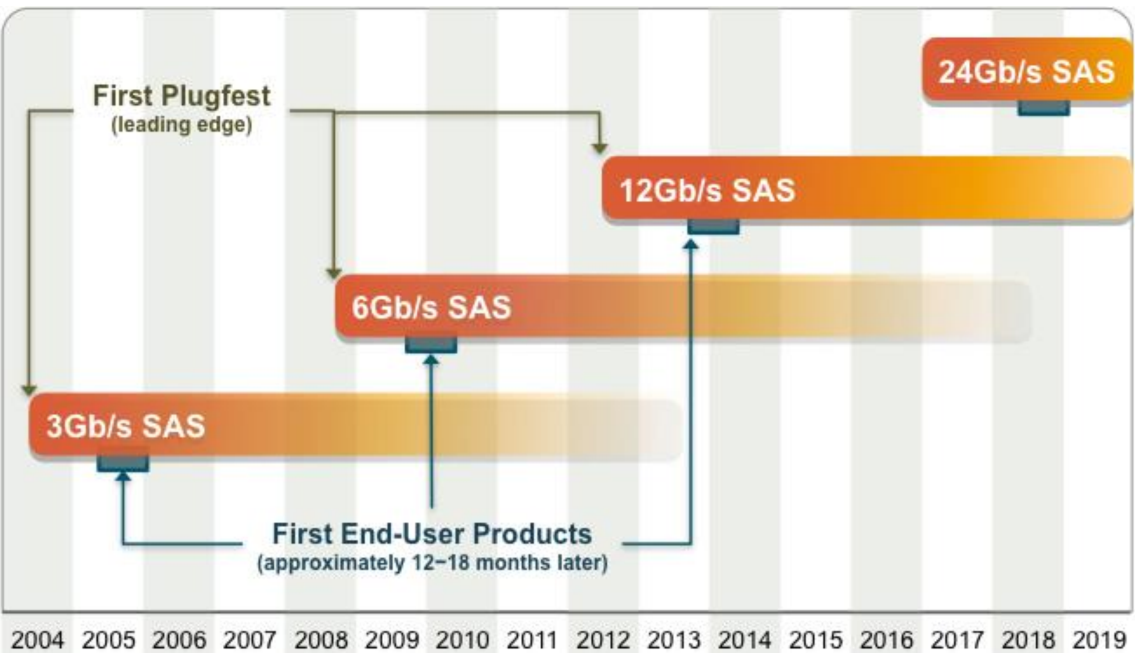


**Figure 1-10: Serial Attached SCSI Roadmap (Source SCSI Trade Association – March 2014)**

### 1.4.2 Infiniband

The InfiniBand protocol is used predominantly for rack-to-rack communication in enterprise data centres and high performance computers (HPC) currently defines a serial bit rate of 14 Gb/s under the FDR (Fourteen Data Rate) scheme, which is set to increase to 26 Gb/s with the introduction of EDR (Enhanced Data Rate) and 50Gb/s with the introduction of HDR (High Data Rate) [27].

The Infiniband roadmap is shown in Figure 1-11.

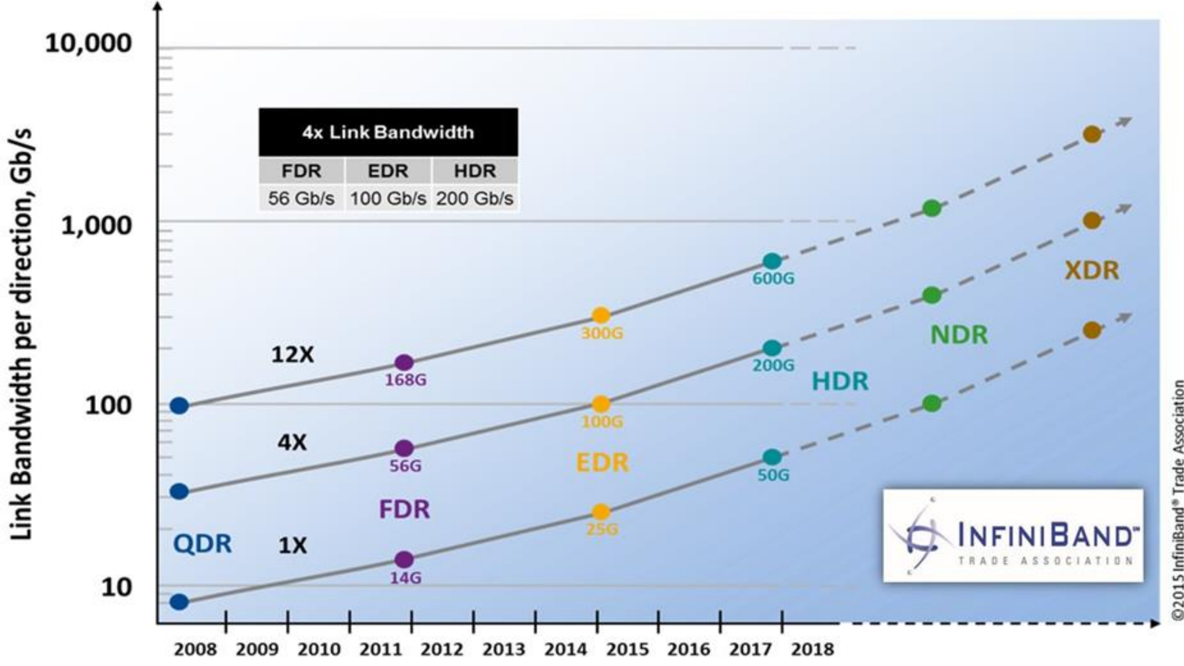


**Figure 1-11: Infiniband™ Roadmap (Source: InfiniBand™ Trade Association)**

### 1.4.3 PCI Express

PCI Express is the dominant CPU expansion bus protocol in servers, computer motherboards and integrated application platforms. PCI Express interfaces can accommodate up to 16 lanes, each capable of supporting high speed traffic. PCI Express 3.0, which is currently implemented, is specified to run parallel channels (1, 4, 8 or 16 wide) at a serial data rate of 8 Gb/s per lane, with PCI Express 4.0 due to increase the per lane bandwidth to 16 Gb/s by 2016 [28].

The PCI Express protocol is also gaining prominence as an alternative drive interface and is one of the preferred data interfaces to Solid State Drives (SSDs) and other flash memory devices, such as memory cards.

**Table 1-2: PCI Express protocol generations**

| Name | Raw bit rate / Gigabaud | Link Bandwidth | Total BW x 16 / | Availability |
|---|---|---|---|---|
| PCIe 1.x | 2.5 | 2 Gb/s | ~ 8 GB/s | 2002 |
| PCIe 2.x | 5 | 4 Gb/s | ~16 GB/s | 2006 |
| PCIe 3.x | 8 | 8 Gb/s | ~32 GB/s | 2010 |
| PCIe 4.x | 16 | 16 Gb/s | ~64 GB/s | 2016 |

### 1.4.4 Fibre Channel

Fibre channel is a communications protocol almost exclusively used between servers and storage systems and is typically restricted to shorter distances of between 2 metres to 20 metres.

**Table 1-3: Fibre Channel protocol generations**

| Name | Raw bit rate / Gigabaud | Link Bandwidth / MB/s | Encoding | Availability |
|---|---|---|---|---|
| 1GFC | 1.0625 | 100 | 8b/10b | 1997 |
| 2GFC | 2.125 | 200 | 8b/10b | 2001 |
| 4GFC | 4.25 | 400 | 8b/10b | 2004 |
| 8GFC | 8.5 | 800 | 8b/10b | 2005 |
| 10GFC | 10.52 | 1200 | 64b/66b | 2008 |
| 16GFC | 14.025 | 1600 | 64b/66b | 2011 |
| 32GFC | 28.05 | 3200 | 64b/66b | 2016 (projected) |
| 128GFC | 4x28.05 | 12800 | 64b/66b | 2016 (projected) |

### 1.4.5 Ethernet

Ethernet is the dominant protocol in data centres for application traffic and has emerged as the backbone of modern data storage infrastructures. This is evidenced by the growth of virtual servers and desktop infrastructure, as well as the rampant adoption of Amazon Web Services.

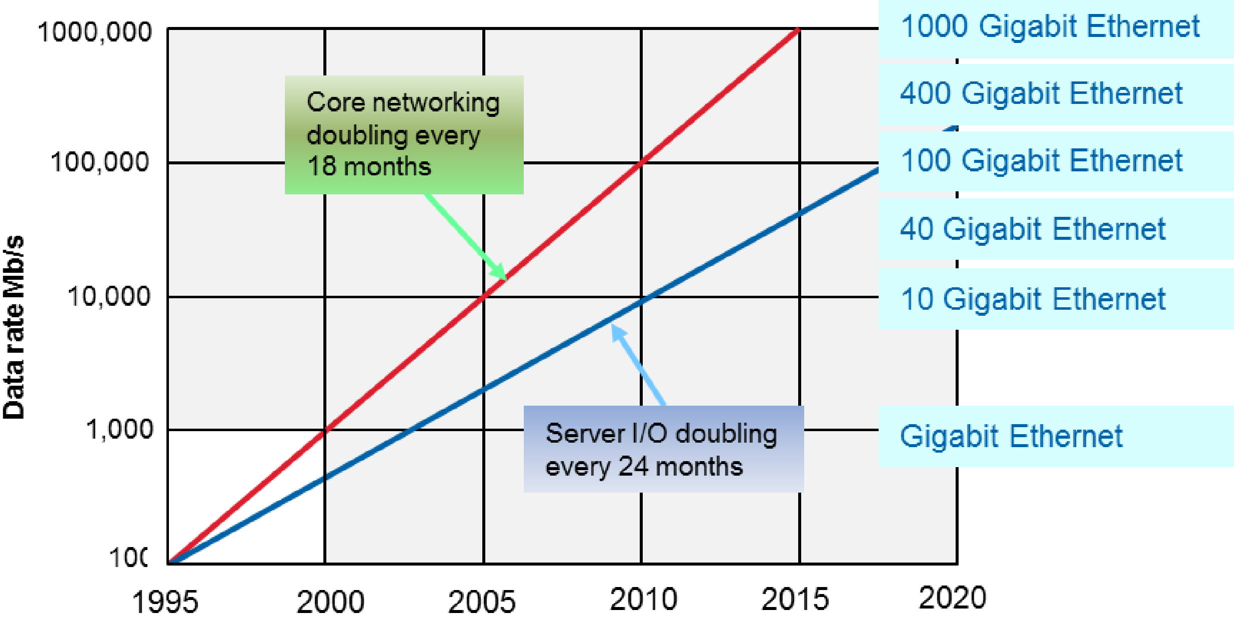


**Figure 1-12: Ethernet roadmap [29]**

The Ethernet standards are governed by the IEEE trade association with most recent emerging standards (e.g. 400 GbE discussed below) specifying optical interfaces and optical links. It was identified that the bandwidth associated with core networking was observed to double every 18 months in accordance with Moore's Law, while the bandwidth associated with high volume x86 servers and computing applications, was doubling every 24 months. In order to accommodate these two separate bandwidth trends, the IEEE P802.3ba Task Force defined objectives for two new wireline Ethernet speeds – 40 Gigabit Ethernet and 100 Gigabit Ethernet. The IEEE working group 802.3ba defines 40 GbE and 100 GbE, electrically while working group 802.3bm defines 40 GbE and 100 GbE optically, based on 4 channels of 10.3125 Gbps NRZ (Non-Return to Zero) and 10 channels of 10.3125 Gbps NRZ respectively.

The Ethernet bandwidth growth rate shown in Figure 1-12 already predicted the need for 400 Gb/s by 2013 and 1 Tb/s by 2015 on the faster core networking trend path. In order to develop standards targeting an Ethernet rate beyond 100 Gb/s, the IEEE 802.3 Industry Connections Ethernet Bandwidth Assessment Ad Hoc was created in 2011. The working group 802.3bs is responsible for 400 GbE standardisation with many options being developed to achieve these aggregate bandwidths.

Electrical implementations of 400 GbE under consideration include:

- 16 channels of 25 Gb/s NRZ, in which 16 separate electrical signals are modulated between two signal amplitude levels at a rate of 25 Gbd

- 8 channels of 25 Gbd PAM4 (4 level Pulse Amplitude Modulation), in which 8 separate electrical signals are modulated between four signal amplitude levels at a rate of 25 Gbd
- 8 channels of 50 Gbps NRZ, in which 8 separate electrical signals are modulated between two signal amplitude levels at a rate of 50 Gbd

Optical implementations of 400 GbE under consideration include:

- 8 channels of 25 Gbd PAM4, in which 8 separate optical signals are modulated between 4 signal amplitude levels at a rate of 25 Gbd
- 8 channels of 50 Gbd NRZ, in which 8 separate optical signals are modulated between 2 signal amplitudes at a rate of 50 Gbd
- 4 channels of 100 Gb/s NRZ, in which 4 separate optical signals are modulated between 2 signal amplitude levels at a rate of 25 Gbd

The next future standard under consideration is 1 Tb/s (1000 Gigabit Ethernet), which is as yet undefined, though it will most likely be based on N x 25 Gb/s or 50 Gb/s (PAM4 or NRZ) electrically and optically.

#### *1.4.5.1. Lossless Ethernet*

Traditional Ethernet is a “lossy” protocol, which means that data frames can be dropped or delivered out of order during normal operation, whereby the task of reconstructing the full data will be carried out by the higher protocol layers such as TCP/IP layer. As such Ethernet performs poorly in terms of latency compared to higher reliability protocols such as SAS, where packets cannot be dropped. In order to address this, a new Ethernet protocol standard was developed known as Lossless Ethernet [30] that provides features supported by other protocols, such as guaranteed packet delivery.

### 1.4.6 Network Convergence

The convergence of these multiple protocol networks (Ethernet, Fibre Channel, InfiniBand) into a unified protocol would give rise to reduction in equipment overhead and associated cost and power consumption, resulting in a reduction in total cost of ownership (TCO). Recently Ethernet has emerged as a possible candidate for network convergence due to its low cost and widespread adoption. Incremental steps towards network convergence include migration from Fibre Channel to Fibre Channel over Ethernet (FCoE) and the adoption of

RDMA (Remote Direct Memory Access) over converged Ethernet standards (RoCE) for high performance, low latency clustering applications.

## *1.5 International research projects targeting system embedded optical data communications*

During the period of this thesis, the author has been actively involved in a number of international collaborative research and development projects, which have been focussed on advancing the eco-system for system embedded optical interconnect in mainstream data communication environments. These include the SEPIANet project [31], the European PhoxTrot project [5], the European Nephele project [11] and the US HDPuG Optoelectronics project [32].

### 1.5.1 SEPIANet project

In mid 2011, a consortium of European organisations Xyratex, Fraunhofer IZM, ILFA, V-I Systems, Conjunct and TerOpta, entered into a 2.5 year EU funded collaborative research and development project called “System Embedded Photonics in Access Networks” (SEPIANet), which was launched as part of the European Commission’s PIANO+ funding scheme.

The aim of SEPIANet was to develop technology solutions for embedded optical architectures in access network head-end systems in order to allow both a significant reduction in power consumption and increased energy efficiency, system density and bandwidth scalability.

The SEPIANet project culminated in the successful development and demonstration of pluggable optical interconnect technologies for embedded planar glass waveguide OPCBs.

These activities will be detailed in Chapter 5.

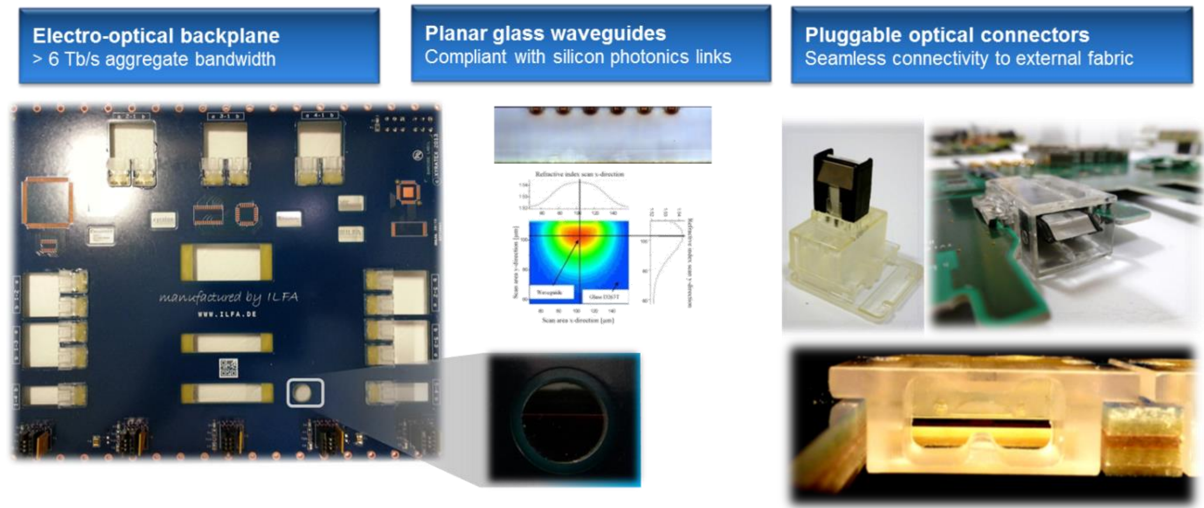


**Figure 1-13: SEPIANet project technology overview**

## 1.5.2 PhoxTroT project

PhoxTroT is an acronym of the full project title: "Photonics for High-Performance, Low-Cost & Low-Energy Data Centers, High Performance Computing Systems: Terabit/s Optical Interconnect Technologies for On-Board, Board-to-Board, Rack-to-Rack data links".

In October 2012, a large consortium of 18 European organisations led by Fraunhofer IZM as project coordinators with Xyratex as lead industrial partner entered into a 4 year European Commission FP7 funded collaborative research and development project called PhoxTrot [5].

The aim of this large scale "Integrated Project" is to develop an entire technology portfolio of cost- and energy-efficient Tb/s-scale on-chip, chip-to-chip, board-to-board and rack-to-rack level photonic interconnect solutions within data centre and HPC architectures (Figure 1-14).

**Figure 1-14: PhoxTrot vision: On-chip, chip-to-chip, board-to-board and rack-to-rack optical interconnect solutions deployed in PhoxTrot demonstration platform**

In particular the PhoxTrot portfolio includes the following key technology areas and platforms:

- Multimode polymers, for use in low-cost, high-performance optical PCB development (Figure 1-15)
- Single-mode polymer and glass waveguide technology platforms for high-end optical PCBs
- III-V material platforms, for use in active transceiver chip-scale circuitry
- Silicon photonics, based on the Silicon-on-Insulator waveguide platforms, which will be exploited for its attractive properties when CMOS-compatible chip-scale optical functions are required
- CMOS electronics and ASICs to deliver high speed and small footprint drive, amplification and routing functions
- Plasmonic interconnects to guide light along metallic stripes

Some of the board-to-board pluggable optical connector and waveguide coupler technologies developed on the SEPIANet project have been enhanced and will be deployed in more advanced demonstration platforms. The PhoxTrot project is due to end in October 2016.

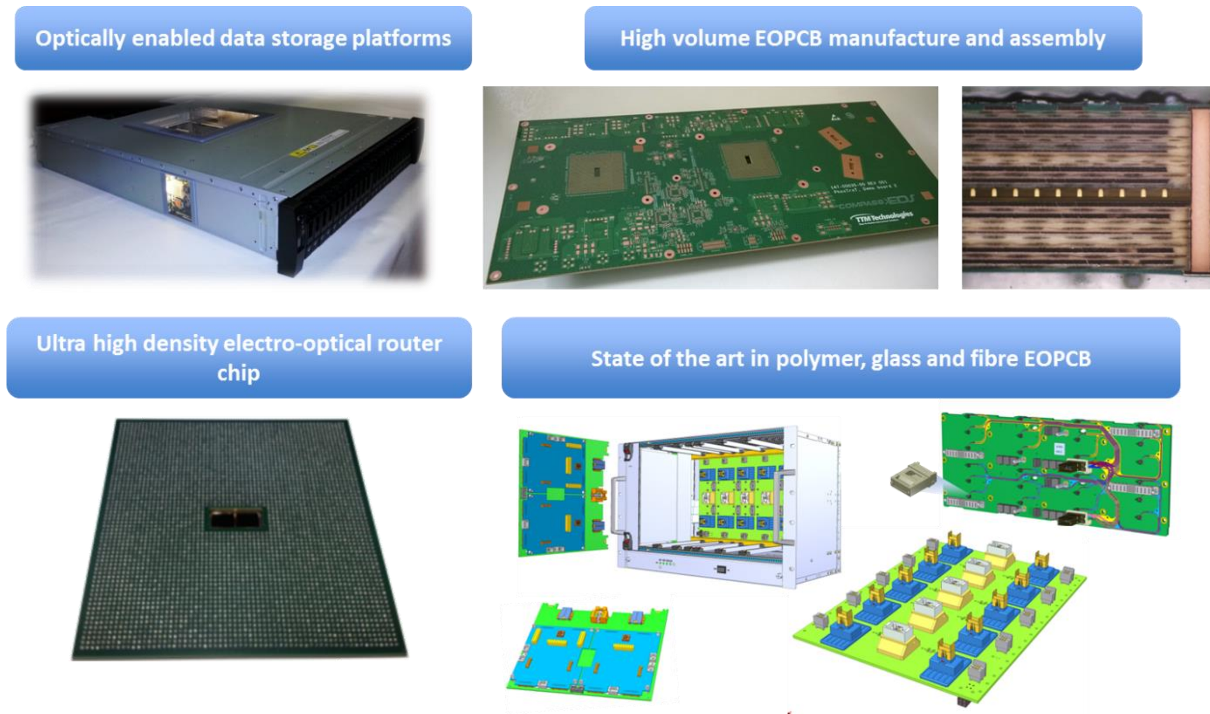


**Figure 1-15: Multimode optical interconnect deployed in systems, chips and OPCBs**

### 1.5.3 Nephele project

Nephele is an acronym of the full project title: "eNd to End scalable and dynamically reconfigurable oPtical architecture for application-awarE SDN cLoud datacentErs".

Nephele is a research project on optical data centre network technologies, supported by the Horizon2020 Framework Programme for Research and Innovation of the European Commission. The three-year project started officially on February 1, 2015 and brings together seven leading European universities, research centers and companies with National Technical University of Athens acting as coordinator.

The aim of the Nephele project is to develop a dynamic optical network infrastructure for future scale-out, disaggregated data centres (Figure 1-16). NEPHELE builds on the enormous bandwidth capacity of optical links and leverages hybrid optical switching to attain the ideal combination of high bandwidth at reduced cost and power compared to current data centre networks. To this end the project brings together research from multiple disciplines spanning data centre architecture protocols, network management software and firmware, optical switching technologies and system embedded interconnect technologies. The end-to-end development path of the project aims to bridge innovative research with near-market exploitation, achieving transformational impact in data centre networks that will pave the way to exascale infrastructures.

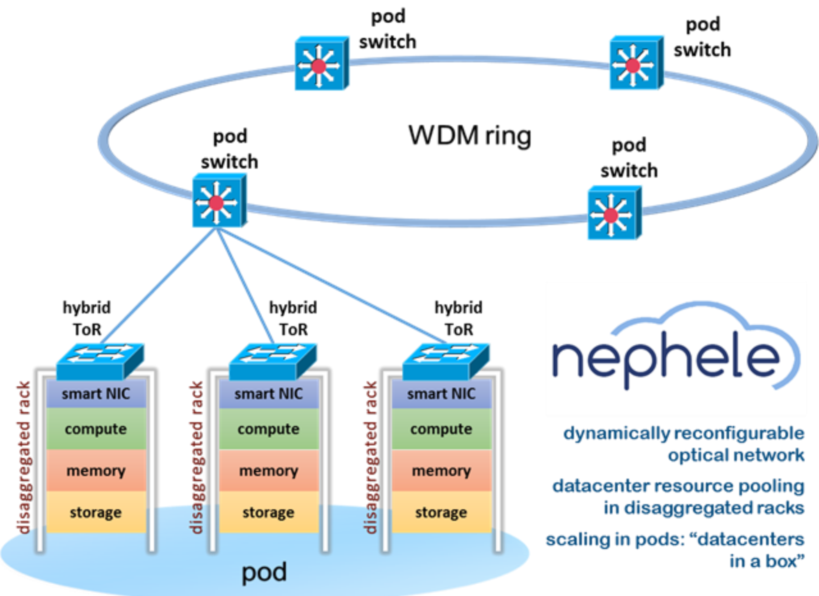


**Figure 1-16: Nephele hybrid electro-optical scale out data centre architecture**

### 1.5.4 HDPuG Optoelectronics project

The High Density Packaging user Group (HDPuG) is an international organisation with a broad membership comprised predominantly of industrial organisations. The purpose of HDPuG is to carry out industrial feasibility projects in packaging and assembly technologies. One of the HDPuG projects, "Optoelectronics", is tasked with evaluating the performance commercial feasibility of new electro-optical PCB technologies for intra-system high speed interconnect with particular emphasis on polymer waveguide based OPCB and connector technologies.
The HDPuG Optoelectronics project, which started in 2010, is coordinated by Cisco and PhoxTrot partner TTM Technologies, with design and test contributions from partners and contributors including Xyratex.
The importance of participation in the HDPuG Optoelectronics project is that it involves many large US organisations including Cisco, Juniper Networks, IBM, Boeing as well as major organisations from other parts of the world including Huawei (China) and Fujitsu (Japan) who have a direct interest in electro-optical PCB technology and could provide an accelerated path to market, depending on the outcome.

Phase 1 of the project involved the fabrication of small stand-alone electro-optical polymer waveguide based printed circuit boards, which were characterised with high speed 25 Gb/s test data. The embedded waveguides in the Phase 1 boards were not connectorised so could only be accessed through butt-coupling with optical fibres.

Phase 2 of the project, which started in 2014, sought to further develop connectorised optical waveguide backplanes and daughtercards with full board-to-board connectivity. Figure 1-17 shows samples developed on the HDPuG Optoelectronics project by different partners.

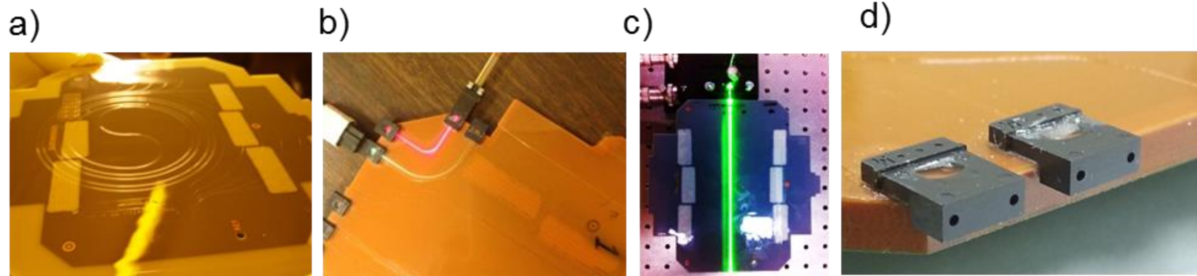


**Figure 1-17: HDPuG demonstrators a) 1.4 metre spiral polymer waveguide sample (source: Dow Corning), b) MT ferrule terminated optical waveguides with visible pink and green light illuminating 2 waveguides (Source: Optical Interlinks), c) OPCB with straight waveguide illuminated with green light (source: TTM Technologies), d) Two MT ferrules assembled onto the edge of a polymer waveguide based OPCB (Source: Optical Interlinks)**

## *1.6 Summary*

In this chapter, the structure of the thesis has been outlined and the research objectives and methodologies employed by the author explained for chapters 2, 3, 4, 5 and 6.

An introduction was provided to the emerging trends in data centre environments with emphasis on data centre systems, architectures and data protocol speeds, which are the main factors promoting a steady migration of optical interconnect technologies deeper down the hierarchal layers of the data centre and ultimately into the system enclosures themselves. A selection of international collaborative research and development projects seeking to advance the deployment of optical interconnect into future data centre environments was also introduced.

The next chapter describes the nascent commercial technology eco-system supporting this system level migration, which includes midboard optical transceivers, optical connectors and ultimately electro-optical circuit boards (OPCBs). The benefits of adopting OPCB technologies in future are explained with reference to the limitations of current electronic interconnect operating at increasing signal frequencies.

# 2 SYSTEM EMBEDDED OPTICAL INTERCONNECT TECHNOLOGY SURVEY

In-system bandwidth densities driven by interconnect speeds and scalable I/O within data storage and server enclosures will continue to increase over the coming years thereby severely impacting cost and performance in future data centre systems. The resulting increase in capacity, processing power, bandwidth density and bandwidth length product in and around data centre subsystems will severely impact design, cost and performance of future assemblies. However this could be mitigated by incorporating embedded optical channels into the backplane, motherboard and peripheral controller circuit boards. System embedded optical interconnect technologies have been the subject of research and development for many years to provide a cost viable "eco-system" to mitigate this impending bottleneck.

This chapter introduces emerging technologies enabling the migration of optical interconnect links into information and computing systems with particular emphasis on data storage and server systems, which form the lower tier, high volume equipment building blocks of modern data centres. In particular the state of the art in board-mounted transceiver technologies, optical connector interfaces and electro-optical circuit board (OPCB) technology will be discussed.

The design constraints associated with the arrangement of high frequency electronic signal transmission lines on PCBs will also be considered and the comparative merits identified of replacing high speed copper traces on the printed circuit board with embedded optical channels.

## *2.1 Introduction*

Prevailing trends in the data server and storage system industry [33], [34] are poised to severely impact the design of future data centre subsystems. Over the last decade, the volume of data being captured, processed, stored and manipulated as digital information has increased exponentially and this trend is set to continue**.** By 2020, 44 ZB of data will be created, of which 13 ZB will need to be stored, however the amount of data that installed capacity will be able to hold will only be 6.5 ZB [6].

The on-going exponential increase in data usage and storage is fuelled by business critical applications, email communications, multimedia networking applications and the emerging "Big Data" environments, where the amount of data exceeds the ability of traditional methods to manage, analyze and understand the meaning behind the data. With this tremendous growth in digital information, applications must become more data-intensive to accommodate the increased role of data analytics for decision management, storage performance requirements and the number and size of files. This means that applications will require more data throughput, higher levels of availability, more storage capacity and better response times to support applications.

Data storage and networking technologies have therefore advanced to address a growing diversity of data management requirements, which is expected to result in an increase not only in the volume and density of data storage devices manufactured to store this information, but also in devices with an increasing variety of physical shapes or form factors. These technologies include storage systems and subsystems, which run on a variety of high bandwidth data communication protocols including SAS, Fiber Channel, PCI Express, Infiniband and Ethernet.

As reported in Chapter 1, interconnect speeds based on the SAS point-to-point bus protocol, the dominant protocol governing data storage devices within data storage systems, will increase to over 24 Gb/s by 2017 [25], while the Infiniband protocol used predominantly for rack-to-rack communication in the switched fabrics inherent to enterprise data centres and high performance computers is projected to provide 50 Gb/s per lane by 2018 [27] with other protocols following similar bandwidth trends.

The exponential increase in system bandwidth and density required to satisfy this demand will impose unmanageable cost and performance burdens on future data centre technologies. In particular, the integration of more data intensive applications, such as servers, and the reduction in size and the increase in the number of high

speed ports of peripheral storage devices, such as hard disk drives, will cause the density of printed channels on the data storage midplane to go up, while the increase in data communication speeds will further expose the system to some of the fundamental constraints incurred when higher frequency data is conveyed along electronic channels. Many of these constraints can be mitigated to some degree, however at a mounting cost to the overall system design [35]–[38], [13].

The resulting performance bottleneck within the system could however be substantially reduced by conveying high speed data optically instead of electronically with the conversion point from electrical to optical interconnects (transceiver location) migrating ever closer to the on-board processing complexes, whether these are CPUs on a server blade, memory modules or data storage network switches.

This requires that optical channels be incorporated into the system, first in the form of discrete optical cables, but ultimately embedded into the system PCBs themselves.

## 2.2 *System embedded optical interconnect technologies*

The migration of optical interconnect into low cost, high volume data communication enclosures is now being made possible by the emergence of a new technology family for system embedded optical interconnect, which includes mid-board mountable optical transceivers and very high density parallel optical connectors and interfaces [39].

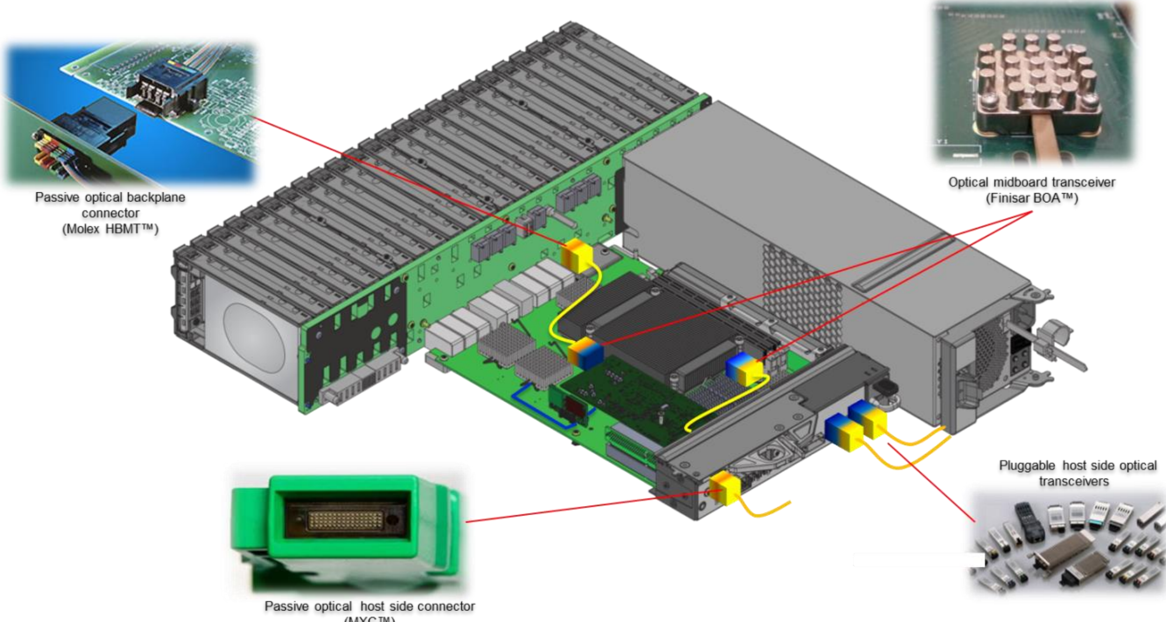


**Figure 2-1: Generic data storage enclosure with passive and active optical connectors**

### 2.2.1 Host side pluggable transceivers

Host side pluggable optical transceivers, which can be plugged and unplugged from the user accessible host side of the ICT system have been commercially deployed in ICT system for over 20 years [40].

The most common type of host side pluggable transceiver is the Small Form Factor Pluggable (SFP) transceiver, which supports one bidirectional channel. The SFP standard defines the package, electrical interface and optical interface of SFP transceivers.

XFPs (Extended Form Factor Pluggable transceivers are a higher speed version of SFPs, which can accommodate 10.3125 Gb/s data rates and include a Clock and Data Recovery (CDR) circuits in a larger package than the SFPs. With the introduction of SFP+, a 10 Gb/s capable transceiver in the original smaller SFP

package, but without intrinsic CDR circuitry, XFPs became for the most part obsolete from around 2013 onwards.

The increase in I/O density requirements started to give rise to parallel optical transceiver pluggable transceiver variants, the most common of which is the Quad Small Form Factor Pluggable transceiver (QSFP), which supports 4 bidirectional optical channels.

Pluggable transceivers with even higher port counts were developed including CFP, which supports 12 bidirectional channels, but these did not enjoy the same level of commercial deployment due for the most part to prohibitively high costs.

Active Optical Cables (AOC) are fixed assemblies comprising two host port pluggable optical transceiver assemblies connected by a fixed optical cable, which cannot be detached from the transceivers. The only accessible interface is the pluggable electrical interface connecting the transceiver to the host port of the ICT system. By removing the ability of the user to separate the optical cable from transceiver modules, the optical power margin can be optimised for the dedicated link in question and the characteristics of the optical transmitter, receiver and cable connection will be decided purely by the vendor. AOCs can support single bidirectional channels (adopting the SFP electrical interface and cage form factor), quad bidirectional channels (adopting the QSFP electrical interface and cage form factor) or 12 bidirectional channels adopting the CFP electrical interface and cage form factor). Host side pluggable transceivers are said to be “field replaceable”, that is a user, such as a data centre technician, may remove the transceivers from or attach the transceivers to the system during operation, without the need to disassemble or power down the system.

### 2.2.2 Passive optical connector interface

#### *2.2.2.1. Host side passive fibre optical connector interfaces*

Host side pluggable transceivers could be described as active optical connectors, in that they are pluggable connectors which also incorporate the optical conversion circuitry.

In contrast, passive host side optical connectors are interfaces, which are not immediately part of a transceiver [41]. They will convey optical signals from somewhere within the system in question either from a transceiver located somewhere within the system, which is not field replaceable or directly to an internal fibre-optic infrastructure.

Host side passive optical connectors have the advantage that they can accommodate a substantially higher density of optical channels than active connectors as they do not need to immediately provide space for the transceiver circuitry and packaging associated with each channel, in the same way as an active optical connector. Passive parallel optical interfaces based on the MT standard (Figure 2-2) will typically accommodate up to 6 rows of 12 optical channels per connector ferrule, whereby adjacent channels will have a centre-to-centre separation of 0.25 mm. MT ferrules are designed to house arrays of multimode or singlemode optical fibres. Optical connection between 2 MT ferrule interfaces is established through a physical contact between the 2 MT interfaces whereby each fibre in one MT interface is physically forced against the fibre in corresponding opposite position in the connecting MT interface. Typically the end facet of an MT ferrule will not be properly flat, but slightly rounded due to the polishing processes required as part of the fibre termination process.
In order to ensure that each connecting fibre pair in the connecting ferrules can make full physical contact with each other even when the connecting MT facets are not completely parallel, the fibres are arranged to protrude slightly out of the MT ferrule facet.
MT ferrules are by far the most common parallel optical connector interface.

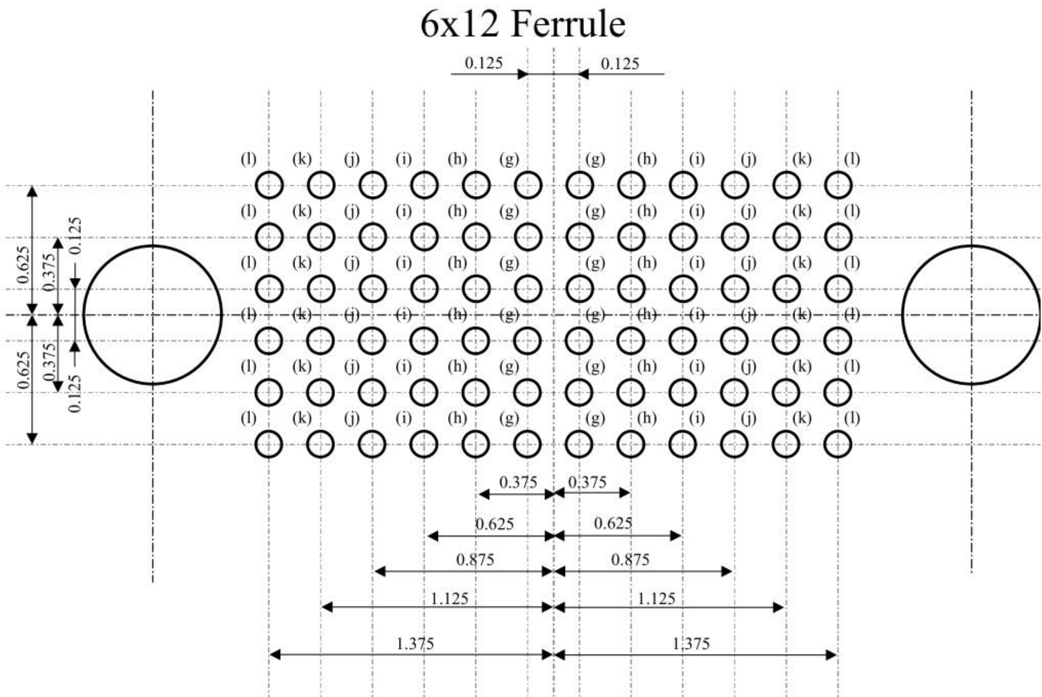


**Figure 2-2: MT parallel optical interface showing the positions of six rows of 12 small circular fibre channels relative to two large circular mechanical alignment slots on the left and right hand side**

One disadvantage of physical contact connectors is that the size of the optical channel at the connecting interface is the size of the fibre itself which is of the order of the size of typical dust or contaminant particles in a forced air environment. That is, a stray dust particle settling on the connecting interface could potentially block an entire optical channel or even multiple optical channels. In order to reduce the likelihood of this occurring, the physical fibre contact interfaces need to be cleaned after a small number of matings. This is manageable, when the number of connectors is small, but in a high volume environment, such as a data centre, including potentially hundreds of thousands or even millions of such interfaces, it would be prohibitive to enforce this requirement.
A new generation of parallel optical connector was developed by USConec in 2013 in collaboration with Intel and Facebook as part of the OpenCompute project to address the problem of scaling such connectors into future mega data centres. The expanded beam PrizmMT™ ferrules (Figure 2-3b) incorporate microlens arrays into the fibre holding structure to ensure that, at the exposed connecting interfaces, the optical beam width was actually increased to about 3.5 times the size of the multimode fibre aperture, thus making it far less susceptible to contamination. The MXC connector (Figure 2-3c), which formed a key part of the publicity drive surrounding the OpenCompute project houses a PrizmMT ferrule in a plastic shell and clip and is designed for host side access.

**Figure 2-3: a) MT ferrule, b) PrizmMT™ ferrule, c) MXC connector**

### 2.2.2.2. *Passive optical fibre board-to-board connectors*

In order to address the emerging need for board-to-board optical connectivity within the system enclosure, fibre-optic push pull type connectors have become commercially available. The most common application of such connectors is the need for daughtercard to backplane of midplane optical connectivity. A leading commercial product is the HBMT™ connector assembly offered by Molex.

### 2.2.3 Board-mountable optical transceivers

A highly relevant development is the emergence over the past 3 years of parallel mid-board mountable optical transceiver modules that can be mounted at any location on the PCB rather than constrained to the card edge.

This allows transceivers to be placed as close as possible to the electronic signal source (e.g. CPU, ASIC, expander) allowing electronic trace lengths and associated signal attenuation losses to be minimised and signal drive power reduced accordingly. Thus locating the optical engine close to the host chip not only reduces power consumption, but also improves the signal integrity compared to running signals over long copper traces on the host board and requires less post transmission electrical recovery at the target.

This migration of optical transceiver and supporting interconnect technologies into the system enclosure itself allows high speed electrical trace lengths to be minimised while increasing channel densities at the front fascia by over an order of magnitude through the use of passive host side pluggable optical connectors such as MTP or MXC. While current midboard transceiver technologies rely on fibre interconnect, efforts are underway to develop compliant interfaces to allow midboard transceivers to couple directly to waveguides embedded in the printed circuit board (PCB) [42].

The expected impact of this technology in the coming years is reflected in the strategic realignment of the major connector companies since 2010 toward parallel optical transceiver technology through acquisition of know-how, assets and development capability in III-V or silicon photonics based technologies.

**Table 2-1: Strategic purchases by optical connector companies of mid-board transceiver start-ups**

| Date | Vendor | Description |
|---|---|---|
| February 2010 | FCI | FCI purchased MergeOptics optical engine producers [43] |
| May 2010 | TE Connectivity | TE Connectivity purchased Zarlink, producer of parallel optical engines and subcomponents [44] |
| January 2011 | Molex | Molex purchased Luxtera's silicon photonics active optical cable (AOC) business [45] |
| April 2011 | Samtec | Samtec purchased AlpenIO – optical engine producers of high-speed active optical cables, optical engines and custom optical interconnects [46] |
| January 2013 | Volex | Volex purchased active optical technology from AppliedMicro [47] |

Crucially board-mounted parallel optical transceiver modules are becoming increasingly commercially available with major transceiver and connector vendors demonstrating product solutions.

Market survey organisation CIR projects revenues for board-mounted optical transceiver modules of $235 million by 2019 and reaching $775 million by 2020 [48].

The Avago miniPOD family was the first commercially available board-mounted optical transceiver module first supporting 8 Gb/s per lane in compliance with PCI Express applications [49]. This was followed by the first Finisar BOA module supporting data rates of 10.3 Gb/s per lane in order to target 10 GbE applications.

The increasing bandwidth demands rapidly led to a series of board mounted transceivers operating at 25Gb/s per lane [50], [51], [52] to primarily address 100G Ethernet applications.

Table 2-2 summarises the state of the art in board-mounted optical transceivers.

**Table 2-2: Survey of commercial midboard optical engine technologies**

| Vendor | Description | Channel number / bandwidth per channel | Image |
|---|---|---|---|
| Finisar | Board mounted optical assembly (BOA™)<br>Midboard optical transceiver module, with internal right-angled 2x12 optical lensed coupling interface to horizontal 24-way (2x12) MT ferrule and optical ribbon cable | 12 channel duplex<br>28.4 Gb/s per channel | |
| Avago | MicroPOD™ and MiniPOD™<br>Midboard optical transceiver module, with internal 1x12 vertical lensed coupling interface to top pluggable USConec right-angled lensed “Prizm Lighturn” connector cable | 12 channel singlex (separate Tx or Rx modules)<br>10.3125 Gb/s per channel | |
| Samtec | Firefly™ ECUO 38 AWG Fibre Optic Micro Cable Assembly<br>Midboard optical transceiver module, part of broader Firefly™ Microflyover interconnect range with internal fibre ribbon pigtail terminated in external MPO connector [53] | 4/ 12 channel simplex<br>28 Gb/s per channel | |
| Reflex Photonics | LightABLE™<br>Midboard singlex transceiver module with horizontally pluggable MT terminated multi-mode parallel fibre connector. Separate Tx or Rx modules [54] | 12 channel singlex<br>11.2 Gb/s per channel | |
| Ultra Communications | X80-Q Fury SMT Quad Transceiver<br>QFN surface-mount quad transceiver With GaAs VCSEL array and GaAs PIN photodetector (PD) array, glass lens array with | 4 channel duplex<br>12.5 Gb/s per channel | |

| | | | |
|---|---|---|---|
| | lens guide mechanisms for pluggable RVCON™ [55] ribbon fibre ferrule termination (proprietary to Ultra Communications). Attachable to PCB through reflow soldering | | |
| TE Connectivity | Coolbit™<br><br>Midboard optical transceiver module, with internal right-angled 2x12 optical lensed coupling interface to horizontal 24-way (2x12) MT ferrule and optical ribbon cable. Demonstrated at OFC 2013 but not yet commercially available | 12 channel duplex<br><br>28.4 Gb/s per channel | |
| FCI | LEAP™<br><br>Midboard optical transceiver module, with internal right-angled 2x12 optical lensed coupling interface to horizontal 24-way (2x12) MT ferrule and optical ribbon cable. Demonstrated at OFC 2013 but not yet commercially available | 12 channel duplex<br><br>28.4 Gb/s per channel | |

## *2.3 Performance limiters on high frequency electronic channels*

There are a number of factors, which will limit the operational bit rate of the high speed links on a commercial copper PCB [56], [57]. Dielectric absorption and Skin Effect are the key loss mechanisms on a copper trace [58], which cause an increase in signal attenuation with frequency, while electro-mechanical connectors introduce parasitic capacitance and inductance effects, and vias can act as impedance stubs giving rise to partial reflections in the signal path [59], [60]. At signal data rates of 24 Gb/s and higher, additional design measures will need to be taken including the use of lower dielectric loss PCB substrates, skew and loss controlled electro-mechanical connectors and enhanced via control techniques such as back-drilling or buried vias [61]–[63]. In order to mitigate rising crosstalk, copper channels will need to be moved further apart, however due to the spatial constraints of modern enclosure form factors for data communication systems, more complex routing patterns will be required and the number of high speed layers in the midplane PCB increased. The available space on the data storage midplanes is further restricted by the need for milled access slots to allow sufficient air flow through the system, while space on controller and other peripheral cards is consumed by ever increasing component densities as functionality is scaled.

### 2.3.1 Overview

In terms of PCB design and manufacture, the challenge of developing a backplane with such architecture becomes significant and costly. Copper traces have to be impedance controlled, which means that they have to be referenced to a ground plane, such that for each high speed routing layer there also has to be a reference plane, as well as layers for power and slow speed signals. Connectors have to be highly specialised in design and this also carries a significant cost adder. There are also considerable environmental effects which have to be managed and which play an important part in the performance of high speed copper interconnects.

There are a number of factors in a commercially produced copper backplane system which will limit the operational bit rate. The effect of ohmic losses and capacitances inherent in any system, as well as crosstalk and electro-magnetic emissions will put a high demand on the designer as bit rates increase. Dielectric absorption increases attenuation in proportion to frequency such that at bit-rates of 24 Gbps the track length would be limited to no more than a few inches or the use of more expensive substrates would have to be considered.

There are also more specific problems associated with backplane systems which have a requirement for multiple connectors in the signal path. Electro-mechanical connectors introduce parasitic capacitance and inductance

effects and PCB holes used for mounting can act as stubs. Manufacturing tolerances in the connectors cause skew and the environmental effects of system operation cause changes in the substrate. The relative permittivity and loss tangent of an FR4 substrate increase with both temperature and moisture.

Consider then the following operational factors which have a direct bearing on any comparison of optical and copper interconnects:

In scalable systems, the trend toward increased storage capacity, faster data processing and an overall reduction in system size will inevitably result in:

- Increased communication bandwidth within the system defined as the sum data rate of all high speed transmission lines in the system
- Increased density of communication channels defined as the ratio of the number of transmission lines on a board to the board area

As high speed electronic pulses travel along copper waveguides, they are subject to a number of effects which prove detrimental to signal integrity. This leads to a fundamental physical trade-off between the signal data rate, channel density and distance over which signals can propagate before irrecoverable degradation occurs. Some of these effects are now described.

There have been extensive studies on the comparative and projected performance of PCB embedded electronic vs embedded optical channels [38], [64]–[68].

### 2.3.2 Crosstalk

Crosstalk results from the coupling of signal energy from one channel to another. There are two principal mechanisms for crosstalk between adjacent copper traces, namely inductive coupling and capacitive coupling.

***Inductive coupling*** causes adjacent signal channels to interfere with each other across the interaction between each channel's magnetic field. These fields are generated by the movement of charge along the conductive traces.

***Capacitive coupling*** occurs between adjacent channels when signal energy is coupled from one conductive trace to another over the small capacitance which exists between them. The nature of capacitance is such that as the variation in voltage (due to high frequency signals) between the two traces speeds up, they will tend to short circuit.

In order to control crosstalk, one inevitably must place a limit on the separation between high speed channels [69]. This usually requires additional design efforts to optimise the layout of the signal traces and an increased number of high speed layers on the PCB.

### 2.3.3 Impedance mismatch

One of the principal concerns in high speed signal design is to ensure that the impedance between a high speed signal trace and its reference plane is maintained at every point along the trace and matches the impedance at the signal source and termination. This requires strict design control over the properties of the signal trace such as width, height, differential pair separation (if applicable) and distance to the reference plane. Erroneous features along the signal path, which deviate from the accepted geometries will cause localised mismatches in the impedance, which in turn will lead to partial signal reflections. Such features are known in the industry as "stubs" and can take many different forms. The most common type of stub is caused by vias, which are metallised holes in the PCB providing an electrical connection between different layers in the PCB stack-up. Typically, these vias pass right through the PCB from the top layer to the bottom layer the board, even if the required signal path, say between two internal layers passes along only part of the metallised inner surface of the via as shown in Figure 2-4. The result is that the section of the via over which the signal does not propagate causes a localised deviation in track impedance, which in turn gives rise to a partial reflection of the signal energy back along the trace.

Via stubs can be eliminated by physically removing the unwanted section of the via, through back-drilling of the via.

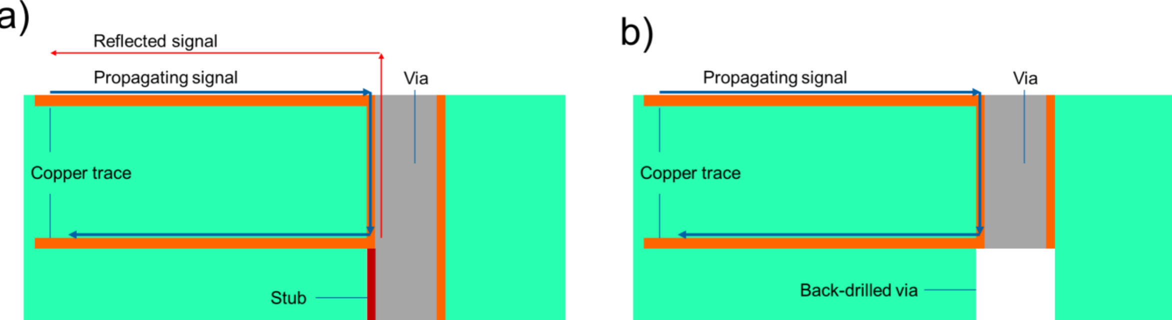


**Figure 2-4: PCB stack-up showing electrical signal propagation between a copper trace on the top layer and a copper trace on an intermediate layer, passing along part of a via. a) fully filled via causes the**

**exposed unused section to create localised impedance mismatch, which gives rise to a partial reflection, b) back-drilling of the via to minimise discontinuity and therefore partial reflection**

### 2.3.4 Signal attenuation

The greatest challenge to conveying high frequency signals over long trace lengths however is signal attenuation. Attenuation will always occur in traces due to the non-zero resistance of the trace. This can be reduced by increasing the cross-sectional size of the trace.

In addition to pure resistance based attenuation, there are strong frequency dependent loss mechanisms, which result in the higher frequency components of the signal experiencing higher attenuation, which in turn will distort the signal. The dominant frequency dependent loss mechanism is Skin Effect, which is the tendency of charge carriers to become more tightly concentrated around the surface or 'skin' of a conductor as the frequency of the signal increases. This concentration of charge into a smaller cross-sectional volume of conductive material, causes it to experience greater electrical resistance and dissipate energy more rapidly into the environment. This therefore places a greater limit on the distance over which high speed signals can propagate before the signal strength degrades beyond acceptable limits. Furthermore, the surface roughness of the copper trace contributes to Skin Effect, in that charge carriers will have further to travel, as they trace out a path along the contour of a rough trace compared to when they travel along a smooth trace.

In order to compensate for this effect, special alternative PCB materials can be used, such as Rogers® [70], the dielectric properties of which will reduce signal energy dissipation along the trace. Active signal conditioning devices such as pre-emphasis or equaliser circuits can also be included along the channel.

The cost and design effort required to remain within these constraints will be driven ever higher as conventional bandwidth requirements increase.

As systems become smaller, denser and faster, some requirements will simply not be physically enforceable. This trend toward greater system integration and transmission bandwidth will inevitably lead to system bottlenecks i.e. areas in which the constraints of electronic transmission cannot be overcome.

The first of these bottlenecks is expected to emerge on the system backplane, which will typically need to accommodate the highest density of high speed channels.

## *2.4 Board-level optical interconnect*

It is proposed that the projected performance bottleneck in ICT systems is mitigated by conveying high speed signals optically rather than electronically even at the system PCB level. In order to achieve this one would need to adopt electro-optical printed circuit board (OPCB) and interconnect technology on the midplane or backplane [71]–[76]. As shown in Figure 2-5 the midplane would thus comprise a PCB with both copper layers for electrical power and low speed electronic signal distribution and one or more optical transmission layers to allow high speed signals to be conveyed optically. There are various types of optical waveguide that can form the basis of such optical transmission layers such as laminated fibre-optic flexible circuits, planar glass optical waveguides or planar polymer optical waveguides. The research has primarily targeted the deployment of planar polymer waveguide technology in electro-optical PCBs and the development of commercially viable pluggable optical connection technologies.

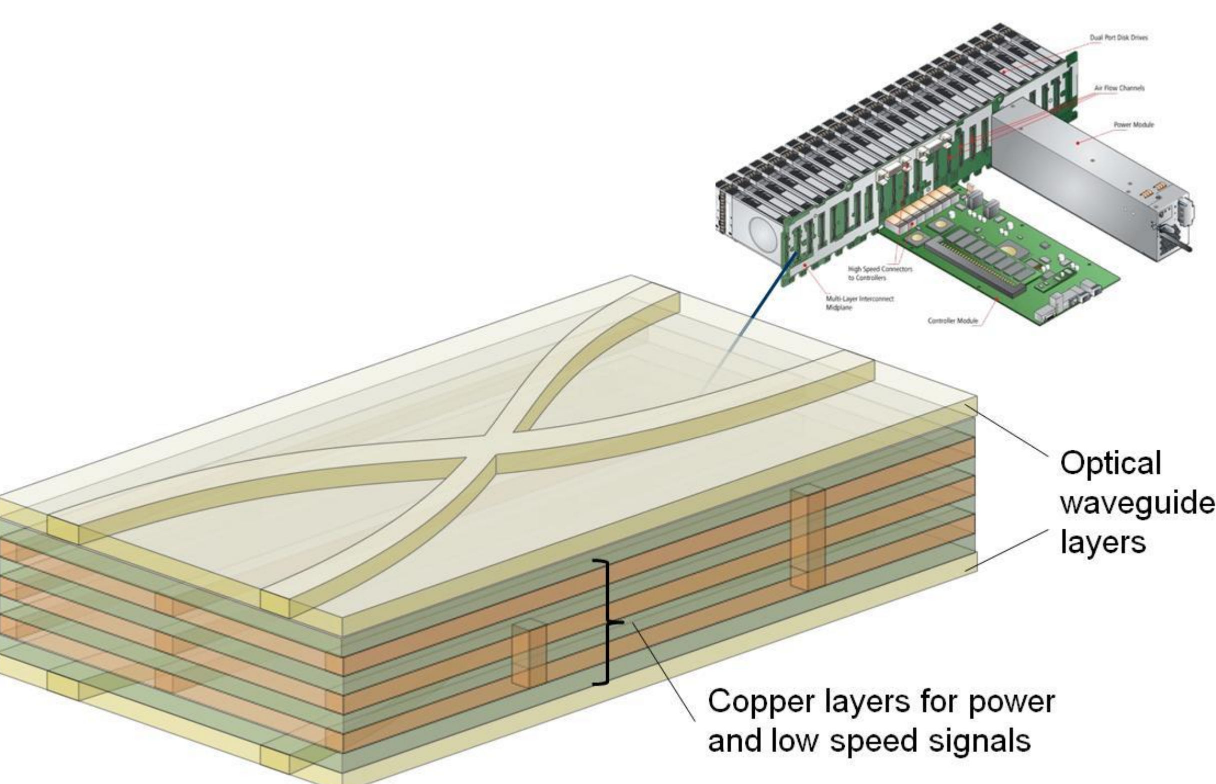


**Figure 2-5: Schematic view of electro-optical midplane architecture comprising copper layers to accommodate power, static and low speed signal distribution and at least one optical transmission layer to convey high speed signals along optical waveguides**

#### *2.4.1.1. Optical circuit board overview*

Although the prospects for commercial proliferation of optical circuit board (OPCB) technology were crippled by the slow-down in the telecoms sector following the stock market crash in 2001, embedded waveguide based OPCB technology advanced substantially during the first 15 years of the 21st Century.

While significant advances have been made in embedding conventional optical fibres onto printed circuit boards [77], focus across Europe and Japan has been on the fabrication of transparent channels (waveguides), which are integrated directly into the printed circuit board substrates. Research by Dangel, Van Steenberge, Penty, Chappell, Shibata, Ishigure, Schröder and Doany [73], [76], [78]–[83] has collectively demonstrated a wide range of waveguide fabrication techniques and materials.

Optical materials are now available, which offer the required resilience to thermal cycling and humidity to allow them to be integrated into PCBs. A variety of waveguide fabrication techniques have evolved, which lend themselves to high volume production (photolithography, batch-processing) or low volume or prototype development (laser direct imaging). Graded index waveguide profiles can now be fabricated in both glass [84] and polymer [85] to offer reduced modal dispersion, crosstalk and radiative losses though the side walls. Waveguide connector termination, considered the final technical barrier to OPCB commercialisation, had reached the stage by 2015 where low termination losses have been demonstrated both in glass [4] and in polymer [81].

There are three primary classes of electro-optical circuit board available at different technology readiness levels: 1) fibre-optic flexible circuits [77], 2) embedded planar polymer waveguides [1], [86] and 3) planar glass waveguides [87] (Figure 2-6). Each interconnect class offers different advantages making them suitable in different applications.

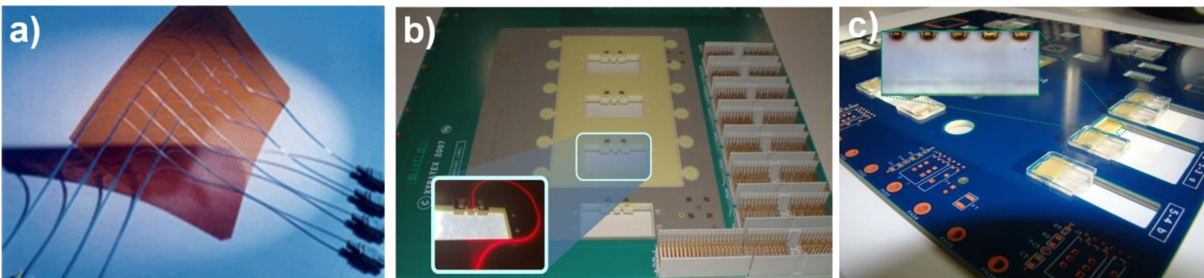


**Figure 2-6: Electro-optical PCB interconnect technologies a) Fibre-optic flexible circuits, b) PCB embedded polymer waveguides, c) PCB embedded planar glass waveguides (Source: Fraunhofer IZM)**

#### *2.4.1.2. Fibre-optic flexible circuits*

Laminated fibre-optic circuits, in which optical fibres are pressed and glued into place on a substrate benefit from the reliability of conventional optical fibre technology. However these circuits cannot accommodate waveguide crossings in the same layer i.e. fibres must cross over each other and cannot cross through each other. Also with each additional fibre layer, backing substrates must typically be added to hold the fibres in place, thus significantly increasing the thickness of the circuit. This would limit the long term usefulness of laminated fibre-optic circuits in PCB stack-ups. At best they can be glued or bolted onto the surface of a conventional PCB. Fibre-optic circuits were deployed on the data storage system demonstrator developed by the author as part of the PhoxTroT project [88] .

#### *2.4.1.3. Planar polymer waveguides*

The incorporation of multi-mode polymer waveguides into PCB stack-ups has been demonstrated extensively over the past decade [39], [73], [76], [89]–[92]. While historically these waveguides have been constrained to stepped refractive index profiles with higher NAs than conventional multi-mode fibres, recent advances in polymer waveguide materials and fabrication techniques have opened the door to graded-index waveguides with excellent ageing properties and better transmissivity at longer infra-red wavelengths [81], [93], [94] making them suitable for short reach on-board or board-to-board interconnect within a system enclosure.

Conventional step-index multimode polymer waveguides would be challenged over longer high speed links, in which modal dispersion would limit performance compared to graded index waveguides. However recent work by Bamiedakis and Chen at the University of Cambridge [95], [96] has shown that long step index polymer waveguides are subject to less modal dispersion than anticipated and can convey a 40 Gb/s signal along a 1 metre waveguide without detrimental dispersion [97]. It should however be noted that due to the spiral test pattern, the higher order modes of propagation are more likely to have been coupled to the radiation modes in a process known as “mode stripping”. As these higher order modes are the strongest contributors to modal dispersion, the result of their being filtered out would provide a “cleaner” signal than would be expected on a straight waveguide of the same length, though at the expense of higher insertion loss. Furthermore, modal dispersion is also strongly determined by the manner in which the optical signal is launched into the waveguide, thus a singlemode launch into the fundamental mode will give rise to the lowest dispersion, while conversely a modally filled launch will give rise to the highest dispersion.

Polymer waveguides exhibit a high transmissivity at shorter communications wavelengths such as 850 nm, but

would be unsuited to convey other operational wavelengths (1310 nm or 1550 nm) over longer distances, due to higher intrinsic absorption losses, though this can be mitigated in some polymer formulations [98]. However they would be suitable for very short reach, versatile, low cost links such as inter-chip connections on a board. They would also be suitable for applications in which certain properties of the polymer such as thermo-optic, electro-optic or strain-optic coefficients could be used to support advanced devices such as Mach-Zehnder switches or long range plasmonic interconnect. A comprehensive overview of polymer waveguide materials, fabrication and devices is provided by Ma and Dalton [99].

Electro-optical circuit boards with embedded polymer waveguide layers were developed by Xyratex in collaboration with IBM Research and Varioprint in 2008 [100]. Planar polymer waveguide connector technologies form the basis of this thesis and are described extensively in Chapter 3. They were also deployed in a joint demonstration system developed by the author at Xyratex in conjunction with Finisar, Vario-optics and Huber+Suhner in 2012 [101].

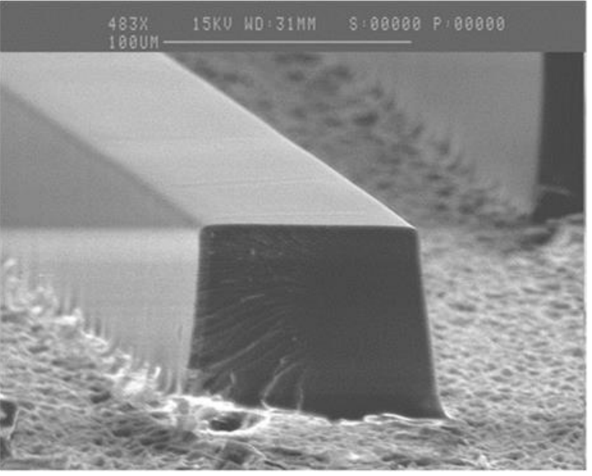


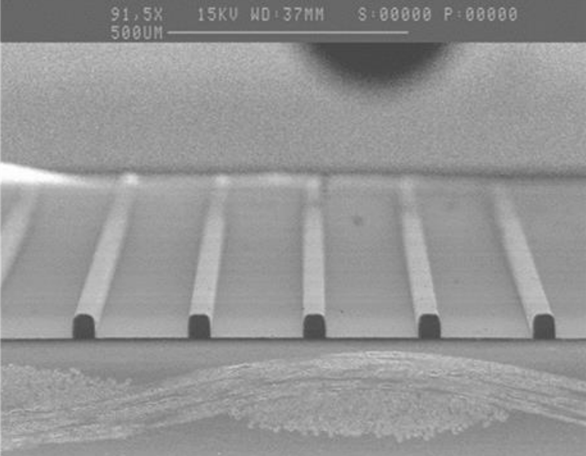


**Figure 2-7: Scanning Electron Microscope image of exposed 50µm x 50µm polymer waveguide core without upper cladding (left) and array of exposed parallel polymer waveguide cores (right)**

### *2.4.1.4. Planar glass waveguides*

Planar glass waveguide technology could combine some of the performance benefits of optical fibres, such as lower material absorption at longer operational wavelengths such as 1310 nm and 1550 nm and lower modal dispersion with the ability to fabricate dense complex optical circuit layouts on single layers and integrate these into PCB stack-ups. OPCBs with embedded planar glass waveguides were developed by Fraunhofer IZM, ILFA

and Xyratex as part of the SEPIANet project [102], with Fraunhofer IZM fabricating the planar glass waveguide foils, ILFA laminating the glass foils into a PCB stack and Xyratex developing both passive fibre-to-board and board-to-board connectors as well as a complete optical backplane connector demonstration platform, characterised at both 850 nm and 1301 nm, which is described in Chapter 5.

For densely populated data communication modules, board embedded waveguides offer crucial advantages over stacked and grouped fibre patchcords or laminated fibres. Primarily, as it would be prohibitive to embed fibre laminates into the PCB substrate directly, they would need to be routed over the board surface, thus strongly reducing the area available for component packages. In comparison optical waveguide layers embedded in the PCB would not limit the layout of component packages on the board surface, though they would incur some additional design constraints to the placement of vias or other through-hole structures.

Preferably, such board-embedded waveguides could be coupled directly to external optical cables, through pluggable interfaces, and would need to thus have comparable interface parameters to those of optical fibre, such as numerical aperture (NA), core size, core shape and refractive index profile in order to minimise coupling loss. OPCBs based on embedded glass waveguides have been demonstrated [103], [104] and would further combine the performance benefits of graded index optical fibers, such as lower material absorption at operational wavelengths used for longer reach fiber optic networks (>1260 nm) and lower modal dispersion as compared to step-index waveguides, with the ability to accommodate dense complex optical circuit layouts on single layers and integrate these into PCB stack-ups. Crucially, this would enable direct, “seamless” optical connectivity from an external fabric to board-embedded optical channels [4]. Although the multimode planar glass waveguides described are graded index and thus have lower modal dispersion than a step-index multimode waveguide, singlemode waveguides will have substantially less dispersion, and so ultimately the low cost fabrication of panel level singlemode glass waveguides is an important goal of fabricators in this area.

#### *2.4.1.5. Free space optics*

There has also been research in the past looking at using free space optics to provide short reach connectivity within system enclosures. In 2006, the EU funded HOLMS (High Speed Opto-electronic Memory Systems) project proposed the use of a 3D free-space optical interconnect module to optically connect different parts of the system [75]. In 2009 Hewlett Packard proposed two schemes for in-system free space optical communications. The first scheme was based on the use of telecentric lenses at the transmitter and receiver side providing an expanded beam solution that allowed very large misalignment tolerance of ± 2 mm. The second scheme made

use of magnetically coupled proximity free space modules [105].

## 2.4.2 Milestones in polymer OPCB research and development

### *2.4.2.1. MT terminated OPCBs*

IBM Research demonstrated the first OPCBs with an MT pluggable optical connector for communication with other boards/instruments. The interconnection of two such boards via a ribbon cable was presented in [106] with OPCBs exchanging data at 120 Gb/s (12x10Gb/s) aggregate bit rate. Also 12.5Gb/s data were successfully transmitted via a 100cm spiral shaped polymer waveguide (0.05dB/cm losses @850nm) formed on a PCB.

### *2.4.2.2. Passive alignment of optical elements on OPCBs*

IBM Research developed a technique for passive alignment of optical elements onto an OPCB by creating structures in a copper layer during manufacturing, which were used as a positional reference for polymer waveguide fabrication and for the formation of mechanical alignment features [107] as shown in Figure 2-8.

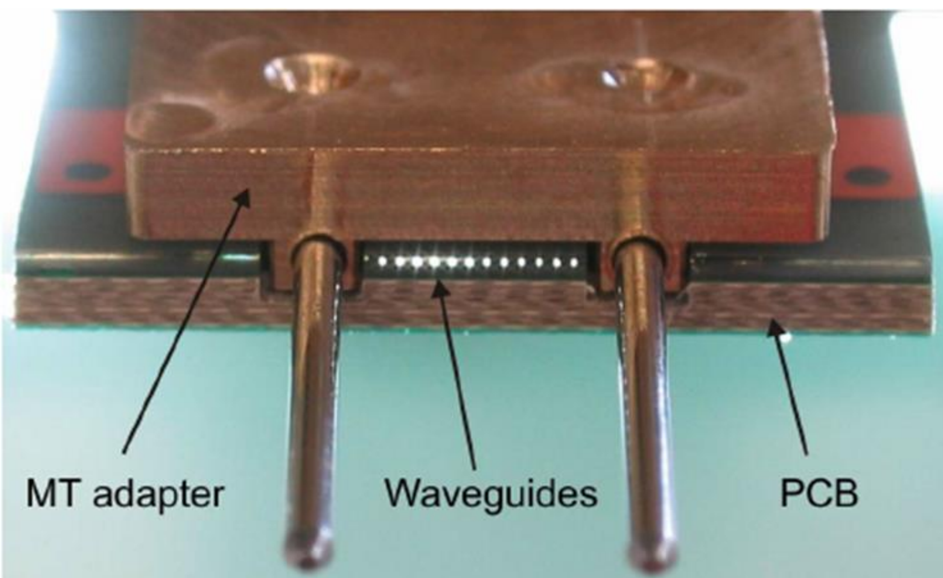


**Figure 2-8: Male MT adapter passively aligned and assembled onto an OPCB (Source: IBM Research** [107])

In this thesis, it will be described how the author has improved on this method by inventing and realising a simpler and more precise technique for passively assembling optical components onto an OPCB, whereby

mechanical registration features are developed in the same process step as the optical waveguide core layer itself. This is described in Chapter 3.

### *2.4.2.3. Direct connection of optical transceiver interfaces to OPCBs*

IBM Research developed a demonstration card with a fixed transceiver interface with MT compliant interface mounted on a flexible substrate, which was butt coupled to the embedded polymer optical waveguide interface of an OPCB [108] as shown in Figure 2-9 a). In Chapter 3 and Chapter 4 of this thesis the author describes two generations of pluggable optical transceiver interface, which allow the optical transceiver interface to be butt-coupled to the embedded polymer optical waveguide interface of an OPCB as shown in Figure 2-9 b).

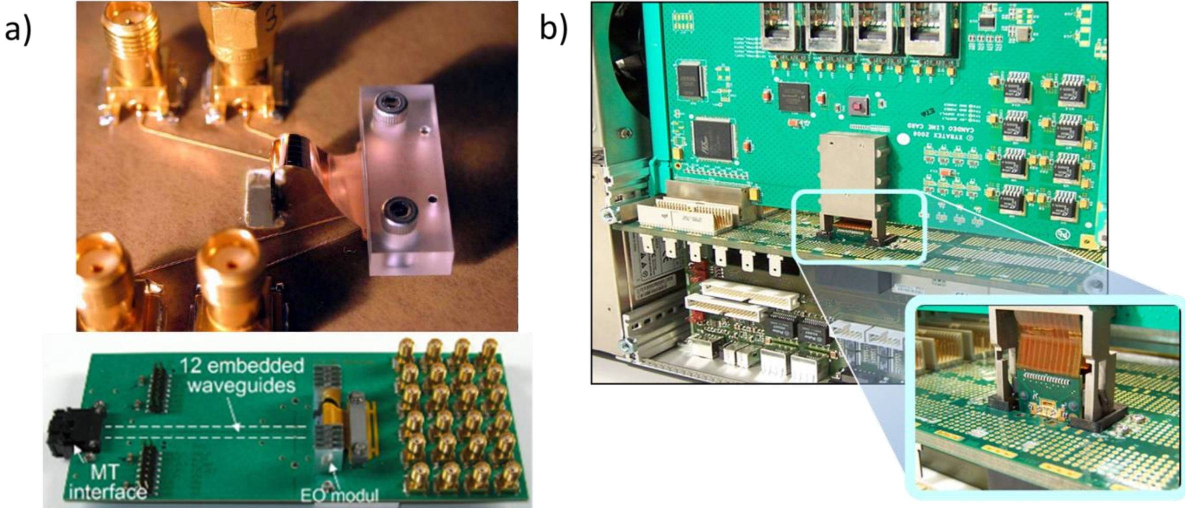


**Figure 2-9: Optical transceiver connectivity to OPCBs a) active optical transceiver interface in a fixed butt-coupled arrangement to embedded polymer waveguide interfaces in OPCBs (Source: IBM Research** [108]**), b) active optical transceiver and pluggable connector mechanism developed by author and described in Chapter 4, allowing optical transceiver interface to be selectively butt-coupled to embedded polymer waveguide interfaces on OPCB during engagement of a daughter card**

### *2.4.2.4. Adiabatic coupling to silicon photonics chips*

Polymer waveguides have been demonstrated to provide low loss coupling between fibres and photonic integrated circuits through the adiabatic or evanescent coupling method, whereby singlemode polymer

waveguides are placed in near physical contact with tapered silicon waveguides and a reciprocal coupling of optical energy occurs between the two [109].

#### *2.4.2.5. Singlemode waveguide fabrication and connectorisation*

IBM have advanced laser direct writing as a preferred means of fabricating singlemode polymer waveguides [110]. Swiss company Huber Suhner demonstrated connectorisation of singlemode polymer waveguides through a passive alignment method based on the principle of using exposed waveguide core features, which were developed in this project [111].

#### *2.4.2.6. Dense OPCB non-blocking architecture*

In the UK, the Centre for Advanced Photonics and Electronics (CAPE), part of the University of Cambridge, has been leading polymer waveguide research and fabrication activities using polysiloxane material provided by Dow Corning. They developed a passive OPCB incorporating 100 embedded waveguides, which is able to connect to 10 other PCBs through a non-blocking architecture [112].

#### *2.4.2.7. Graded index polymer waveguides*

Historically it has been exceedingly difficult to fabricate graded index polymer waveguides with a similar numerical aperture (difference in refractive index maximum in core and cladding) to graded index fibres. It is only recently that the fabrication of high performance graded index polymer waveguides was successfully spearheaded in Japan by Keio University and Sumitomo Bakelite. Keio University demonstrated lower crosstalk between adjacent channels using an OPCB with sixteen parallel graded index MM waveguides as presented in [113] exhibiting more than 5 dB inter-channel crosstalk reduction compared to step index waveguides.

#### *2.4.2.8. OPCBs with integrated optical transceivers*

Furukawa demonstrated 12 channel optical modules coupled directly to PCB embedded polymer waveguides [42] while Hitachi demonstrated an OPCB with dual polymer waveguide layers and a 1 Tb/s, 48 channel optical transceiver assembled onto it, coupling directly to embedded waveguides operating at 20 Gb/s per lane [114].

### *2.4.2.9. Self-written optical waveguides*

Yoshimura at Tokyo University of Technology demonstrated self-writing waveguides by launching 448nm UV light from a fibre into a photocurable optical polymer. As the polymer cured in response to the Gaussian intensity distribution of UV light emanating from the fibre, it produced a corresponding Gaussian dome, which acted to guide and collimate the UV curing radiation. This allowed a long waveguide to be written directly [115].

## 2.5 *Design attributes of embedded optical channels*

There are many design benefits to replacing the high speed electronic copper traces on a data storage midplane with optical traces, however the limits on in-plane bend radius are still unacceptably high for practical design purposes.

### 2.5.1 Density comparison between printed copper and polymer waveguides

The maximum permissible density of copper transmission lines is determined by the crosstalk incurred between electronic channels.

Most PCB layout engineers will adhere to "20*H" design rule [116], whereby the separation between high speed traces must be at least 20 times the height of the high speed traces over the reference plane on an adjacent layer, which is used to provide the return path for the current. This distance between traces and reference plane is commonly annotated as H in PCB trace simulation tools, such as those offered by Polar Instruments [117].

A typical distance H between a high speed trace and a reference plane in modern printed circuit boards will be around 125µm. Therefore the minimum separation between adjacent electronic transmission lines required to convey data at around 24 Gb/s should typically be no less than about 2.5 mm. This helps prevent the coupling of the signal onto adjacent wires and also helps prevent free radiation of the signal from the edge of the PCB [118]. However conventional parallel optical layouts can accommodate centre to centre horizontal separations between optical channels of 250 µm. Furthermore multiple optical layers can be accurately stacked to also allow vertical 250 µm separations between optical channels as demonstrated by Betschon [73].

In Figure 2-10 we compare a cross-section of an electrical PCB with high frequency electrical differential transmission lines arranged as tightly as possible across multiple layers using the 20xH rule and a cross-section of an optical PCB with optical waveguides arranged as tightly as possible across multiple layers with a horizontal and vertical centre-to-centre separation of 250 µm. In the same cross-sectional area occupied by 3 electrical differential channels, 39 optical waveguides can be arranged, thus indicating a 13 fold increase in channel density when multilayer high speed copper traces are replaced with multilayer polymeric optical channels. An increase in channel density will allow both for a reduction of the functional horizontal area of the PCB and a reduction in layer count. It should be noted that optical waveguide pitches as low as 62.5 µm have been demonstrated [72]. Assuming that only the horizontal waveguide pitch is reduced from 250 µm to 62.5 µm, the density improvement would increase by a factor of four. Further scaling in density improvements would also come with reduction in vertical channel pitch, though this would be harder to achieve. Optical fibres have larger

pitch restrictions than waveguides as each fibre core is surrounded by a cladding. The standard cladding diameter for OM2, OM3 and OM4 fibres is 125 µm, though some modern fibres are being introduced with smaller cladding diameters of down to 80 µm. This means that parallel fibre arrays with a channel pitch of 125 µm are possible and have been proposed for development [119]. While waveguides can have smaller pitches than optical fibres, any waveguide interface to fibres would need to fan out to the fibre channel pitch at least in the interface region.

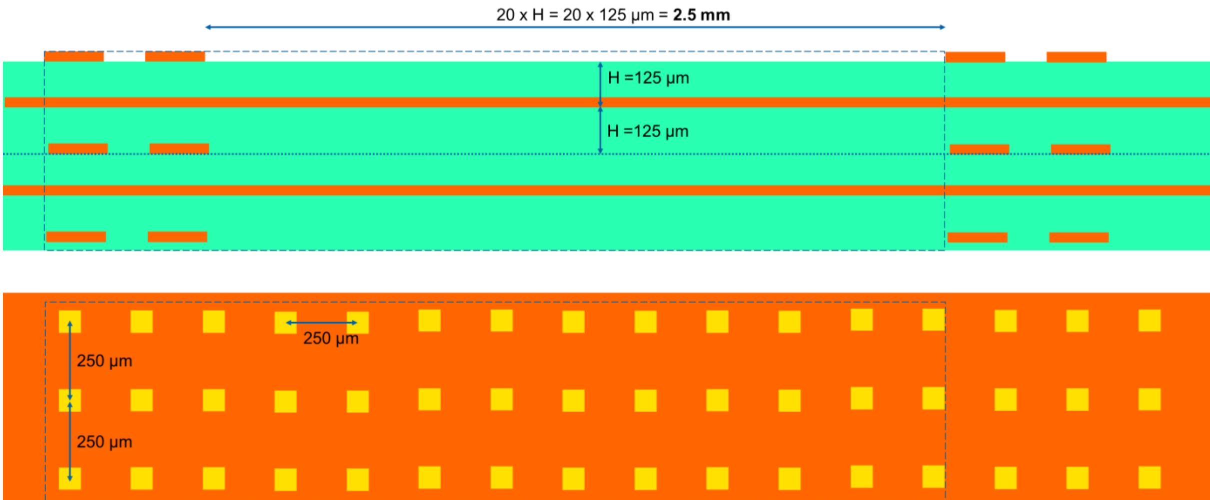


**Figure 2-10: Channel density comparison between 24 Gb/s copper differential trace and polymeric optical channels**

The 20 * H rule is however very conservative. In reality, there are many factors, which influence the minimum separation between adjacent high speed channels. Crosstalk is a cumulative effect and will increase with the distance over which the channels in question remain adjacent. Often the PCB layout designer can limit the distance over which a pair of channels remain adjacent, by moving said channels away from each other after given distances, so much shorter inter-channel separations could be tolerated over shorter distances.
Also, traces on the PCB surface (e.g. microstrip) will be more susceptible to crosstalk from adjacent traces compared to traces on internal layers of the PCB (e.g stripline).

#### *2.5.1.1. Environmental benefits of embedded optical waveguides*

Optical waveguides neither produce nor are affected by electro-magnetic (radio frequency) interference and are therefore not constrained by Electromagnetic Compatibility (EMC) regulations that impose a severe cost burden on the design of high speed copper PCBs. The layout advantages offered by optical waveguides will give rise to

a reduction in functional area and layer count of the PCB. This level of reduction will strongly depend on the application with the more IO intensive applications subject to the greatest potential reduction in PCB volume. This will help to meet the prevailing industrial trends of system integration and miniaturisation. This will also give rise to substantial reductions in PCB waste materials including copper, FR4, solder resist and laminating adhesives.

During the PCB design stages of the system backplane developed as part of the FirstLight demonstration platform described in Chapter 4, a design feasibility study was carried out and it was estimated that by removing the 96 high speed electrical signal traces on a typical 24 drive storage system, the layer count of the PCB in the case of a 4U midplane could potentially be reduced by over 50% and the open area available for airflow increased by over 20%. The reduction in thickness and material area can give rise to a corresponding reduction in total PCB material by over 60%, assuming the mechanical rigidity of the PCB will not be compromised beyond the requirements of the system. However separate reinforcing structures, such as metal bracing structures can be put in place in order to provide the necessary mechanical support to the PCB if required.

This is not a reliable estimate as it is highly dependent on the immediate functional and design requirements of the system, which evolve year on year as the products advance. However given that the product requirements evolution is dominated by increasing data rates and increasing numbers of high speed links with tighter routing restrictions, these estimates can be considered a conservative lower bound on the material savings possible.

### *2.5.1.2. Crossovers*

The most crucial advantage of laying out optical waveguides in lieu of electronic transmission lines is that while copper traces can only be routed across each other by redirecting at least one trace along a different PCB layer through the use of bridging vias, optical waveguides can intersect each other on the same layer. This means that the number of layers devoted to optical transmission is only limited by spatial constraints. To minimise crosstalk between crossed waveguides, the waveguides should be crossed orthogonally, however a range of crossing angles can be tolerated depending on the loss budget of the waveguides [120]. Figure 2-11 shows both orthogonal and non-orthogonal waveguides.

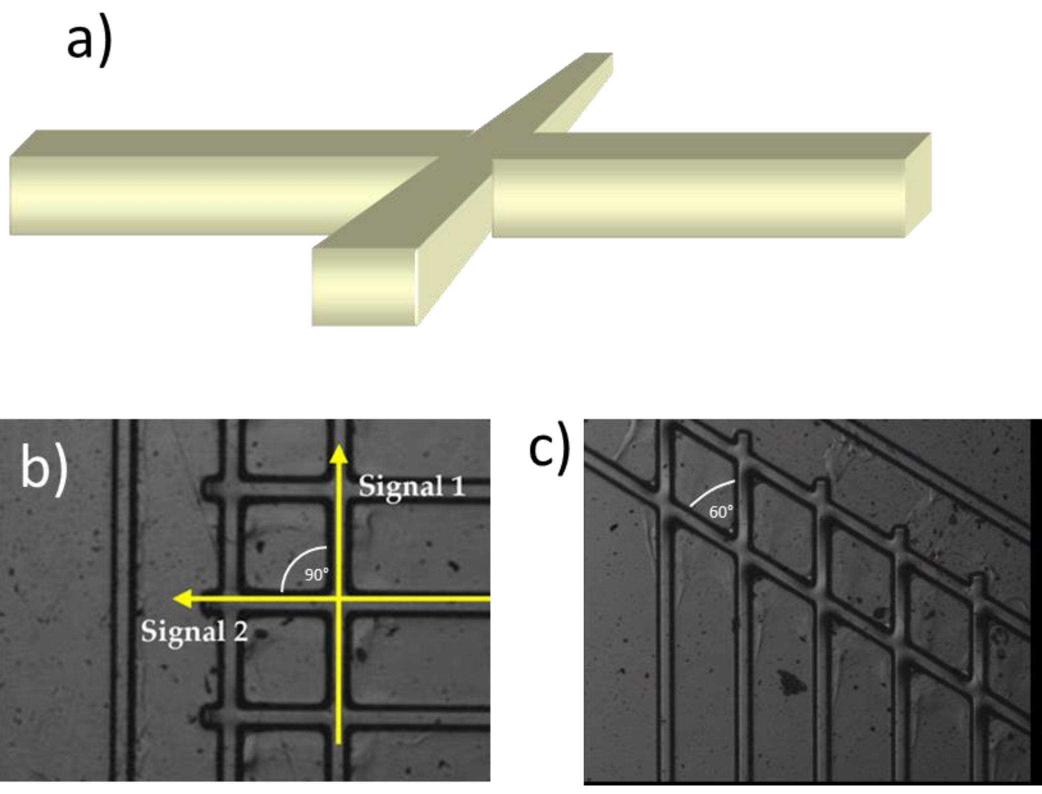


**Figure 2-11: Waveguide crossovers a) Schematic view of waveguide crossover, b) Photo of 50µm x 50µm polymer waveguides patterned photolithographically with multiple 90° crossovers (Source: Exxelis Ltd), c) Photo of 50µm x 50µm polymer waveguides patterned photolithographically with multiple 60° crossovers (Source: Exxelis Ltd)**

#### *2.5.1.3. Passive high speed splitters*

Optical waveguides can be split into multiple branches to allow division of the signal power along each waveguide branch. The number of branches into which a waveguide can be divided would depend on the available power budget and the loss characteristics of the waveguide. However passive division of high speed electronic traces cannot be reliably achieved and would usually require an active device such as a crosspoint switch. Figure 2-12 shows images of 1x2 splitters.

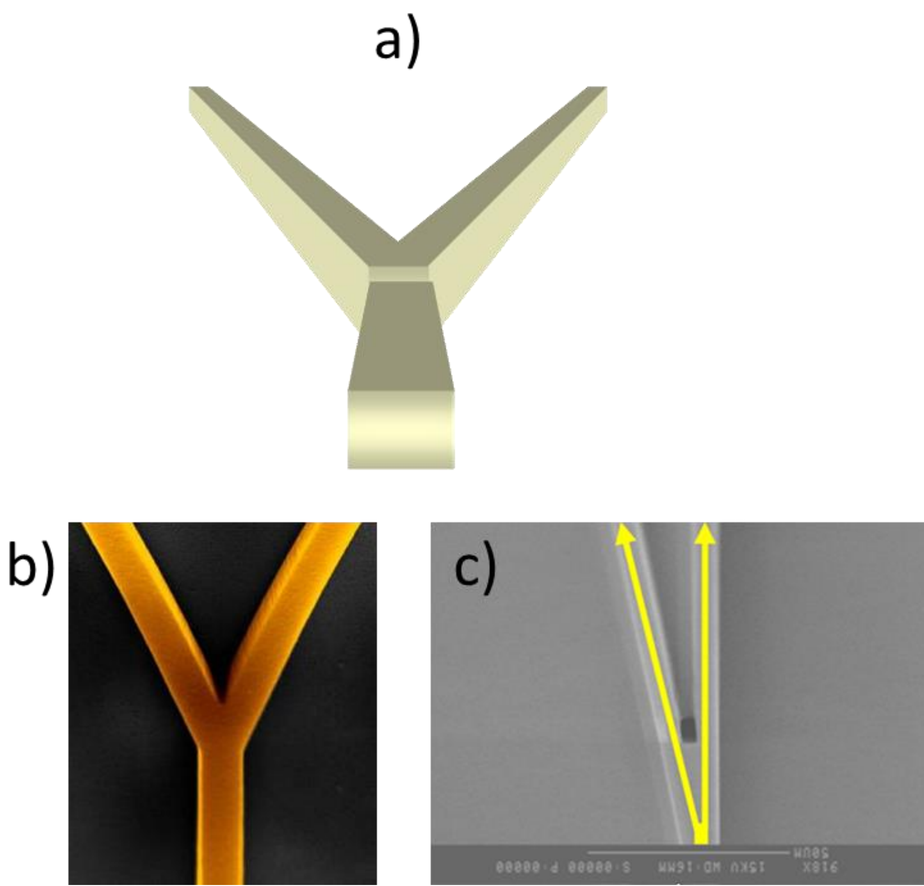


**Figure 2-12: Waveguide splitters a) Schematic view of 1 x 2 waveguide splitter, b) Photo of 50µm x 50µm polymer waveguides with 1 x 2 symmetric Y-branch splitter (Source: IBM Research), c) Photo of 50µm x 50µm polymer waveguides with 1 x 2 splitter (Source: Exxelis Ltd)**

### 2.5.2 Bend radius routing constrictions

In spite of the substantial design benefits associated with embedded polymer optical architectures one severe constriction has yet to be overcome, namely the minimum bend radius in the plane of the interconnect layer. Figure 2-13 shows a photo of a 90° polymer waveguide arc with a tight 10 mm bend radius illuminated with

visible 650 nm light. The light is scattering strongly from the bend. The minimum in-plane bend radius of a waveguide will be determined by many factors including critical angle, difference between core and cladding refractive indices, core size and shape and size wall roughness. Characterisation of optical propagation losses in printed multimode polymer waveguides had shown that in-plane bend losses due to scattering of higher order modes will realistically restrict the bend radius to more than 15 mm [120]. In practical terms, it would be very difficult and severely limiting to accommodate a bend radius that large within a high density PCB interconnect layout. Advances have been made in further reducing the in-plane bend radius of polymer waveguides. Hendrickx and Steenberge at the University of Ghent demonstrated the fabrication of in-plane micromirrors through laser ablation to create very tight in-plane bends [121]. Research by Bamiedakis in 2013 showed that the minimum in-plane bend radius could be reduced further by removing the cladding on the outer bend side, and thus temporarily providing a higher index contrast between the core and air, which served to better confine the light as it propagated along the bend [122].

In contrast, out-of-plane bend radii can be much smaller. It has been demonstrated that polymer multimode waveguides fabricated on flexible substrate can tolerate out-of-plane bend radii as tight as 1.5 mm without damage or substantial loss in reliability [93]. The cause of this strong discrepancy in bend limitations between in-plane and out-of-plane bends is based on the fact that the main scattering mechanism for high order modes is the roughness of the waveguide side walls in the direction of the bend. The side wall roughness of in-plane bends depends on the resolution achievable by the waveguide fabrication method. For instance in the case of photolithographically patterned waveguides, the roughness profile of the in-plane side walls will depend on the mask resolution, whereas in the case of laser direct imaging it will depend on the translation step resolution of the writing laser or substrate stage. However the surfaces of the top and bottom of the waveguide, which constitute the out-of-plane bend side walls, will be significantly smoother as their profiles are not directly dependent on patterning resolution, but rather on other factors such as material viscosity, layer thickness uniformity, thickness control techniques (e.g. Doctor blading) and surface tension.

The constraints on in-plane bend radius could in future be effectively mitigated by refinement of manufacturing techniques, further flexibility in crossing angles and in particular adoption of reliable in-plane mirrored bends.

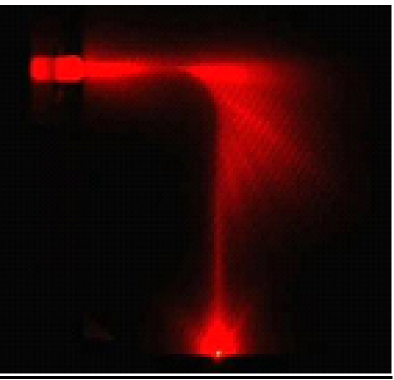

**Figure 2-13: Photo of polymer waveguide 90° arc bend illuminated with 650 nm visible light**

### 2.5.3 In-plane vs out-of-plane optical coupling to OPCBs

There has been investigations into out-of-plane coupling methods whereby integrated or discrete deflection elements were assembled onto the embedded waveguides in order to allow light to be launched and extracted normal to the plane of the waveguides [123]–[125]. However the coupling losses incurred by the deflection element, including Fresnel and scattering losses at the multiple interfaces are inherently higher than would be achieved by directly coupling to the waveguide interface collinear to the waveguide axis i.e. an in-plane coupling scheme . As will be detailed in Chapter 3, an in-plane coupling scheme between peripheral devices and the embedded waveguides was pursued in the research reported in this thesis, due to the respective orientations of the connector elements and the daughtercards on which they were mounted and in order to minimise the number of coupling interfaces.

## *2.6 Summary*

This chapter has provided a survey of commercial technologies, including midboard transceivers and optical connectors, which have in recent years brought system embedded optical interconnect closer to reality. It was described how integrated optical waveguides could in future mitigate or solve many of the mounting technical and performance challenges associated with propagation of high frequency electronic signals along traditional copper PCB transmission lines. OPCB technology is however not yet mature enough for commercial deployment, and the absence of a viable method of board-to-board connectivity is considered one of the final barriers to adoption.

The next chapter will detail the first efforts to tackle this challenge with the design and development of a suite of novel pluggable active OPCB connectors, which incorporate a parallel optical transceiver along with a mechanism to enable direct board-to-board pluggable connectivity to an OPCB backplane. These efforts included the invention and deployment of a method to assemble optical components onto an OPCB passively and with sufficient precision to ensure low loss coupling to the waveguide interface.

# 3 FIRST GENERATION PLUGGABLE ACTIVE OPTICAL CIRCUIT BOARD CONNECTOR FOR POLYMER WAVEGUIDE BASED OPTICAL CIRCUIT BOARDS

## *3.1 Introduction*

In order for OPCB technology to become commercially viable, it is crucial that the embedded waveguides can be terminated and that connector technologies are developed that allow either other PCBs or external optical cables to connect directly to the optical waveguides in the OPCB.

This will require both novel optical connector technologies and the ability to reliably and accurately align and assemble components onto OPCBs with respect to the optical waveguides therein. In the past such assembly could only be achieved through active alignment, however this would be unsuitable for high volume OPCB assembly due to the impact on equipment cost, assembly time and board yield.

### 3.1.1 Objective

- Develop a novel method of connecting peripheral devices orthogonally to a polymer waveguide OPCB, which precludes the need for embedded mirrors
- Develop active transceiver based connector system
- Develop novel, low-cost method of assembling components (optical or mechanical) to an OPCB to allow suitably accurate alignment to multimode optical waveguides embedded in the PCB substrate
- Develop bespoke mechanical coupling elements to form part of the waveguide connector receptacle
- Develop test platforms to characterise these novel technologies

### 3.1.2 Methodology

An in-plane optical connector interface concept will be designed and a number of iterative prototypes developed to prove the concept and evaluate its suitability for commercial deployment. The design will incorporate a high density parallel optical interface and would need to accommodate high speed serial modulation rates of over 10 Gb/s per channel. To this end a parallel optical transceiver circuit incorporating such an interface will also be developed to form part of an active connector mechanism. This mechanism will reside on the edge of pluggable peripheral devices to allow those devices to be optically plugged and unplugged to and from an optical printed circuit board. A suitable programming interface will be developed to allow user configuration of critical transceiver control parameters such as channel enable, laser modulation current, laser bias current and receiver squelch.

Peripheral test cards will also be designed to accommodate the connector prototypes. These test cards will serve as a conduit for external serial test data to the transceiver channels in the connector. Finally a complete proof of concept demonstration platform will be constructed, comprising a test chassis, single board computer, test cards,

prototype connectors and an OPCB to allow comprehensive optical and mechanical characterisation of the connectors. The test chassis will be 19” wide in order to be compliant with standard data centre racks. This will enable the chassis to be mounted into a data centre rack at a later stage with the option to be connected to fully functional storage and server enclosures in a real data centre system.

A crucial part of this OPCB connection system will be the optical waveguide receptacle for the pluggable active in-plane connector, which is fixed to the board and meets the tolerance requirements to enable the optical connector interface to align with the embedded waveguide interface.

To this end, a novel low-cost method will be invented and developed of assembling components (optical or mechanical) to an OPCB such as to allow suitably accurate alignment to multimode optical waveguides embedded in the PCB substrate. A suite of receptacles will be designed and developed, which incorporates features to allow passive accurate alignment of the receptacle onto the board waveguides. In addition, compliant features will be designed in the board itself to accommodate the accurate mechanical registration of the receptacle.

The final iteration of the prototype will allow a lens array to be accurately fixed to the OPCB and form part of a dual lens expanded beam coupling solution.

## *3.2 OPCB connectivity in data storage enclosures*

### 3.2.1 Electro-optical midplane architecture

It is proposed that the projected performance bottleneck in data storage systems is mitigated by incorporating electro-optical printed circuit board (OPCB) and interconnect technology on the midplane [71]–[76]. As shown in Figure 3-1 the midplane would thus comprise a PCB with both copper layers for electrical power and low speed electronic signal distribution and one or more optical transmission layers to allow high speed signals to be conveyed optically. As discussed in the previous chapter there are various types of optical waveguide that can form the basis of such optical transmission layers such as laminated optical fibre webbing, planar glass optical waveguides or planar polymer optical waveguides. The efforts described in this thesis focus on the deployment of planar polymer waveguide technology in electro-optical PCBs and the development of commercially viable supporting connection technologies.

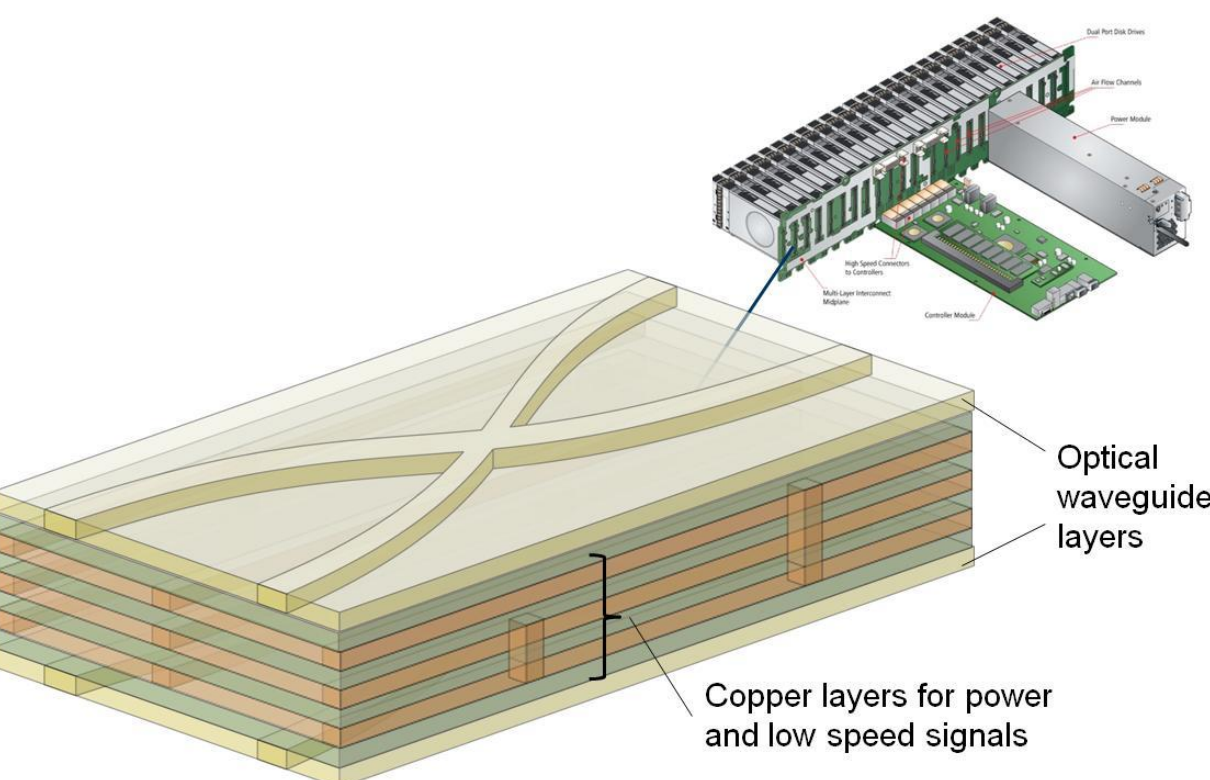


**Figure 3-1: Schematic view of electro-optical midplane architecture comprising copper layers to accommodate power, static and low speed signal distribution and at least one optical transmission layer to convey high speed signals along optical waveguides**

### 3.2.2 Midplane interconnect requirements

A conventional data storage array enclosure is shown in Figure 3-2. The enclosure comprises 24 hard disk drives, two controller modules (only one of which is shown populated) and two power supply modules (only one of which is shown). The disk drives are located on the "front end" and plug into a multilayer midplane PCB from one direction, while the controller modules and power supply modules plug into the midplane from the other direction. The midplane serves as the interconnect backbone of the system accommodating power distribution, static control lines, low speed control busses and high speed transmission lines. The midplane will also typically require slots to allow air flow through the system, which further reduces the amount of functional area that can be used for interconnect (traces) or components.

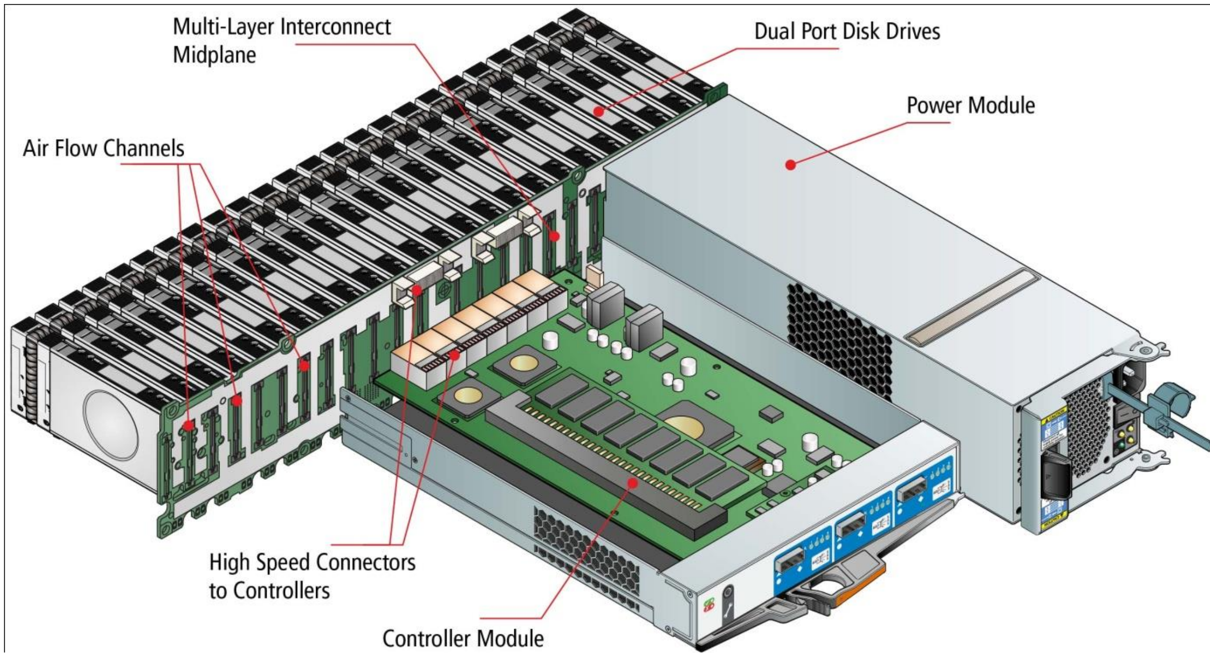


**Figure 3-2: Data storage array with 24 drives, two controller modules (one shown) and two power supply modules (one shown)**

In a conventional data storage, integrated or computer system, the midplane / backplane and its peripheral line cards (such as disk drives, controller modules, server blades) are arranged in a mutually orthogonal configuration. The peripheral line cards must be pluggable to the midplane i.e. they can be manually connected to or extracted from the midplane. In most cases there is a requirement that peripheral devices are also "hot pluggable", which means they can be connected or disconnected from a system during operation without

compromising the system. The midplane in data storage enclosures must be passive, that is to say there can be no active devices on the high speed signal paths on the midplane. The reason for this is that as the midplane is not a field replaceable unit and, the critical failure modes on the midplane should be minimised by ensuring that all high speed data links are passive. This would therefore preclude the use of signal conditioning circuitry such a pre-emphasis or decision feedback equalisers on the midplane, and more pertinently it also prevents the deployment of optical transceiver circuitry on the midplane.

### 3.2.3 Electro-optical PCB connection scheme

As the midplane is passive in the proposed application model, then this would require that the optical transmitter and receiver devices are situated on the peripheral line cards (such as controller modules and disk drives), and that the line cards are both electrically and optically pluggable to the electro-optical midplane (Figure 3-3).

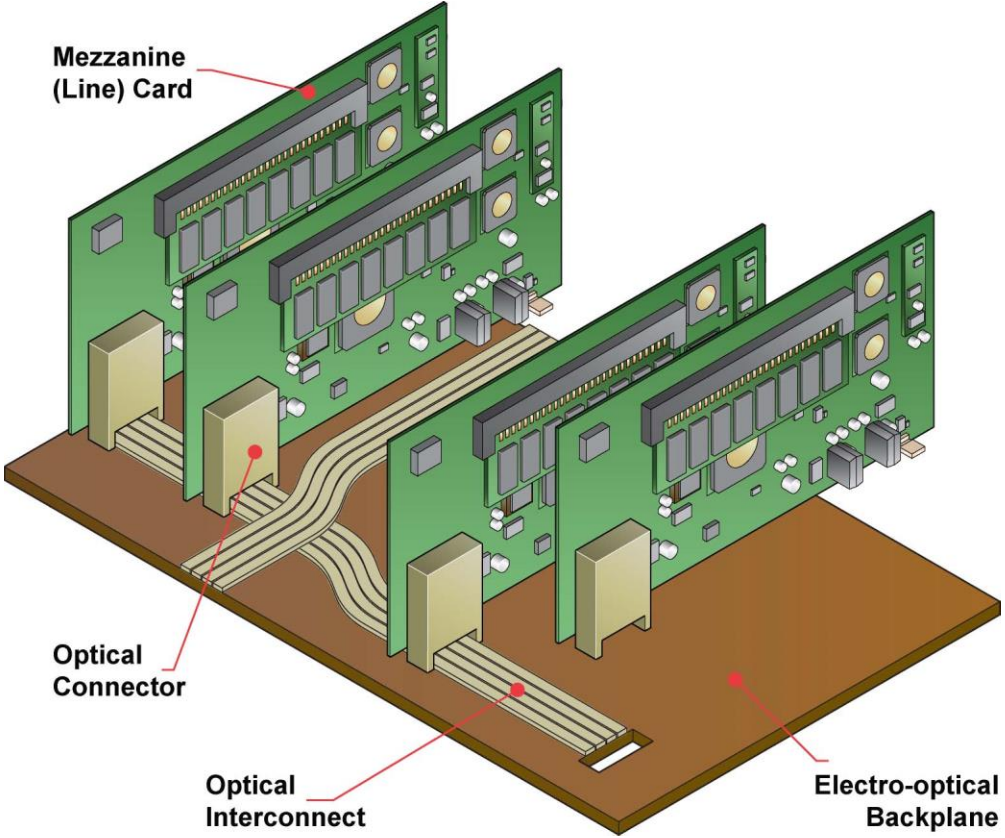


**Figure 3-3: Schematic view of line cards with optical edge connectors plugging in orthogonally to an electro-optical backplane with embedded optical interconnect**

In a cost-effective system, the optical transmitter and receiver devices would most likely comprise VCSELs and

PIN photodiodes.

In the most elementary configuration, whereby the VCSELs and photodiodes are mounted directly onto the line card, the optical axes of their respective emitting and receiving areas are orthogonal to the plane of the line cards and would lie parallel to the plane of the electro-optical midplane to which the line cards would connect.

Thus an optical channel can be conveyed by an optical waveguide in the midplane if the optical transmitter and receiver interface on the line card are drawn into direct physical contact with the waveguide interface on the midplane (Figure 3-4). This would be satisfied by a butt-coupling engagement scheme, which would eliminate the need for intermediary optical interfaces on the backplane, such as 45° optical deflection structures, and thus minimise the number of interfaces incurring optical losses as well as additional assembly costs.

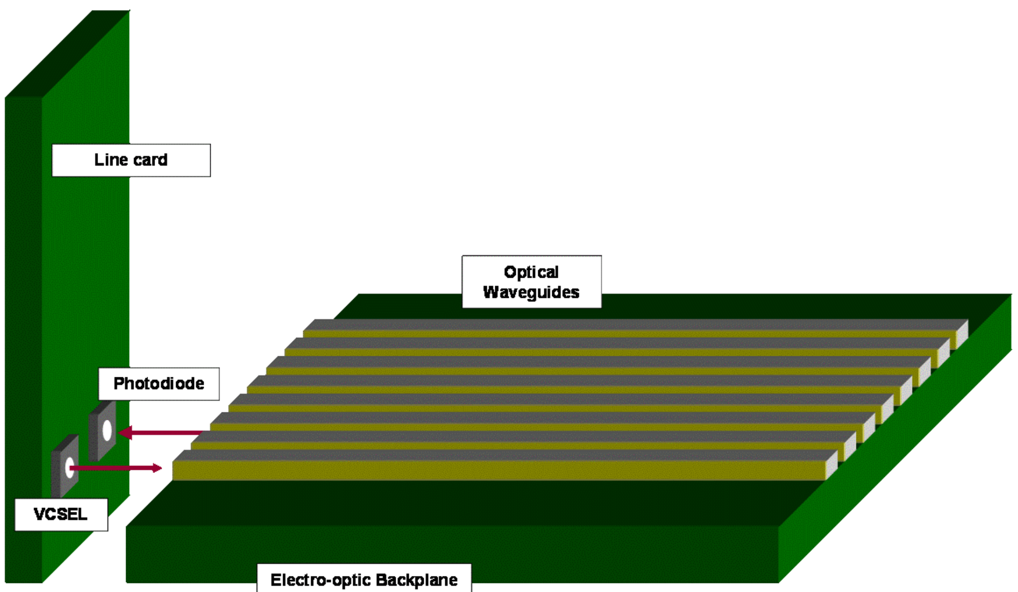


**Figure 3-4: Butt-coupled optical connection between surface emitting and receiving devices on the line card and planar optical waveguide interface on electro-optic backplane**

### 3.2.4 Electro-optical PCB connection concept

Research into lateral misalignment tolerances [126] of a VCSEL to a multimode planar waveguide have shown that direct butt-coupling of an 850nm VCSEL to multimode polymer waveguides would be subject to lateral misalignment tolerances of approximately 0.9 of the waveguide core size, the maximum lateral misalignment being defined as the displacement from the optimum alignment position, at which the transmissivity through the waveguide drops by 1 dB from the maximum transmissivity. Given the dimensions and layout densities inherent in conventional parallel multimode waveguides, the interface between the line card and optical midplane would

be subject to very stringent alignment tolerances (±22.5 µm for a 50 µm core width).

These tolerances would have to be met continuously throughout the connection cycle and would therefore require that the physical optical interface remain immune to typical movements arising between the line card and midplane, which will arise in a typical data centre enclosure such as system enclosure vibrations, air-flow and PCB deformation. The optically connected interface would therefore require a level of detachment from other non-optical midplane connecting components, such as mechanical guide-rails and electrical (power/low speed signal) connectors, as these are generally subject to far lower assembly tolerances than could be accommodated optically. The cost of widespread precision assembly on all midplane interface components would be prohibitively high.

In order to address these challenges, the author invented and patented a pluggable optical PCB (OPCB) connector concept [127] whereby the optical transceiver interface is housed on a platform, which is detached from the rigidity of the supporting line card by a mechanically flexible bridge. The optical interface platform contains mechanical features which will engage with compliant structures on the optical midplane allowing the optical interface to be drawn into precise physical alignment with the waveguide interface on the OPCB (be it a midplane or backplane). Finally, a mechanical engagement mechanism is required to raise and lower the optical interface platform into compliancy.

The method comprised a two stage engagement process: a first stage of coarse alignment whereby the daughtercard is inserted into the midplane providing the necessary electrical and mechanical connections, and a second stage of higher precision alignment whereby the optical connection is asserted. Thus the requirement for precision component assembly is restricted to the line card optical interface and the compliant receptacle on the OPCB midplane / backplane.

## 3.3 Electro-optical PCB demonstration platforms

In order to evaluate the viability of this technology concept in a high speed data communication system environment, the author designed and developed two generations of active connector and corresponding demonstration platforms:

The Storlite demonstration platform (Figure 3-5a) comprised a first generation active pluggable optical backplane connector, two test line cards and separate passive optical and electrical PCBs [128]. The Storlite platform will be described in this chapter.

The FirstLight demonstration platform (Figure 3-5b) comprised a second generation active pluggable optical backplane connector, four test line cards and an electro-optical backplane [89], [129].

The FirstLight technologies will be described in Chapter 4.

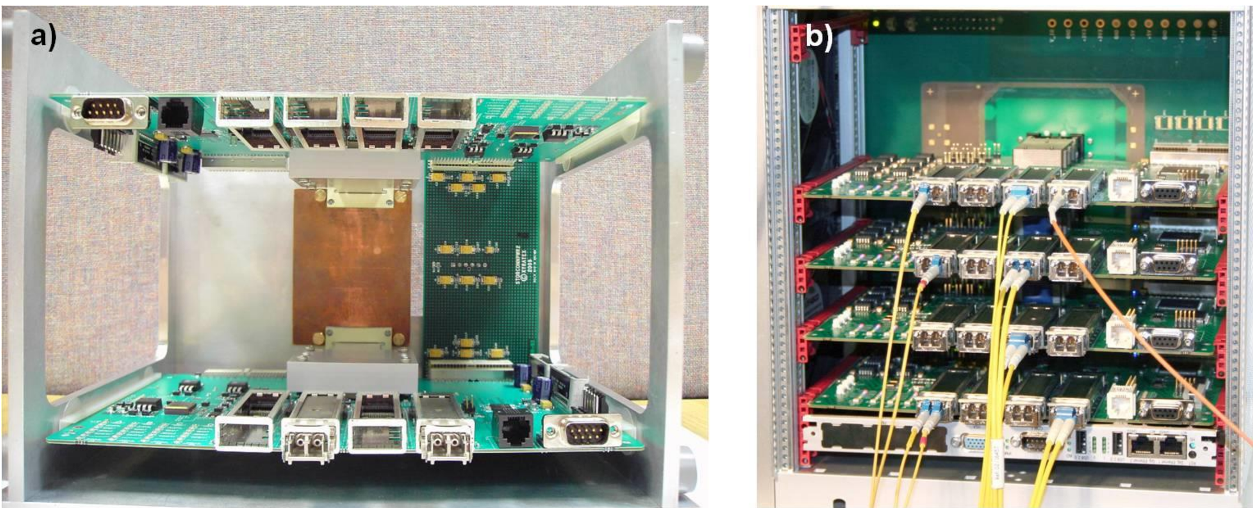


**Figure 3-5: a) Storlite Demonstrator - First generation optical backplane connector platform, b) FirstLight Demonstrator - second generation optical backplane connector platform**

## *3.4 Storlite – first generation passive optical backplane system with active pluggable connectors*

### 3.4.1 Storlite project summary

During the Storlite project, a first generation optical backplane connection system was developed, which allowed for repeatable docking and undocking of an active optical interface housed on a line card to the optical interface of a passive optical backplane containing planar multimode polymer waveguides. The connector comprised a parallel optical transceiver circuit, a self-aligning optical interface and a connector mechanism. The transceiver was capable of supporting data rates of 10.3 Gb/s on each of four duplex channels. The transceiver circuit was constructed on a flexible material to enable the optical interface to mechanically float with respect to the line card, thus allowing the critical optical connection to remain relatively immune to displacements between line card and backplane. A manual connection mechanism controlled the engagement and disengagement of the transceiver with the optical backplane.

Finally an optical backplane demonstration system was constructed comprising a passive optical printed circuit board, a passive electrical backplane, two line cards and the proposed pluggable optical connector (Figure 3-6). The system was successfully characterised with respect to 10.3 Gbps board-to-board test data exchange.

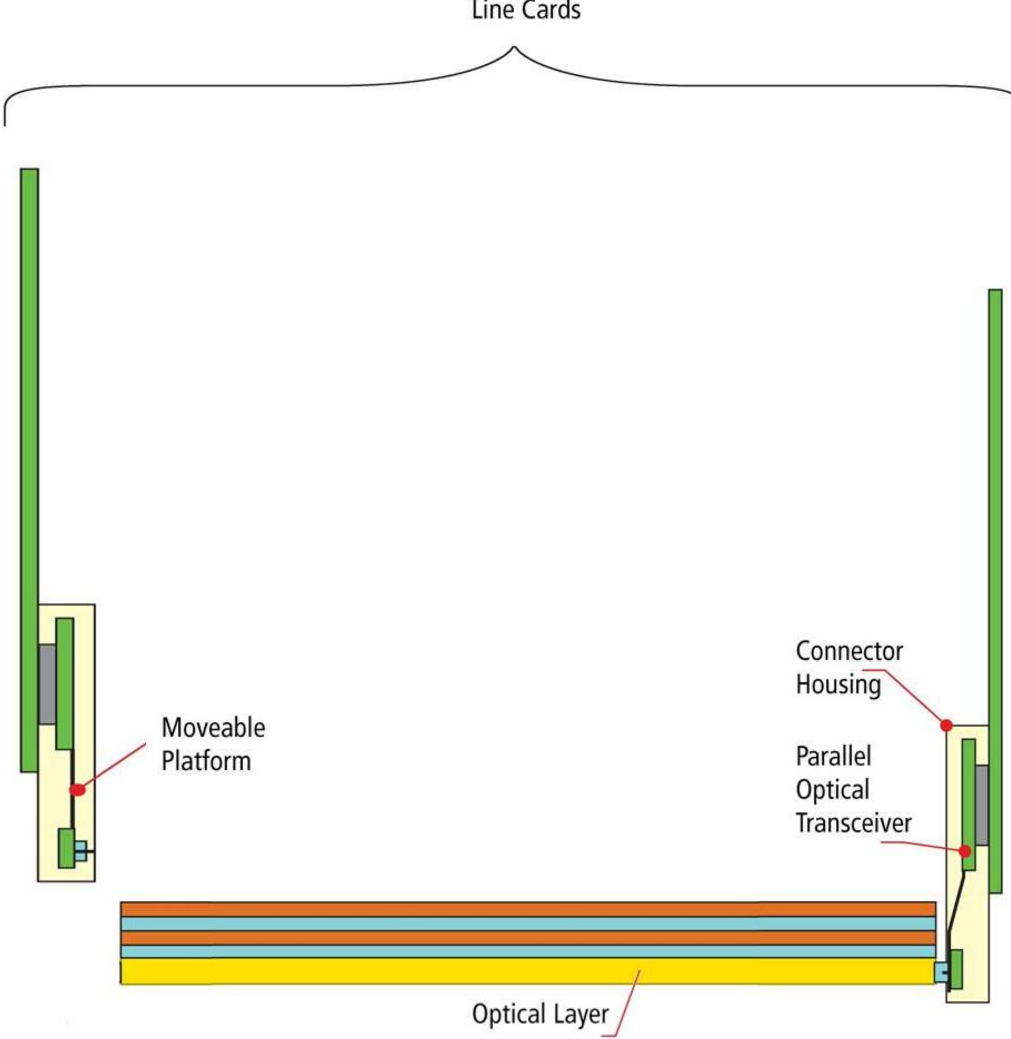


**Figure 3-6: Storlite optical backplane connection scheme**

### 3.4.2 System design overview

The author carried out the design, development and characterisation of the active pluggable optical connector and of all circuit boards in the Storlite platform. Table 3-1 contains the formal designations for each circuit board type designed for the Storlite system.

**Table 3-1: Storlite card designations**

| Card designation | Description |
|---|---|
| StorConn2 | Quad 10 Gb/s parallel optical transceiver circuit |
| StorConnTest2 | 10 Gb/s test daughtercard |
| StorConnOpt2 | Passive optical backplane |
| StorConnPwr2 | Passive electrical backplane |

The Storlite platform is populated by two 10 GbE LAN test daughtercards (StorConnTest2) plugging into a passive optical backplane (StorConnOpt2) and a separate passive electrical power backplane (StorConnPwr2). Each test daughtercard supports one active pluggable optical connector, which includes a quad 10 Gb/s parallel optical transceiver circuit (StorConn2).

### 3.4.3 StorConn2 - Storlite quad 10 Gb/s parallel optical transceiver

A quad parallel optical transceiver circuit was designed to accommodate the proprietary connection technique. The Storlite optical transceiver circuit was constructed on flexible laminate PCB; the sections housing the optical interface and the base section with the electronic array line-card connector were supplemented with rigid FR4 layers, while the bridge section between these sections was exposed and flexible. This arrangement served the critical purpose of ensuring that the rigid platform supporting the optical interface was free-floating with respect to the line-card, thus allowing the optical interface to be manipulated into precise alignment with the polymer optical waveguides embedded in the passive optical PCB irrespective of the relatively coarse alignment of the line-card.

Figure 3-7 shows an oblique view of the quad transceiver circuit mounted on a partially flexible and rigid substrate. The transceiver circuit comprised three sections: a base section, a flexible bridge section and a moveable optical interface. The circuit was constructed on a flexible laminate substrate, which was reinforced with rigid FR4 layers in the base section and optical platform leaving the intermediary bridge section flexible. The base section allowed for the electrical connection of the transceiver to the peripheral line card by means of a high speed electronic array connector (Mezzanine GIG Array®) capable of supporting data rates up to 11 Gb/s.

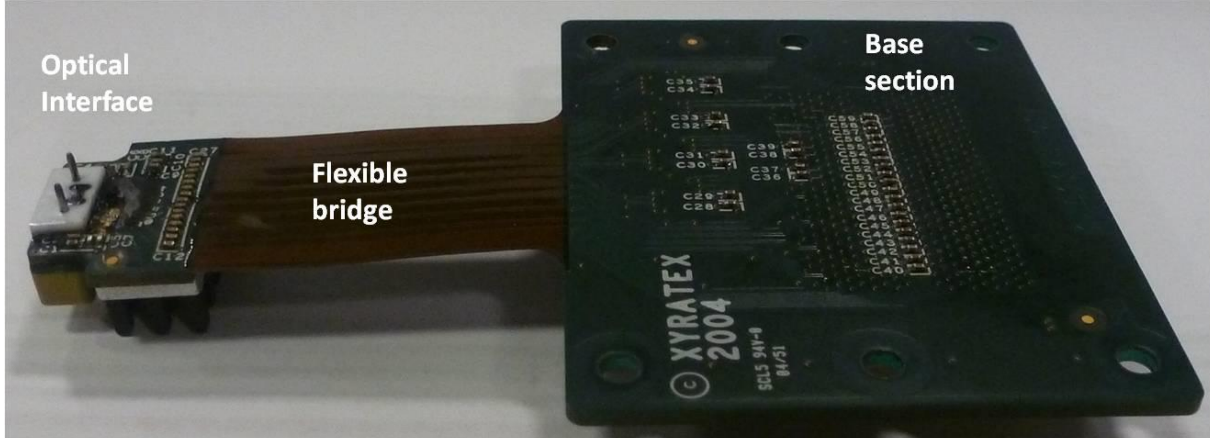


**Figure 3-7: Storlite parallel optical transceiver circuit**

Figure 3-8 shows an annotated side view of the Storlite transceiver circuit and its key components. A polymer resin block (Figure 3-8a), which contained two MT pins of diameter 0.7 mm and with a centre to centre separation of 4.6 mm was attached to the bottom of the optical interface section, with the two pins protruding up through compliant holes in the PCB. These pins form part of the opto-mechanical interface (Figure 3-8b), which comprises the lasers and receivers, GRIN micro-lens array and male MT compliant plug. A heatsink (Figure 3-8c) is provided below the power consumptive quad VCSEL driver and quad transimpedance amplifier / limiting amplifier (TIA/LA) array ICs. The ICs are mounted on the top side of the PCB as shown and covered in epoxy. Each IC has a thermal dissipation pad on its base. In order to maximise heat transfer from the ICs to the heatsink on the bottom of the PCB, two holes were milled in the PCB under each IC's thermal pad section. Two copper slugs were then carefully assembled into the holes in order to create a thermal bridge between the ICs and the heatsink. An intermediary flexible section (Figure 3-8d), in which the board was not reinforced with FR4 layers, was required to provide a level of mechanical detachment between the optical interface section and the base section. The base section in turn was rigidly connected to the StorConnTest2 test daughtercard by means of the GigArray® high speed electronic array connector offered by FCI (Figure 3-8e).

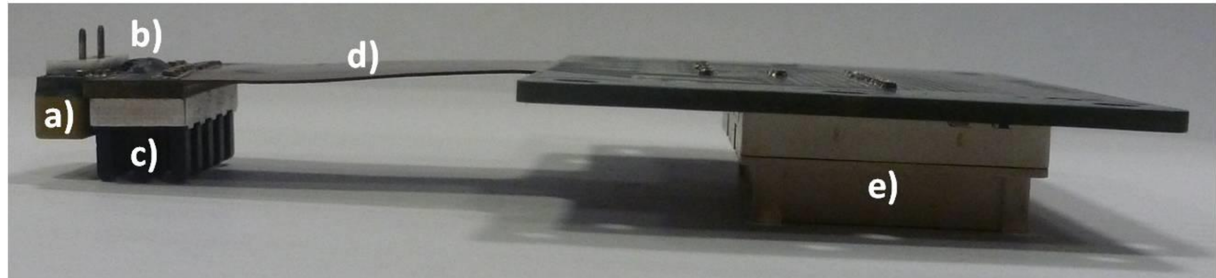


**Figure 3-8: Storlite parallel optical transceiver side-view: a) MT pin holder, b) Opto-mechanical interface, c) Heat sink, d) Flexible bridge, e) GigArray® high speed connector**

#### *3.4.3.1. Transceiver circuit functional description*

Figure 3-9 shows a functional diagram of the transceiver circuit. The Base section of the transceiver contains a high speed electronic array connector, which conveys electronic signals to and from the transceiver circuit from and to the StorConnTest2 test daughtercard PCB on which the transceiver is mounted. The differential signals are conveyed from the Base section across a flexi-rigid bridge to the electro-optical conversion circuitry on the Optical Interface section. The Optical interface comprises a four channel transmitter array of four Vertical Cavity Surface Emitting Lasers (VCSELs) operating at a wavelength of 850 nm, and a four channel receiver array of four positive intrinsic negative (PIN) photodiodes sensitive to a wavelength of 850 nm. There are two reasons for the choice of 850 nm sources: 1) Commercial VCSEL arrays for data communication were only available at this wavelength. 2) The optical polymer material to be used in the OPCB is substantially more transmissive at this wavelength than at the longer wavelengths associated with laser sources used for telecommunications. The PIN photodiodes are comprised of GaAs, which will allow them to generate a current in response to a wide range of wavelengths. An important metric for photodiodes is the Responsivity, which is the ratio of electrical current generated to the optical power received. GaAs PIN photodiodes have a suitably high responsivity at 850 nm. A four channel VCSEL driver chip converts differential electronic signals from each of four high speed differential transmission lines into a varying current across the VCSEL diode, which in response generates corresponding optical signal pulses. On the receiver side, the PIN photodiodes are held under reverse bias so they do not conduct electricity, save for a negligible dark or leakage current. When a photon of sufficient energy arrives at and is absorbed in the depletion region, it will form an electron-hole pair, which is swept away by the reverse bias field, thus generating a current which is proportional to the number of photons received. Thus varying optical signal pulses received on any of the four photodiodes are converted to an electric current, which is then passed to a trans-impedance amplifier (TIA), the purpose of which is to convert current

back into an electronic differential signaling voltage of proportional amplitude. Each of the four differential signal lanes is then conveyed across the flexible bridge to a 10.7 Gb/s adaptive receive equaliser designed to reduce signal jitter on high speed signals, by reducing inter-symbol interference and adapting the signal to frequency dependent skin effect and dielectric losses incurred over the transmission lines from the TIA. The differential signal output from equaliser is then conveyed to the high speed electronic array connector to the StorConnTest2 test daughtercard. A detailed introduction to the operational principles of optical transmitters and receivers is provided by DeCusatis [130].

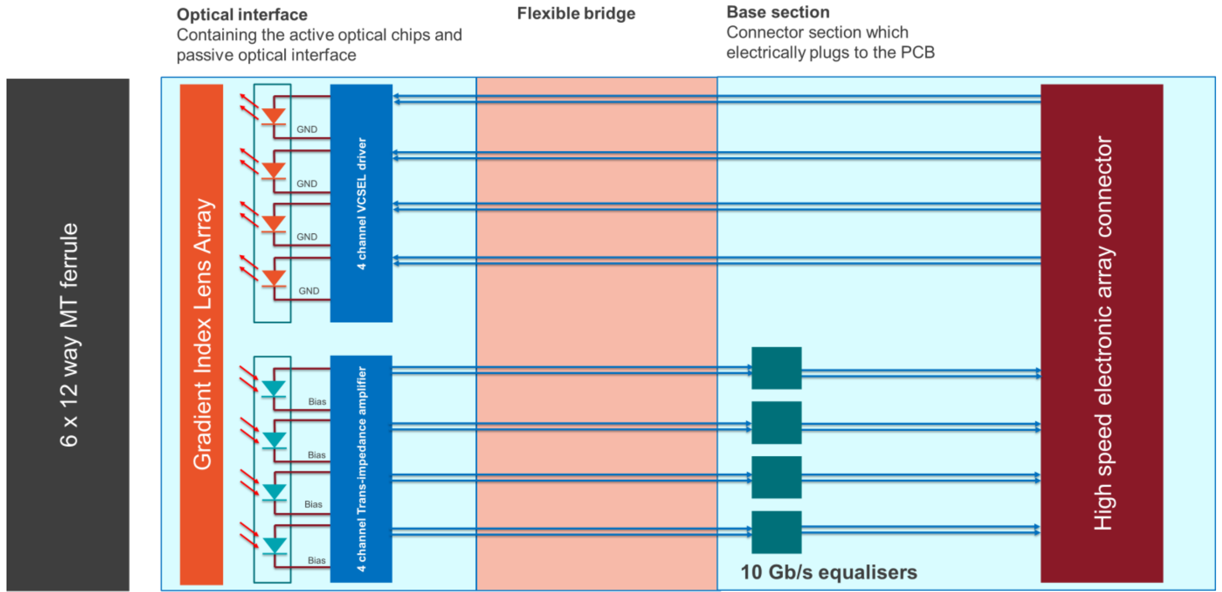


**Figure 3-9: Functional diagram of quad transceiver circuit**

### *3.4.3.2. StorConn2 circuit hardware design*

The principal feature of the *StorConn2* circuit design allows for both a robust electronic connection to the host board and a quasi-free floating optical connection to the optical backplane. The reason for the latter lies in the fundamental requirement that the host board optical connector be directly connectable to the optical backplane precluding the need for intermediary optical transmission media such as external fibre ribbons. In view of the fragility of the optical backplane paraphernalia and the criticality of the optical connection, the quasi free-floating feature coupled with the self-aligning mechanism of the connector allows for greater optical alignment tolerance and mechanical sensitivity. The interface should be able to tolerate relative lateral translation between the daughtercard and backplane of ±2 mm, which would account for the most severe operational vibrations that the system may undergo, though, in practice, given the firm mechanical connection of the daughtercard to the

backplane through electrical connectors and guide rails, it is unlikely that translations of that magnitude would occur.

An electro-optical connector with a floating mechanism was later developed by Fujitsu to address the same issue of decoupling precise optical alignment tolerances to greater electrical alignment tolerances [52].

### *3.4.3.3. Raw card outline*

The StorConn2 raw card consists of a Kapton Polyimide based flexible PCB sandwiched between two rigid extremities of the card. In both the Base and Optical head sections, the flexible material is reinforced by FR4. The Bridge section is mechanically flexible.

Figure 3-10 shows the physical shape of the card and the principal dimensions of its outline.

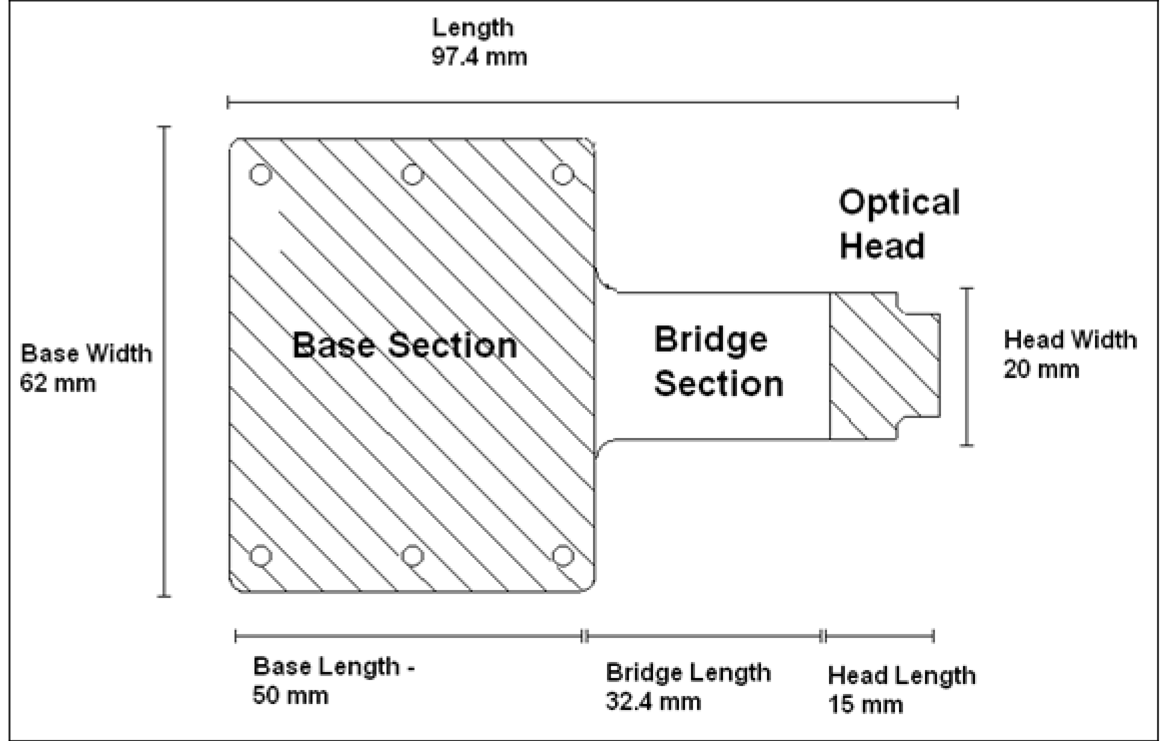


**Figure 3-10: StorConn2 raw card outline dimensions**

### *3.4.3.4. PCB layer stack-up*

Figure 3-11 outlines the complete layer stack-up of this circuit board.

StorConn2 has five electrical layers, of which two are signal layers (layers one and five), two are ground layers (layers two and four) and one is a power layer (layer three).

The high speed differential tracks are located on the top layer (layer one) and bottom layer (layer five) and designed with a controlled differential impedance of 100 Ω ± 8%. As these differential tracks are on the surface layers and referenced to an adjacent ground layer, they are referred to as microstrip type. Differential traces buried between layers are known as stripline and result in reduced EMI compared to microstrip as all electric field lines are coupled to planes on surrounding layers and thus confined within the PCB. Stripline design, however also requires more PCB layers and more vias than microstrip, which leads to a greater probability of creating impedance mismatch points (e.g. vias) and greater difficulty in achieving the required differential impedance of 100Ω [131].

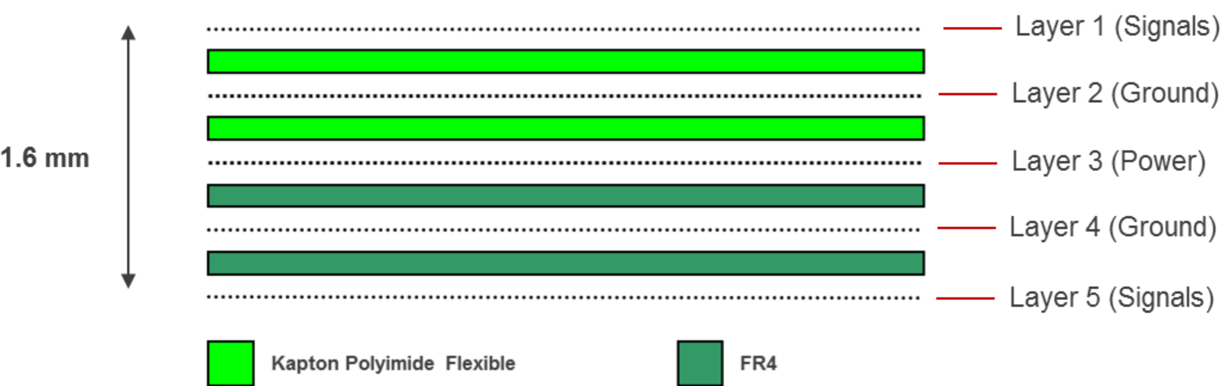


**Figure 3-11: StorConn2 PCB layer stack-up**

### *3.4.3.5. Opto-mechanical interface*

The optical interface comprised a section of the parallel optical transceiver, which established the direct physical connection with the optical waveguides embedded in a passive optical backplane. The active transmit and receive elements – VCSEL die array and PIN photodiode die array (Figure 3-12a) – were bonded directly to the transceiver PCB over appropriate thermal distribution structures.

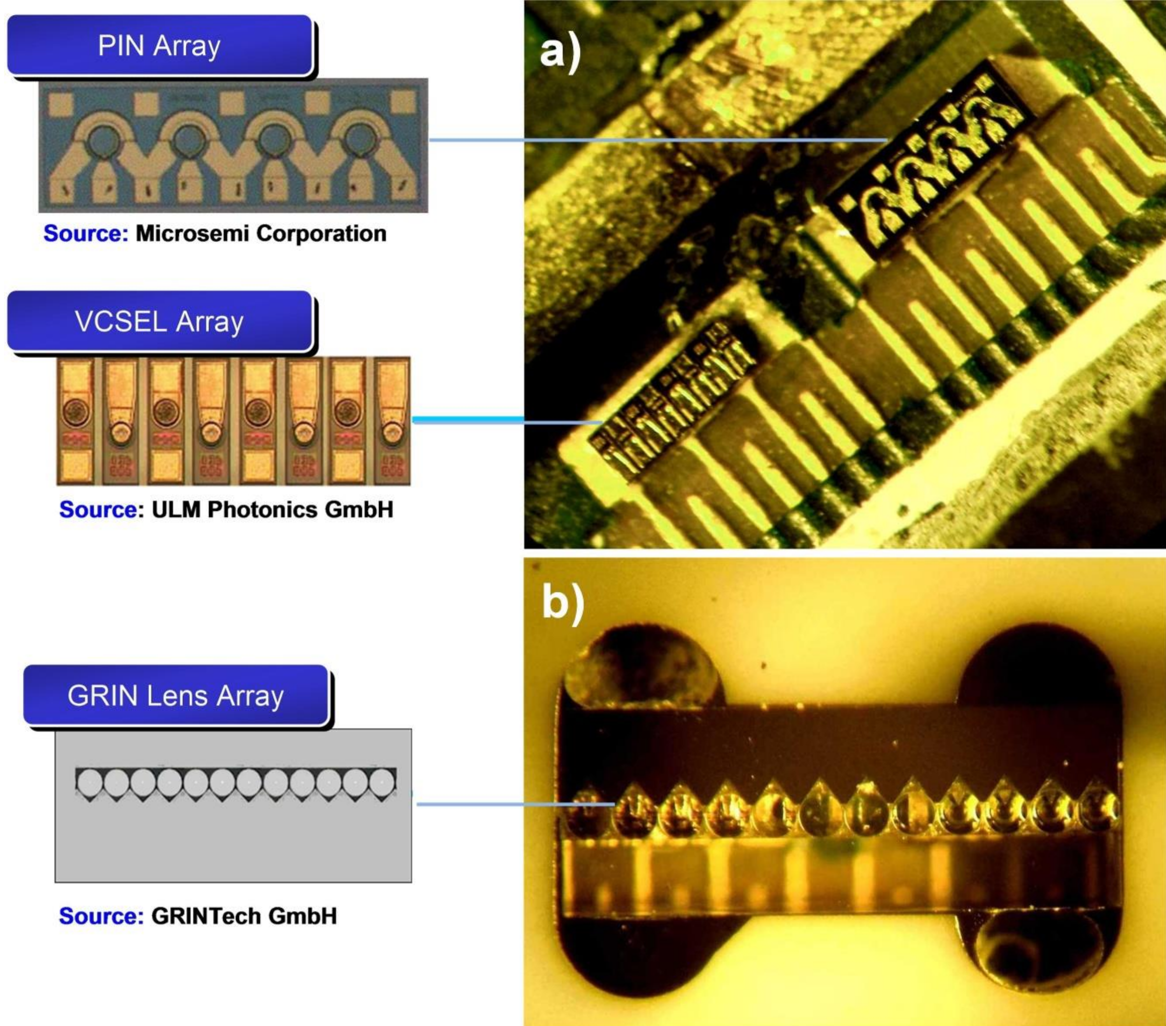


**Figure 3-12: Optical interface a) Quad VCSEL and PIN photodiode die array aligned and assembled onto PCB, b) GRIN micro-lens array in ceramic holder aligned over active areas of VCSEL and PIN photodiodes**

VCSELs were chosen as a suitable laser transmitter on the basis of cost, suitability of emission orientation and wavelength compatibility with optical polymer transmission [72], [73], [76], [79], [90], [132]–[139]. A quad die array VCSEL from ULM Photonics, ULM850-10-TT-C0104U, was chosen with a package die size of 1000 µm x 350 µm x 150 µm, nominal emission wavelength of 850 nm and maximum data rate of 11 Gb/s. Each VCSEL had a circular active area of 7 µm diameter and an emission divergence angle of 25°.

A coplanar PIN photodiode quad die array from Microsemi, LX3045, was chosen with a package size of 1200 µm x 450 µm x 203 µm, nominal receive wavelength of 850 nm and active circular receive aperture size of 75 µm.

The high speed electronic connections were made between the VCSEL and PIN photodiode die and the PCB

through wire bonds between copper lands on the PCB and contacts on the top surface of the die.

After the VCSEL and PD die were attached and bonded, a GRIN micro-lens array (Figure 3-12b) in ceramic holder was aligned over active areas of VCSEL and PIN photodiodes using the MT pins protruding from the optical interface as a reference.

#### *3.4.3.6. GRIN micro-lens array and MT compliant support structure*

In selecting an appropriate surface to mate physically with the waveguide interface, the author initially considered a geometric micro-lens array; however, following discussions with vendors, the author decided that that the ability to collimate the divergent output of a VCSEL into a waveguide at close proximity carried too high a risk of exceeding available manufacturing tolerances and introducing undesirable Fresnel losses and spherical aberrations.

The author then identified that a gradient index micro-lens array (GRIN lens) would be suitable as it could be designed to image the circular optical output from the VCSEL onto the point of intersection with the waveguide by determining, through simulation, the optimum GRIN lens length. Likewise the lens could image the output of the waveguide onto the active area of the PIN photodiode. GRIN lens technology has been deployed as a suitable coupler to micro-waveguide structures on silicon-on-insulator chips [140], but, to the knowledge of the author, has not been used in this way before to couple directly to macroscopic optical waveguides.

The GRIN lens array and support structure could serve both to protect the fragile die underneath and to counter the optical divergence from the VCSELs and the waveguides. In addition, the planar ingress / egress surfaces of the GRIN lenses made them suitable for butt-coupling to flat surfaces.

The author designed a custom GRIN lens array and support structure, which included two MT compliant slots and a GRIN lens array supported at an offset from the datum of the two MT pins (Figure 3-13). The purpose for this offset was to accommodate the need for compliancy with an MT compliant optical waveguide receptacle (described in section 3.4.6), which sits over the waveguide interface and on which, therefore, the datum between the two compliant MT pins cannot be in line with the waveguide array on the optical PCB. The offset was chosen to be compliant with that of a 6x12 way MT ferrule (Figure 3-14), which includes six rows of 12 fibres distributed evenly around the datum of the two MT pins. The offset between the GRIN micro-lens array and the datum between the two MT slots in the GRIN lens holder was chosen to match the offset between the first or sixth row in the 6x12 MT Ferrule and the datum of the two MT slots. This way a commercial fibre-optic patch-

cord terminated in a 6x12 MT ferrule could be used for stand-alone testing of the transceiver and the optical PCB waveguides.

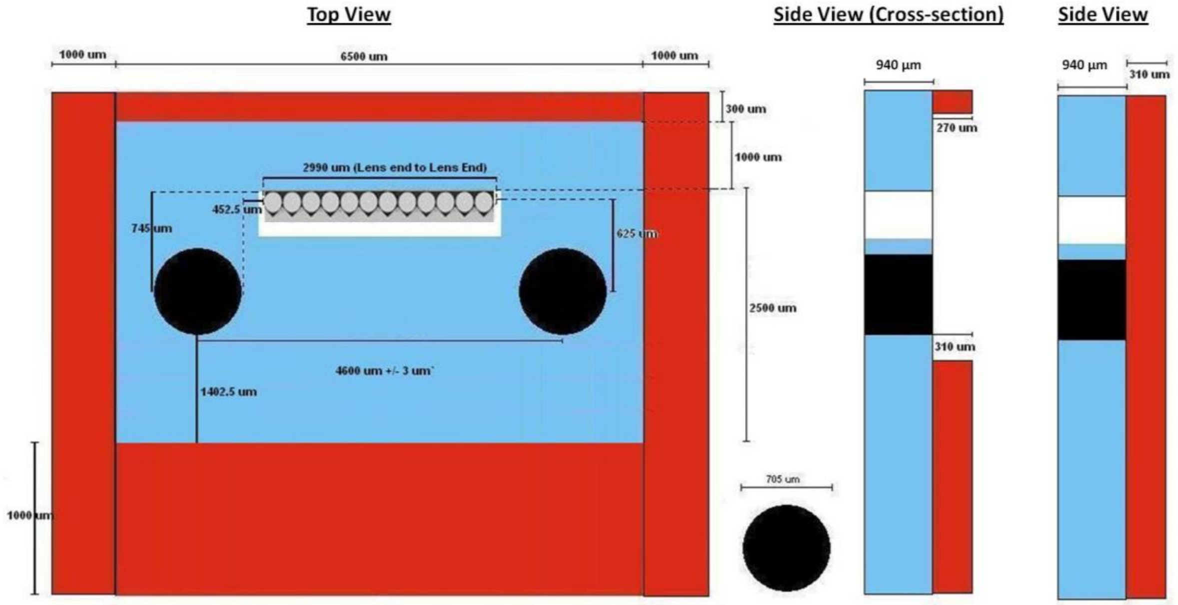


**Figure 3-13: Physical drawings of Gradient Index Micro-lens array and custom ceramic holder**

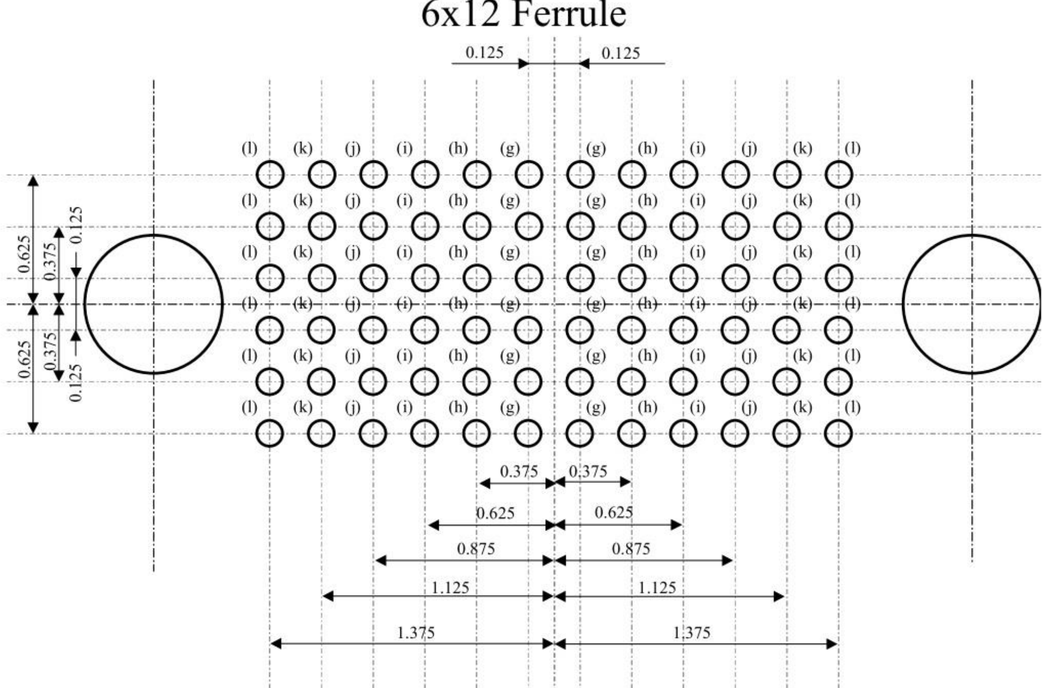


**Figure 3-14: 6x12 MT ferrule interface. The lowest (or highest) row of optical channels is offset from the datum of the MT pins by 625 µm**

The author liaised with a German GRIN lens fabricator GRINtech GmbH to determine the optimum lens parameters and support stand-offs for the optical interface.

The GRIN lenses each had a diameter of 240 µm and were held at an inter-lens pitch of 250 µm by a V-groove array. The ceramic holder incorporated two slots to accommodate the MT pins protruding through the PCB. The positional alignment of these slots to the lens array, mirror the alignment of the MT pins to the VCSEL and PIN arrays, thus providing initial self-alignment of the lens array to the associated photonics. The central four lenses were not used.

The surface plane of the ceramic holder is supported on platforms on all four sides leaving a clearance in the central area under the lenses for the VCSEL / PIN die and their respective bonding areas. The front platform is shorter than the others by 30 µm in order to accommodate the thickness of the high speed copper traces passing to the arrays.

The same GRIN lens had to be used to both image the divergent VCSEL output into the 70 µm x 70 µm waveguide aperture and image the output of the waveguide into the receive aperture of the PIN photodiode. Therefore the length of the GRIN lens axis was chosen to provide the best compromise between the imaging requirements of both the VCSEL to waveguide axis and the waveguide to PIN photodiode axis. GRINTech carried out a ZEMAX ray tracing simulation to determine the optimum GRIN lens length ZL (Figure 3-15) to satisfy both requirements. The simulation took account of the different die thicknesses of the VCSEL and the photodiode and therefore the different distances between the emitting VCSEL aperture and the planar GRIN lens ingress point and the GRIN lens egress surface and the PIN photodiode receiving aperture. The VCSEL die was 50 µm thinner than the PIN photodiode die. As both die were mounted on the same level PCB surface, this meant the receiving aperture of the PIN photodiode was 50 µm closer to the GRIN lens surface than the emitting aperture of the VCSEL. Figure 3-15 shows three positions for source and receiving element in order to highlight how the image point will change when the source is misaligned with respect to the GRIN lens. The simulation also took into account the NA of the waveguide based on the refractive indices of core and cladding and the NA of the VCSEL based on its nominal divergence of 25°. Based on the results of this simulation, an optimum GRIN lens length of 0.944 mm was chosen. It should be noted that the launch profiles from both waveguide and VCSEL were assumed in the simulation to have a Gaussian profile. The output of a waveguide could be assumed to be Gaussian, if the waveguide geometry (length, bends) is such that it will promote sufficient mode mixing. However it should be noted that VCSEL source profiles are not Gaussian in nature, as they will typically exhibit a characteristic dip in the centre of the profile, therefore this simulation could be further optimized in future.

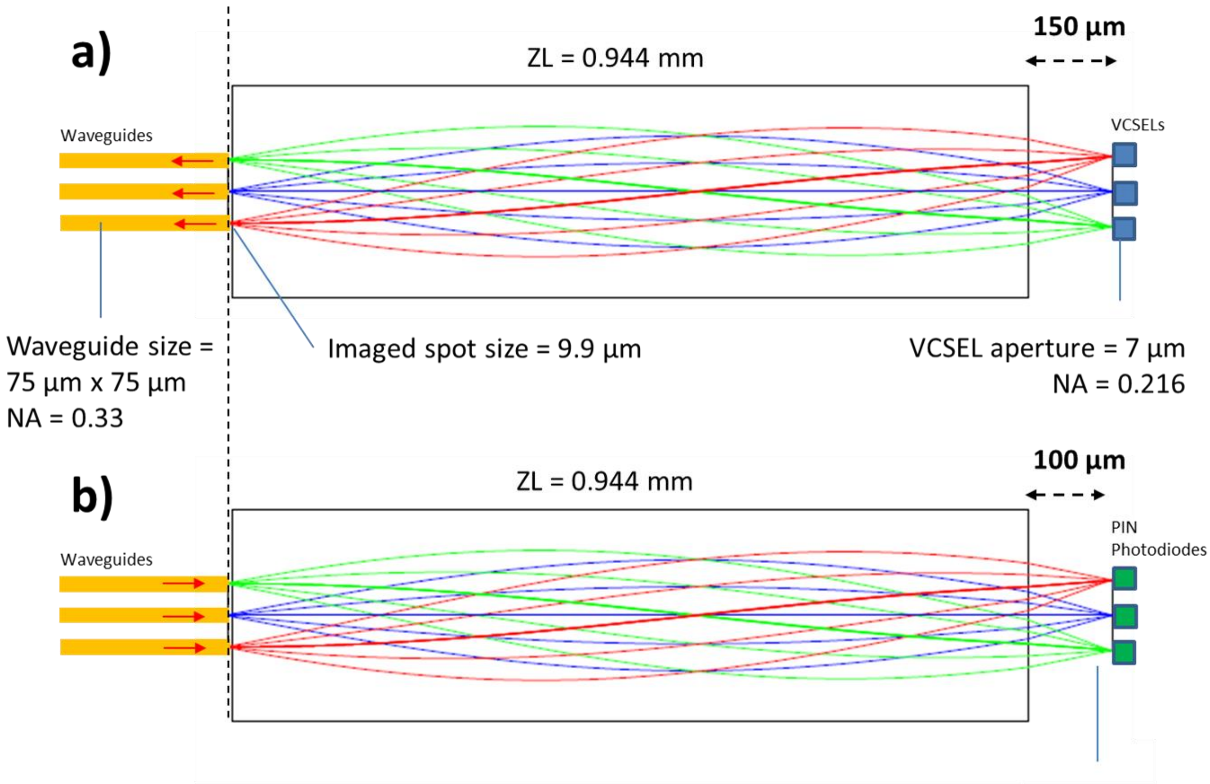


**Figure 3-15: ZEMAX simulation provided by GRINTech of optimum GRIN lens length (ZL) required to enable: a) imaging of the 75 µm x 75 µm waveguide output into the 75 µm circular aperture of the PIN photodiode and b) imaging of the VCSEL output into the waveguide**

The GRIN lens design was sent to GRINtech GmbH who fabricated the custom ceramic assembly incorporating a 1 x 12 GRIN lens with part number **GT-LFRA-025-12-048-NC**.

*3.4.3.7. StorConn2 component layout*

Table 3-2 lists the main components of the StorConn2 card, while Figure 3-16 shows their positions. The power dissipation of all active components is shown except the PDs in order to extrapolate the power dissipation and efficiency of the complete module.

**Table 3-2: Component layout and description**

| Item | Supplier | Part No. | Description |
|---|---|---|---|
| **Optical Head** | | | |
| VCSEL Driver | Primarion (Acquired by Infineon Technologies in April 2008) | PX6514 (obsolete) | 10 Gb/s data rate per channel (40 Gb/s aggregate data rate)<br>Power dissipation = 570 mW (all on) |
| TIA / LA | | PX6524 (obsolete) | 10 Gb/s optical receive system, Transimpedance amplifier & Limiting amplifier per channel<br>Power dissipation = 650 mW (all on) |
| Passive heatsink | Xyratex inventory | custom | Attached to PCB onto copper slugs forming thermal contact with VCSEL Driver / Optical Receiver die |
| VCSEL array (chip) | ULM Photonics | ULM850-10-TT-C0104U | 10 Gb/s Vertical Cavity Surface Emitting Laser array (1 x 4)<br>Power dissipation = 20 mW (max) |
| PIN array (chip) | Microsemi | LX3045 | 10 Gb/s PIN photodiode array (1 x 4) |
| MT pin block | Extec | custom | Precision ceramic mount incorporating to standard MT pins |
| GRIN lens array and holder | GRINTEch | GT-LFRA-025-12-048-NC (custom) | 1 x 12 GRIN lens array embedded in custom ceramic holder to accommodate MT pins |

| **Flexible Bridge** | | | |
|---|---|---|---|
| No active or passive components | | | |
| **Base** | | | |
| Signal equalisers | Maxim | MAX3805 | 10.7 Gb/s Adaptive Receive Equalisers to improve signal integrity of 10 Gb/s differential signals received from TIA/LA<br>Power dissipation = 135 mW per device. Total = 4*135 mW = 540 mW |
| Mezzanine connector | FCI Connect | 55740c | Copper array connector receptacle supporting differential signals up to and exceeding 10 Gb/s data rate and low speed control signals. |
| Fan connectors | | standard | Attached to PCB onto copper slugs forming thermal contact with VCSEL Driver / Optical Receiver die |
| **Device efficiency** | | | |
| **Total estimated power dissipation** | | 570 mW + 650 mW + 20 mW + 540 mW = 1780 mW = **1.78 W** | |
| **Aggregate bandwidth** | | 8 unidirectional channels * 10.3 Gb/s = **82.4 Gb/s** | |
| **Efficiency** | | 1780 mW / 82.4 Gb/s = 21.6 mW / Gb/s = **21.6 pJ/bit** | |

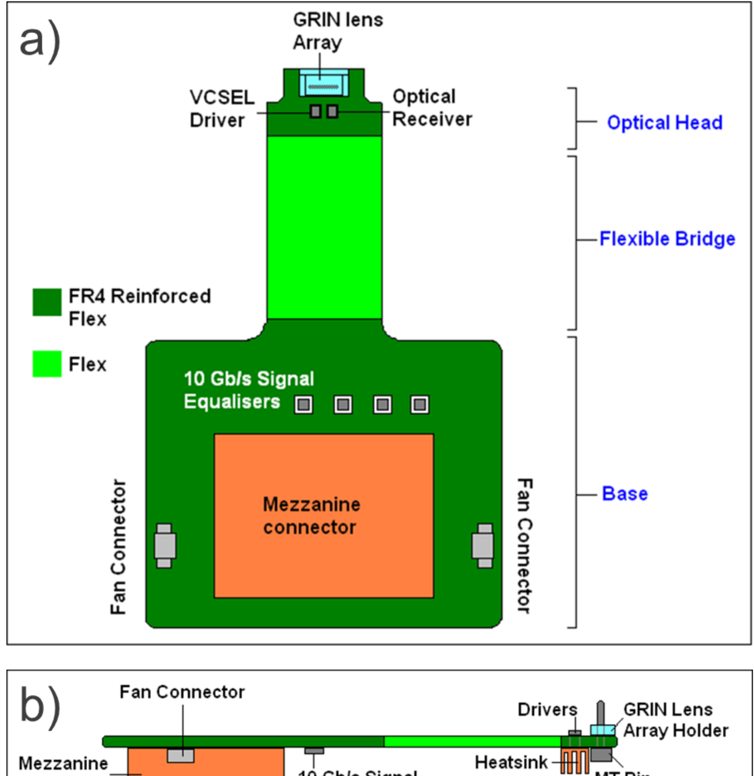


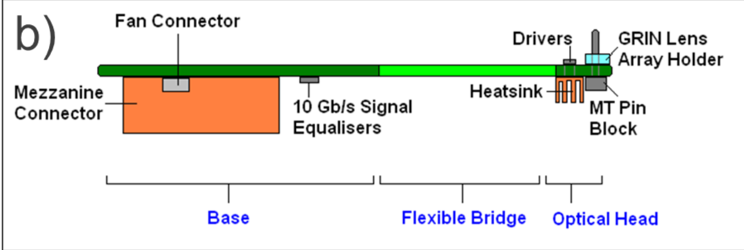


**Figure 3-16: StorConn2 component layout: a) top view, b) side view**

### *3.4.3.8. StorConn2 card assembly process*

The complete assembly of the StorConn2 card required, in addition to the conventional process of card population with standard passive and active components, the direct wire-bonded assembly of the VCSEL drive and TIA/LA array chips in raw die form, the integration of novel thermal dissipation structures and high-precision passive alignment processes for both chips and opto-mechanical structures.

The assembly process and sequence is outlined below.

### *3.4.3.9. Mezzanine connector placement*

The **GIG-ARRAY®** high speed mezzanine connector (Figure 3-17) is designed to provide high speed differential (10 Gb/s) and single ended electrical connection between two parallel boards. The connector utilises Ball Grid Array (BGA) for solder attachment to the PCB. For initial alignment, the plug housing has a

chamfered lead-in that captures and guides the receptacle cover and to assure proper mating orientation, both parts are keyed appropriately.

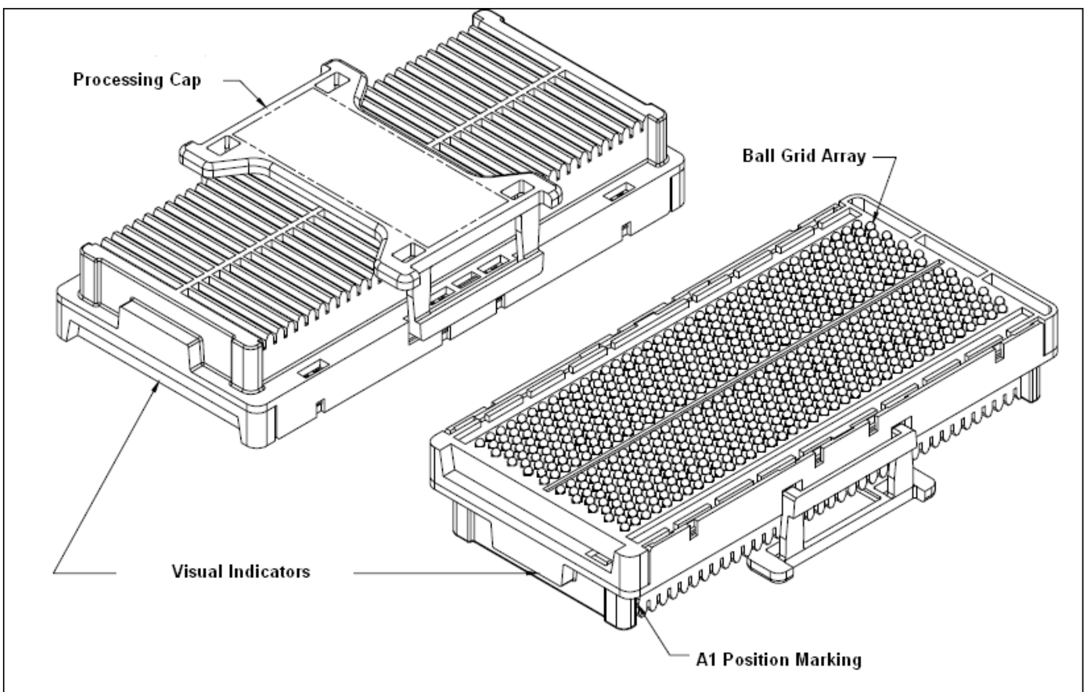


**Figure 3-17: Gig array receptacle (*Source: FCI Interconnect - GS-20-016 Rev B)***

### *3.4.3.10. Equaliser chip assembly*

Maxim **MAX3805** is a 10.7 Gb/s Adaptive Receive Equaliser designed to improve jitter on high speed signals, by reducing inter-symbol interference and adapting the signal to frequency dependent skin effect and dielectric losses.

The Maxim **MAX3805** chip was packaged in a thin "quad-flat no-leads" (QFN) surface mount package with a thermal pad on the base and had to be assembled onto the board by means of a solder reflow process (Figure 3-18).

Subsequently, all passive components including resistors and capacitors were solder-attached to the card in the standard manner.

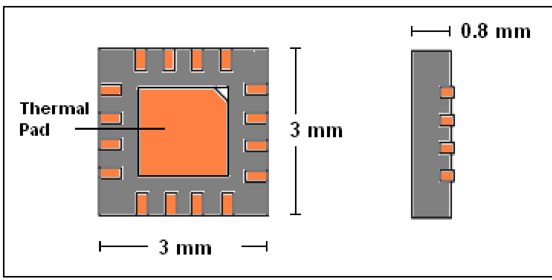


**Figure 3-18: Equaliser Chip – 10.3 Gb/s**

### *3.4.3.11. Thermal dissipation structures*

The VCSEL driver and optical receiver die dissipated a large amount of power relative to their size. In order to effectively dissipate heat from the area, solid, thermally-conductive structures were put in place to form direct thermal contact with the base of the die and dissipate the heat into the heatsink on the opposite side of the card. These cylindrical structures (slugs) traversed the card through hollow vias directly under the die and served as a more effective alternative to thermal vias.

The slugs alone were inserted into the hollow vias and fastened to the side walls with a thermally conductive epoxy.

The heatsink was then attached to the base of the card with the protruding slugs inserted into compliant slots in the heatsink housing.

### *3.4.3.12. Optical driver die assembly*

Primarion® **PX6514** was a four-channel VCSEL driver designed for various 4 x 10 Gb/s parallel optics applications. It consisted of a DC-coupled amplifier with selectable modulation and bias currents optimised for driving commercially available, common cathode VCSELs from a single +3.3 V supply. Primarion® **PX6524** was a four-channel TIA/LA optical receiver designed for various 4 x 10 Gb/s parallel optics applications. It consisted of a DC-coupled trans-impedance amplifier and an AC-coupled differential limiting amplifier.

The Primarion **PX6514** VCSEL driver and **PX6524** optical receiver die were positioned over the vias now populated with the thermal dissipation slugs and attached to the top surface of the slugs with a thermally conductive epoxy, thus establishing a thermal channel to the underlying heatsink. The die pads were then wire-bonded to the compliant pads on the PCB.

### *3.4.3.13. MT pin block assembly*

The MT pin block is a custom high-precision holding piece for two MT pins (Figure 3-19). In accordance with the NTT standard; these pins require a diameter of 0.7 mm and are mounted onto a planar surface at a pitch of 4.6 mm. The block will be fastened to the underside of the optical head with the pins protruding through a compliant slot in the PCB. The block is composed of machinable ceramic, which satisfies the requirements for thermal resilience and mechanical stability.

The MT pin block was attached to the base of the card such that the pins and the adjoining alignment shelf protruded up through the compliant slot in the PCB.

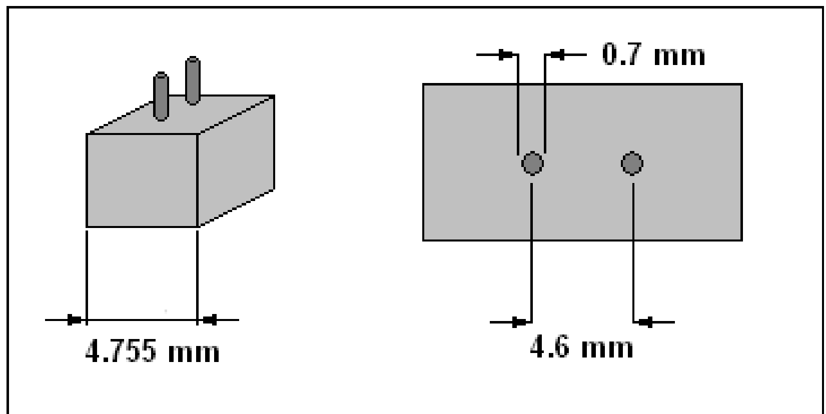


**Figure 3-19: MT pin block**

### *3.4.3.14. VCSEL array and photodiode array assembly*

The ULM Photonics **ULM850-10-TT-C0104U** quad (four channel) VCSEL array, capable of sustaining modulation speeds of 10 Gb/s per channel at a nominal wavelength of 850 nm, was presented in an unsealed chip form with the cathode / anode bond pads and the active area located on the top surface.

The Microsemi **LX3045** quad (four channel) GaAs coplanar PIN photodiode array, capable of processing optical signal modulation of 10 Gb/s at a detection wavelength of 850 nm, was presented in an unsealed chip form with the cathode / anode bond pads and the active area located on the top surface.

The quad VCSEL array and quad PIN photodiode array die were visually aligned through the microscope on the chip placement and wire bonding machine relative to the MT compliant alignment pins protruding through the board from the MT pin block to within a lateral positional accuracy of 5 µm (Figure 3-20).

The bases of the VCSEL array and PIN array were then secured in place over the thermal vias on the PCB with a thermally conductive epoxy.

The cathode and anode pads on the VCSEL and PIN arrays were attached to the recipient bond pads on the PCB with bond wires.

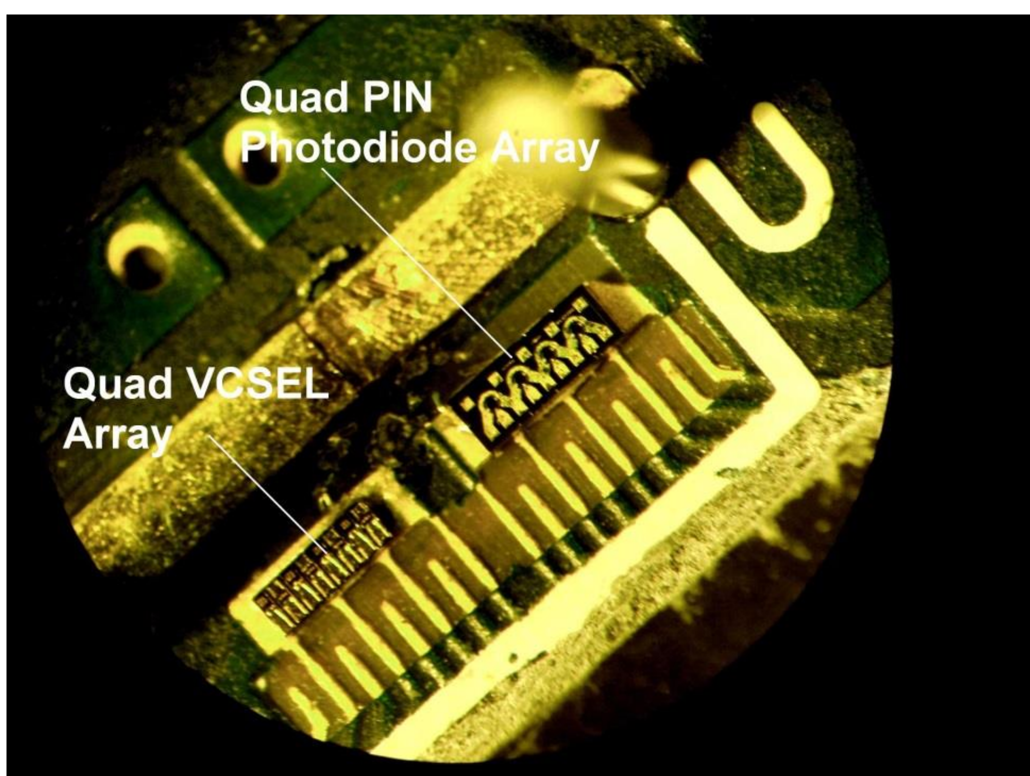


**Figure 3-20: Photo of VCSEL array and photodiode assembly prior to wire bonding (photo through microscope)**

### *3.4.3.15. GRIN lens assembly*

The GRIN lens array, embedded in a ceramic holder, was lowered over the die with the MT pins protruding through the two compatible slots in the ceramic. The MT slots which are slightly larger than the MT pins allow for initial self-alignment of the lenses to the active areas of the underlying die. As these are imaging lenses, the images of the underlying die were clearly visible, allowing for reasonably precise visual alignment of the centres of the active areas with the centres of the lenses (Figure 3-21).

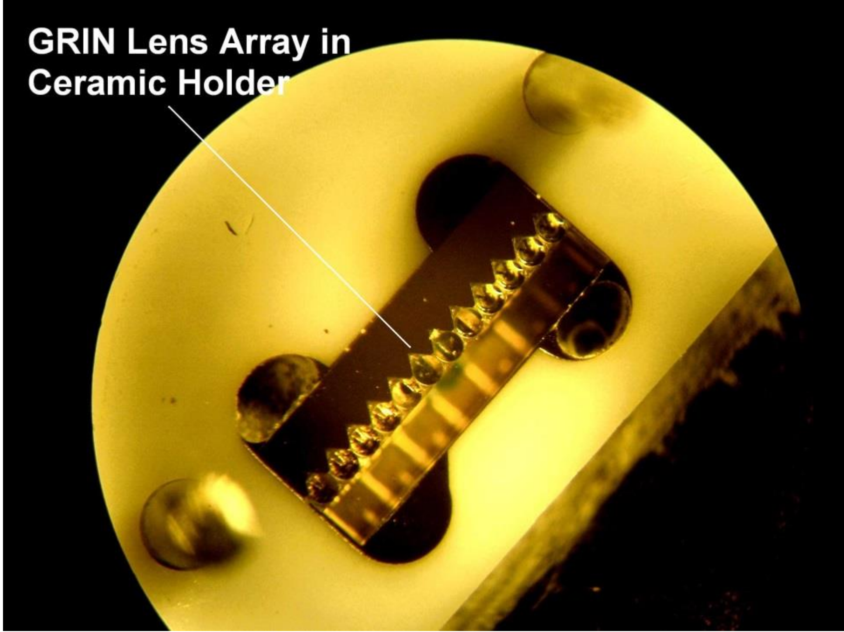


**Figure 3-21: GRIN lens assembly (photo through microscope)**

### 3.4.4 Storlite connector mechanism

The author developed a full requirements specification for the connector concept outlined below and worked with Xyratex mechanical engineer Chris Smith who rendered the design concept into a viable mechanical model and system.

#### *3.4.4.1. Connector composition*

The parallel optical transceiver is housed within a mechanical casing incorporating a pivoted lever, which is driven by the cam handle. The pivoted lever is bifurcated to provide a balanced tension to the photonic interface section of the transceiver PCB. Thus, by operating the cam handle, the photonic interface can be retracted and elevated as required.

#### *3.4.4.2. Engagement mechanism*

The connector mechanism served two key functions:

1. Retraction of the photonic interface to protect the salient mechanical alignment structures during the test daughtercard insertion and retraction process (Figure 3-22a).

2. Elevation of the photonic interface into engagement with compliant receptacle on optical PCB (Figure 3-22b).

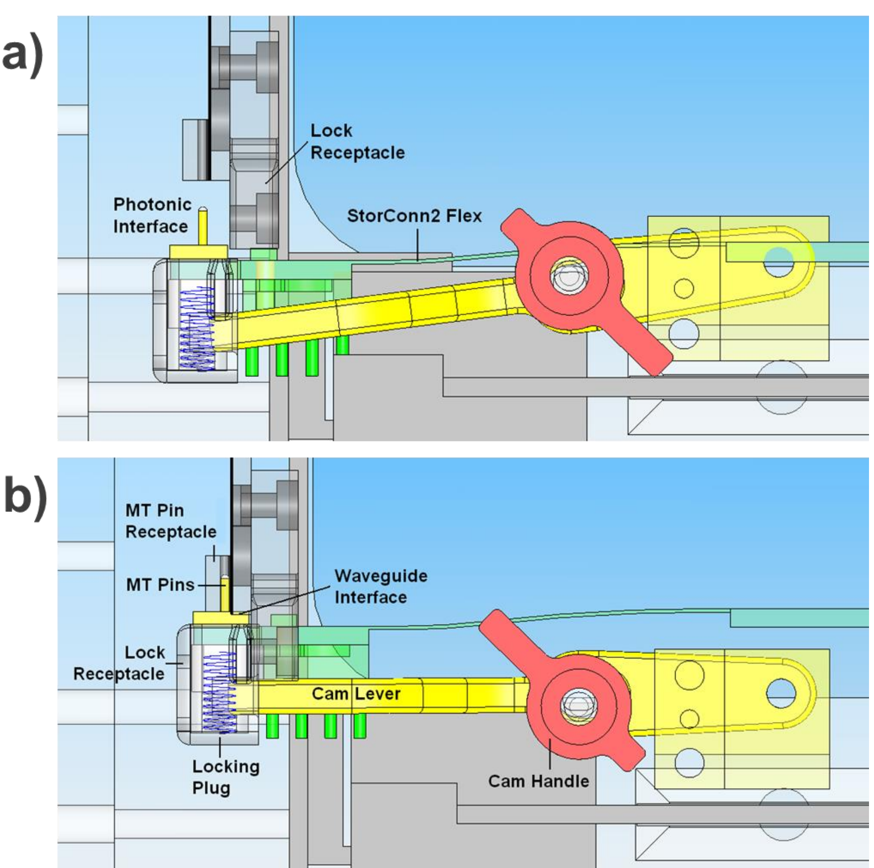


**Figure 3-22: Storlite connector CAD model - side view a) optical interface retracted, b) optical interface elevated**

Figure 3-23 shows photos of the connector with photonic interface retracted (Figure 3-23a) and photonic interface elevated (Figure 3-23b).

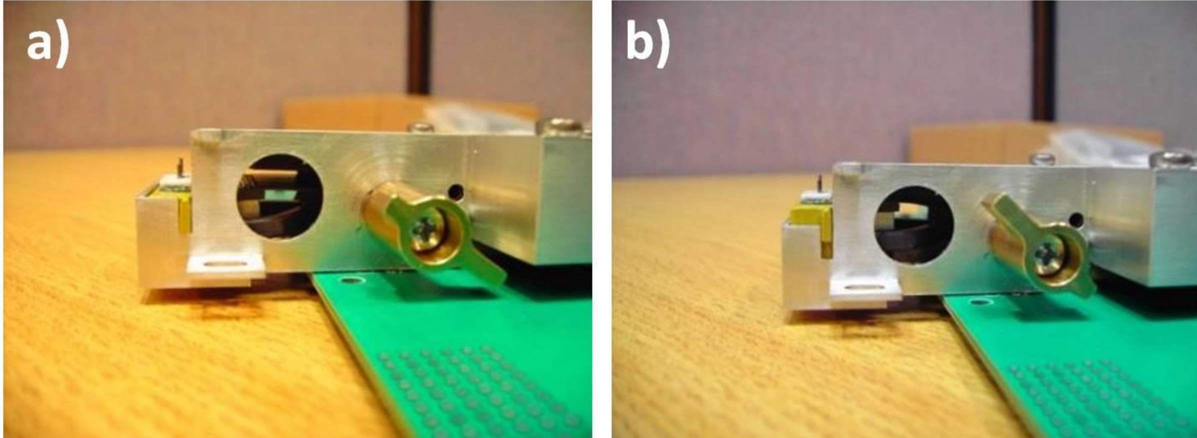


**Figure 3-23: Photo of Storlite connector - side view a) photonic interface retracted, b) photonic interface elevated**

The connector assembly including the transceiver board (Figure 3-24a) was electronically plugged to the line-card with the GigArray® connector (Figure 3-24b) and fastened at the screw points. The test daughtercard was guided within the chassis along standard guide rails to the point of engagement with both the optical backplane and a separate electrical backplane. At the first level of optical engagement, guide features on the connector housing brought the optical interface into coarse alignment with the compliant receptacle.

At the second level of engagement, the cam handle was manually turned and the optical interface elevated, such that the MT alignment pins on the interface are made to mate with the MT slots on a compliant receptacle, which is mounted to the optical PCB, thus forcing the optical array into a precise butt-coupled arrangement with the waveguide array on the backplane.

All precision assembly on the design is thus restricted to the optical interface and the compliant receptacle.

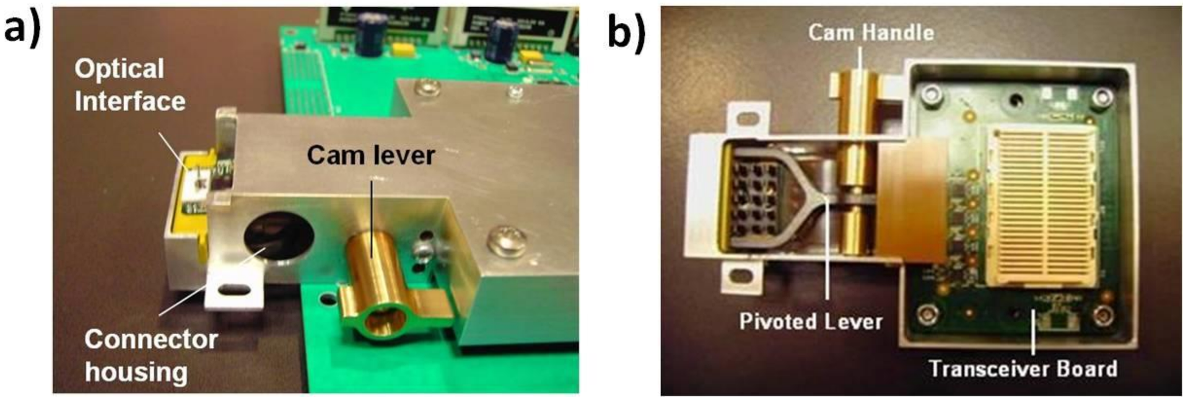


**Figure 3-24: a) Storlite connector with retracted optical interface, b) Storlite connector bottom view showing bifurcated pivoted lever mechanism and electronic high speed GigArray® connector**

A full exploded view of the Storlite connector is shown in Figure 3-25.

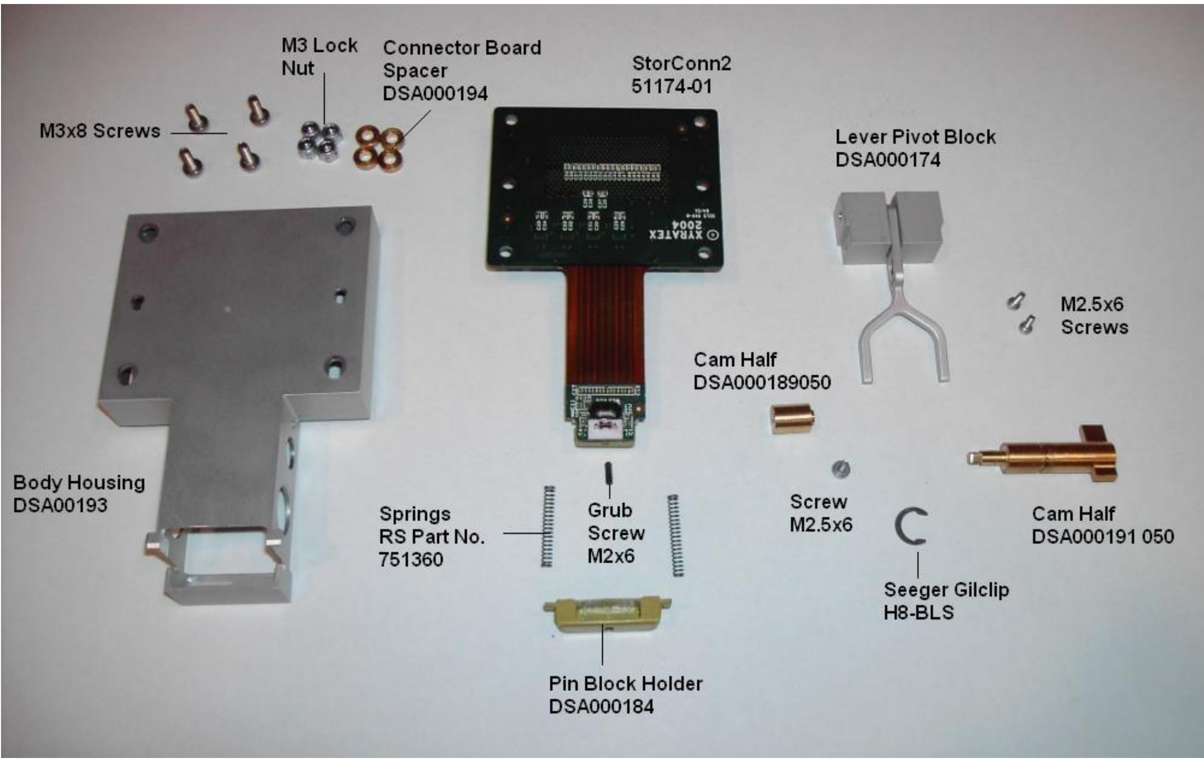


**Figure 3-25: Storlite connector exploded view**

The fully assembled Storlite connector attached to the test line card is shown in Figure 3-26.

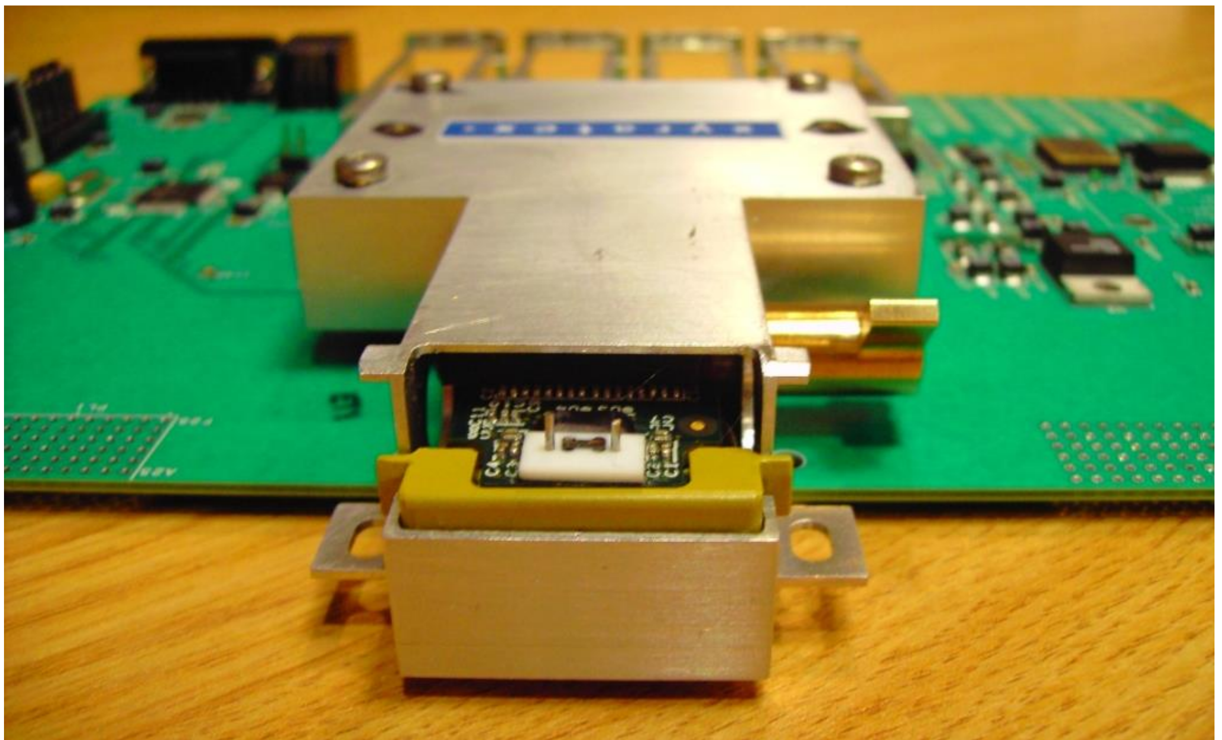

**Figure 3-26: Storlite connector attached to line-card - front view**

## 3.4.5 StorConnOpt2 - Storlite optical waveguide backplane

The *StorConnOpt2* optical backplane is a purely passive optical printed circuit board with one optical layer and no electronic layers.

The optical waveguide parameters for the experimental evaluation of the proposed interface and coupling method were designed to be compliant with the optical interface on the Storlite transceiver. As a result, the waveguide interface on the StorConnOpt2 OPCB could also be tested with MT patch-cords.

### *3.4.5.1. Optical waveguide design layout and dimensions*

The design of the StorConnOpt2 OPCB contained four groups of 12 straight parallel waveguides with a centre to centre pitch of 250 µm and waveguide cross-section of 70 µm x 70 µm. The waveguides extended over a distance of 100 mm from one physical ingress point at one end of the PCB to the egress point at the other end (Figure 3-27a).

The waveguide core and cladding were composed of a cross-linked polymer acrylate, Truemode®, which was deposited on a rigid FR4 substrate and patterned lithographically as will be described in section 3.4.5.2.

The refractive indices of the polymer core and cladding layers were 1.556 and 1.5264 respectively at 850 nm, thus yielding a N.A. of 0.302 for the waveguide.

As the waveguide N.A. was larger than that of the imaged VCSEL, good power coupling was ensured between the active optical transmitter on the line-card and the recipient waveguide on the optical PCB (Figure 3-27b).

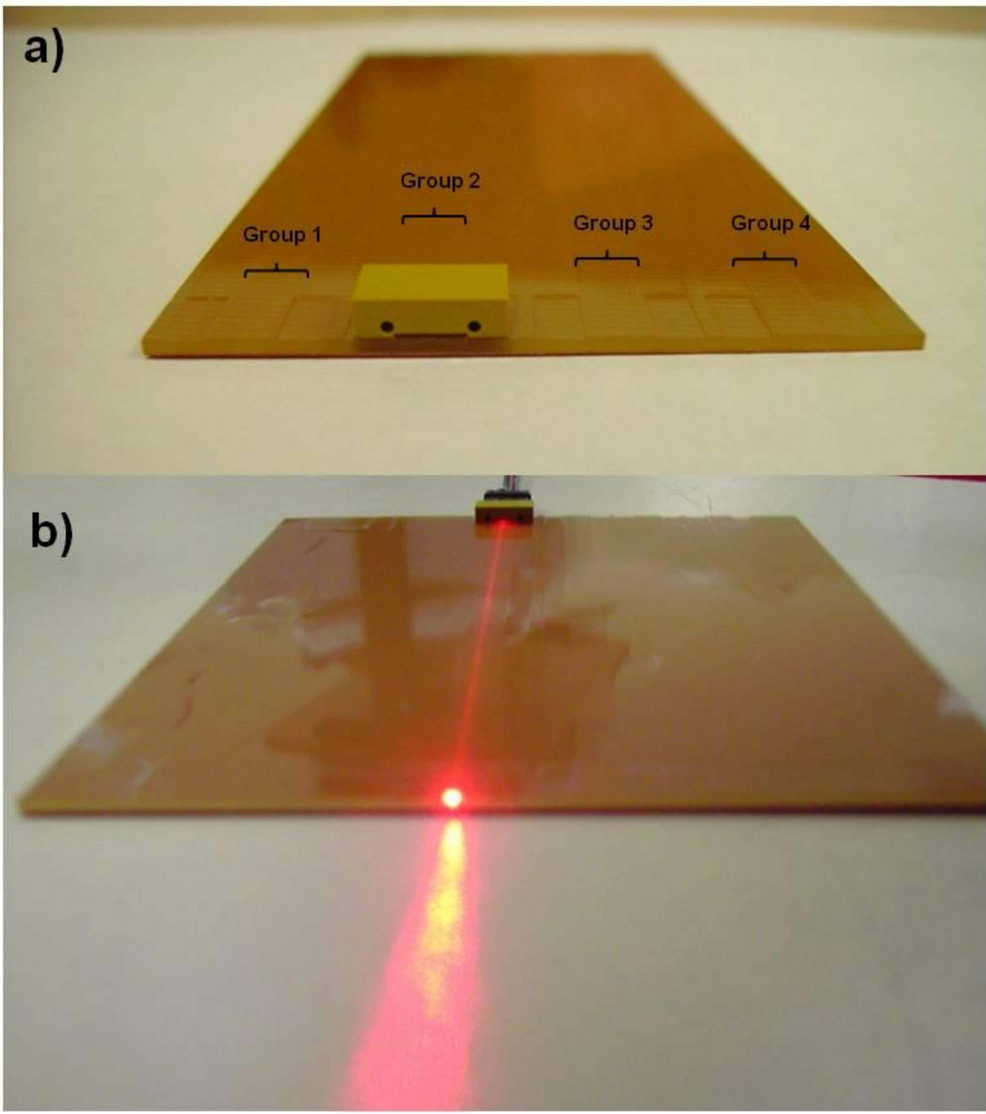


**Figure 3-27: Storlite optical waveguide PCB: a) four groups of 12 parallel straight waveguides. MT compliant receptacle with MT pins slots offset from waveguide array by 625 µm (centre to centre), b) Single waveguide illuminated with 635 nm visible light to demonstrate effectiveness of passive alignment method and devices.**

### *3.4.5.2. Truemode® optical waveguide material*

The waveguide core and cladding were composed of a highly cross-linked polymer acrylate, Truemode®, which offers sufficient thermal stability and humidity resistance to withstand conventional PCB manufacturing conditions [74][141] and has been prototyped successfully for optical backplane applications [82], [139], [142], [143].

The Truemode® optical polymer comprised a proprietary polymer acrylate and photoinitiator blend. The photoinitiator catalyses a crosslinking reaction in the polymer when exposed to UV light, which increases the photosensitivity of the polymer blend and allows the polymer to be cured solely through exposure to curing radiation without the need for thermal curing steps. Figure 3-28 shows a photo of part of the waveguide interface with four parallel waveguides back-illuminated in a shadowgraph. The reticule grating pitch is set to 20µm.

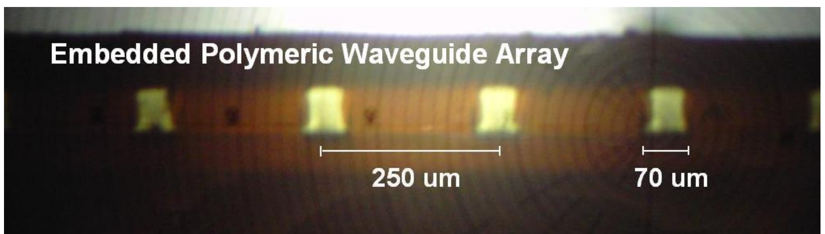


**Figure 3-28: Polymeric optical PCB interface**

**Table 3-3: Optical waveguide testbed attributes**

| | |
|---|---|
| Material | UV-curable polymeric acrylate (Truemode®) |
| Propagation loss @ 850nm | 0.04 dB/cm |
| Heat degradation resilience | 350° |
| Waveguide size | 70 µm x 70 µm |
| Waveguide centre to centre pitch | 250 µm |
| Refractive index core | 1.556 |
| Refractive index cladding | 1.526 |
| Numerical aperture | 0.302 |

### *3.4.5.3. Two-step photolithographic fabrication process for Storlite waveguide OPCB backplane*

The Storlite passive optical backplane comprised an FR4 substrate with a polymeric multimode optical waveguide layer patterned onto it. The fabrication process (Figure 3-29) detailed below included a first photolithographic step in the core layer to pattern the core waveguide features and a second photolithographic step in the upper cladding layer to create a clearance in the upper cladding allowing mechanical access to core features peripheral to the signal waveguides to enable high precision passive mechanical alignment and assembly

of the connector receptacle with respect to the waveguide interface.

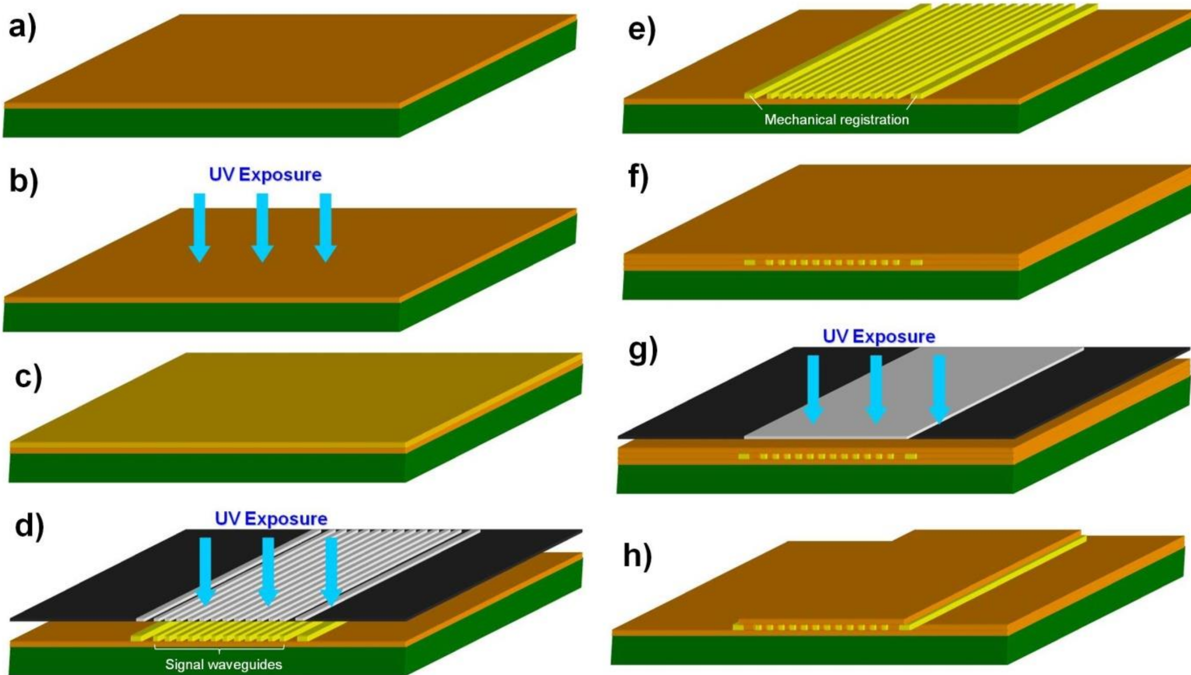


**Figure 3-29: Fabrication process for Storlite optical waveguide PCB**

Truemode® is a negative polymer material, in which those portions of the material exposed to curing radiation will become insoluble to developer (solvents), while those portions not exposed to light can be removed with a developer. Therefore the photolithographic mask will allow light to pass through in those areas, which are to form the hardened waveguide structures.

The two-step photolithographic fabrication process is detailed as follows:

a) A viscous liquid variant of the photosensitive Truemode® optical polymer acrylate is spin-coated onto the surface of an FR4 substrate.

b) Ultraviolet light is applied to the liquid polymer coating. As the polymer is photosensitive, crosslinking occurs and the polymer is cured. This polymer has a refractive index of 1.5264 at 850 nm in its cured state and comprises the cladding variant of the polymer, which surrounds the higher refractive index core channels to form the waveguides. This layer comprises the lower cladding layer and has a thickness of 150 µm.

c) A viscous liquid higher refractive index variant of the Truemode® polymer acrylate is spin-

coated onto the surface of the cured lower cladding layer

d) A photolithographic mask is aligned over the deposited core polymer layer and UV light is passed through the mask exposing only those areas, which are to form the 12 signal waveguides and the two peripheral mechanical registration waveguides. This polymer has a refractive index of 1.556 at 850 nm in its cured state and comprises the core variant of the polymer.

e) Uncured areas of the core layer are removed with developing agent and washed away leaving only the salient signal and mechanical registration waveguides. This layer comprises the core layer and has a thickness of 70 µm.

f) A liquid cladding variant of Truemode® polymer is spin-coated over the patterned core layer to a thickness sufficiently large to ensure that the core features are completely covered.

g) A second photolithographic mask is aligned over the deposited upper cladding layer. The mask includes a broad clearance allowing the upper cladding over all signal waveguides to be exposed to UV light and cured. The clearance is chosen such that at least the outer edges of the peripheral mechanical registration waveguides and a region outside these outer edges are not exposed to UV light and cured.

h) The uncured areas of the upper cladding area are removed with developing agent and washed away. The uncured areas include the outer edges of the peripheral mechanical registration waveguides and a clearance region around them. This will provide the means to passively align and assemble the MT compliant optical receptacles onto the optical PCB edges.

### 3.4.6 Storlite optical waveguide receptacle

The author designed an MT compliant optical waveguide receptacle that can be passively aligned onto the waveguide interface. My designs were rendered into CAD models by Xyratex mechanical engineer Chris Smith and used to fabricate the waveguide receptacle piece.

The optical waveguide receptacle was a custom high-precision unit housing two MT compliant slots. In accordance with the 6x12 MT interface requirements shown in Figure 3-14, these slots require a diameter of 0.7 mm and are mounted onto a planar surface at a pitch of 4.6 mm.

The receptacle unit contained two mounting feet, which are critical to the passive self-alignment process outlined in section 3.4.7 below.

Figure 3-30 shows the mechanical drawings for the optical waveguide receptacle, which was fabricated out of Torlon®, an amorphous (non-crystalline) engineering thermoplastic.

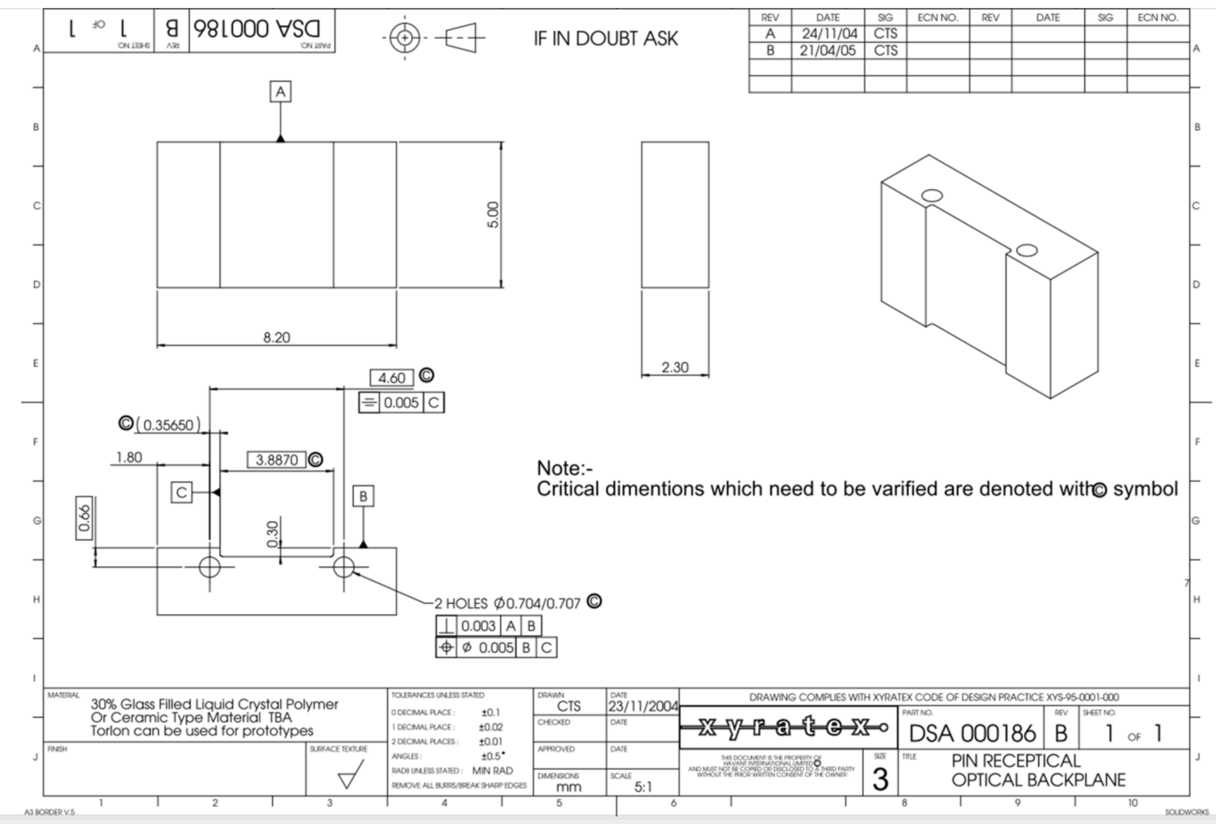


**Figure 3-30: Mechanical drawings of Storlite waveguide receptacle**

## 3.4.7 Passive optical assembly process

### *3.4.7.1. Patented method of high precision passive alignment onto optical PCBs*

The author is the lead inventor of a patent "Optical printed circuit board and manufacturing method" [144], which outlines a method of highly precise assembly of components onto OPCBs with respect to optical waveguides using a passive self-alignment process. Figure 3-31 shows two images from the published patent. The key premise of the patent is that additional features are fabricated in the optical core layer on either side of the optical waveguide or group of waveguides. While the waveguide cores themselves need to be covered in a lower refractive index cladding, the outer edges of the additional registration stubs must be mechanically exposed, that is the upper cladding around the outer edges must be removed. The height of typical multimode core structures are of the order 50 µm, which is sufficient to mechanically engage a high precision component with compliant registration features passively. This component can then be aligned laterally very accurately with respect to the waveguides themselves. Accurate vertical alignment is achieved through the use of the top surface of the bottom cladding on which the component rests.

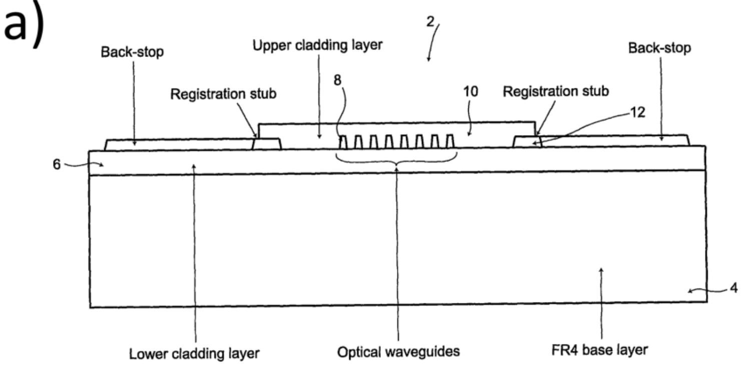


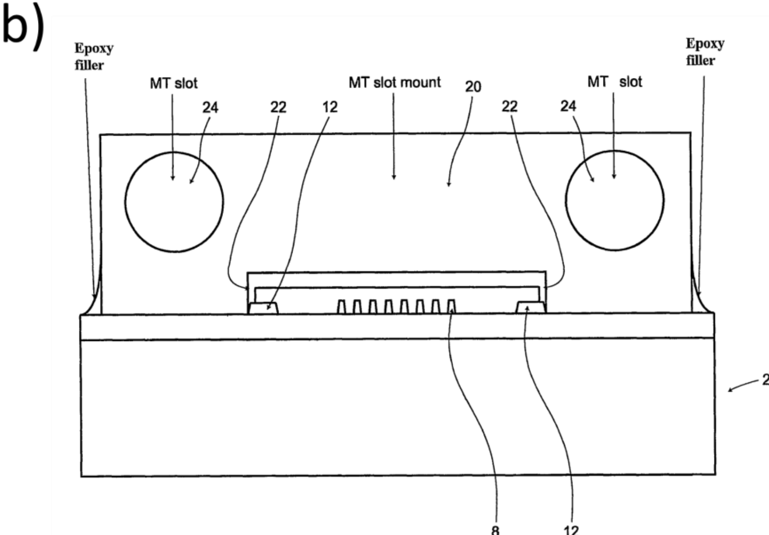


**Figure 3-31: Diagrams from patent "Optical printed circuit board and manufacturing method"** [144]**: a) optical waveguide receptacle socket showing exposed registration stubs, b) optical waveguide receptacle assembled into optical waveguide receptacle socket**

This method of passive self-alignment has been adopted in recent years and forms the basis of modern waveguide termination schemes for both multimode and singlemode polymer waveguides [145] [146]. Another method of self-alignment of photonic components on OPCBs takes advantage of the surface tension effects of liquid solder on copper pads on a PCB [147]. This effect has been successfully exploited to produce accurate self-alignment of singlemode optical components with a misalignment of the order of 1µm demonstrated [148].

### *3.4.7.2. Alignment and assembly of optical waveguide receptacle*

The process of alignment and assembly of the Storlite optical waveguide receptacle is shown in Figure 3-32.

The mounting feet on the Storlite optical waveguide receptacle unit are passively inserted into the exposed registration recesses on the optical backplane *StorConnOpt2*, such that the following conditions are met.

- Condition one
  - the bases of the mounting feet rest against the top surface of the lower cladding. It is crucial that the exposed top surface of the lower cladding is completely clean and free of any debris or residue. Preferably the OPCB fabricator will provide a protective cover over these recesses, which can be removed only prior to the assembly process.
- Condition two
  - the inner edges of the mounting feet are in contact with the outer edges of the registration stubs
- Condition three
  - the front face of the unit is flush with the polished waveguide interface (Figure 3-32c)

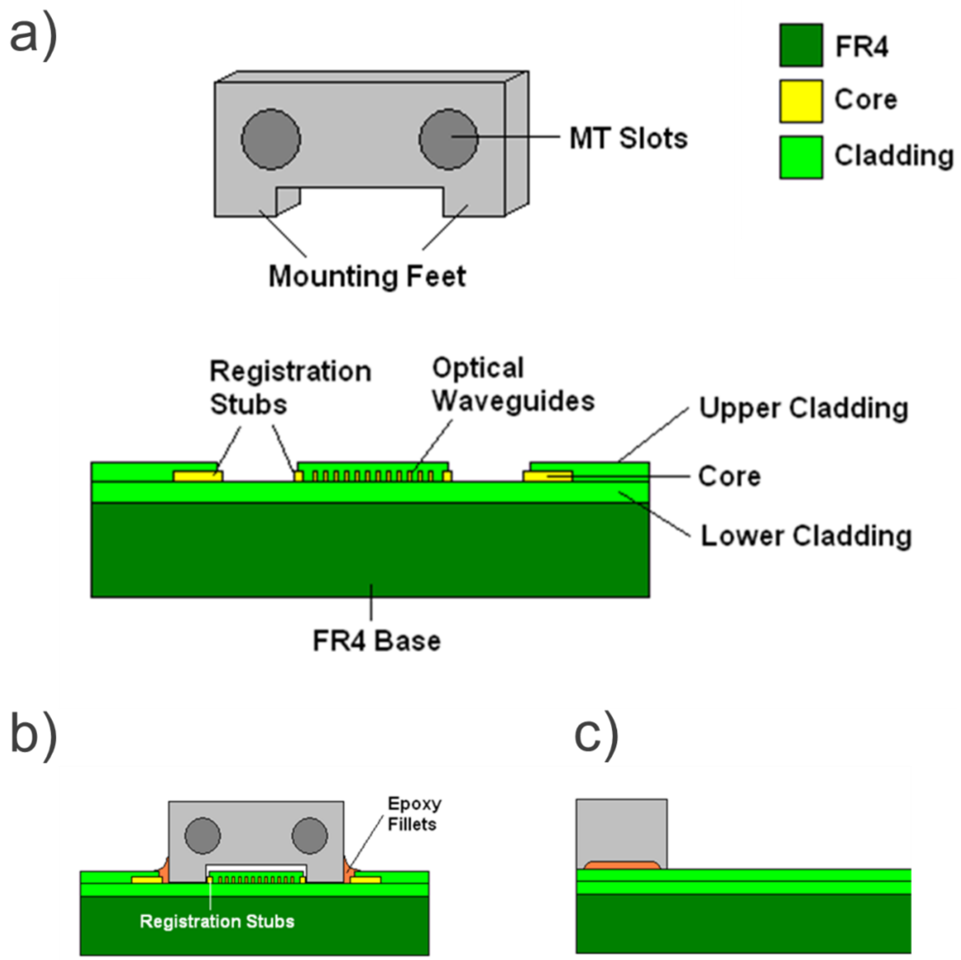


**Figure 3-32: Passive self-alignment assembly procedure: a) Storlite waveguide receptacle and OPCB waveguide interface with mechanical recesses and exposed registration structures, b) Waveguide receptacle assembled onto OPCB (view facing waveguide interface), c) Waveguide receptacle assembled onto OPCB (side view showing front face of unit flush with waveguide interface)**

The waveguide interface of a given waveguide group on the Storlite optical backplane contains 12 waveguides and two mechanical registration features on either side of the 12 communications waveguides. As can be seen in Figure 3-33 the mechanical registration waveguides are effectively very wide waveguides, the outer edges of which serve as mechanical registration points for the waveguide receptacle. The receptacle could thus be passively aligned to a very high accuracy with respect to the communication waveguides. Once aligned and assembled, the datum of the MT pin slots in the receptacle are offset from the communication waveguides by 625 µm centre to centre. This offset was deliberately chosen to be compliant with the offset in a commercial 6x12 MT ferrule between the MT pin slots and the lowest (or highest) row of fibres, thus allowing the waveguides to be tested with a 6x12 MT ferrule on which only the lowest row of fibres is populated and used for

measurement. As previously mentioned, the reason the transceiver optical interface, and waveguide interface were not made compliant with a standard 1x12 MT ferrule interface, whereby the row of optical channels is inline with the two MT pin slots (no offset), was that it would not have been possible for the MT pin slots in the receptacle to have been placed in line with the waveguide interface. As is evident from Figure 3-33, this would have required that clearance sections be diced immediately around the communication waveguides to allow each wing of the receptacle containing an MT pin slot to be lowered in line with the embedded waveguide array. Given dicing tolerances on FR4 of ~100 µm, it would not have been possible to achieve this around the full set of 12 waveguides as required.

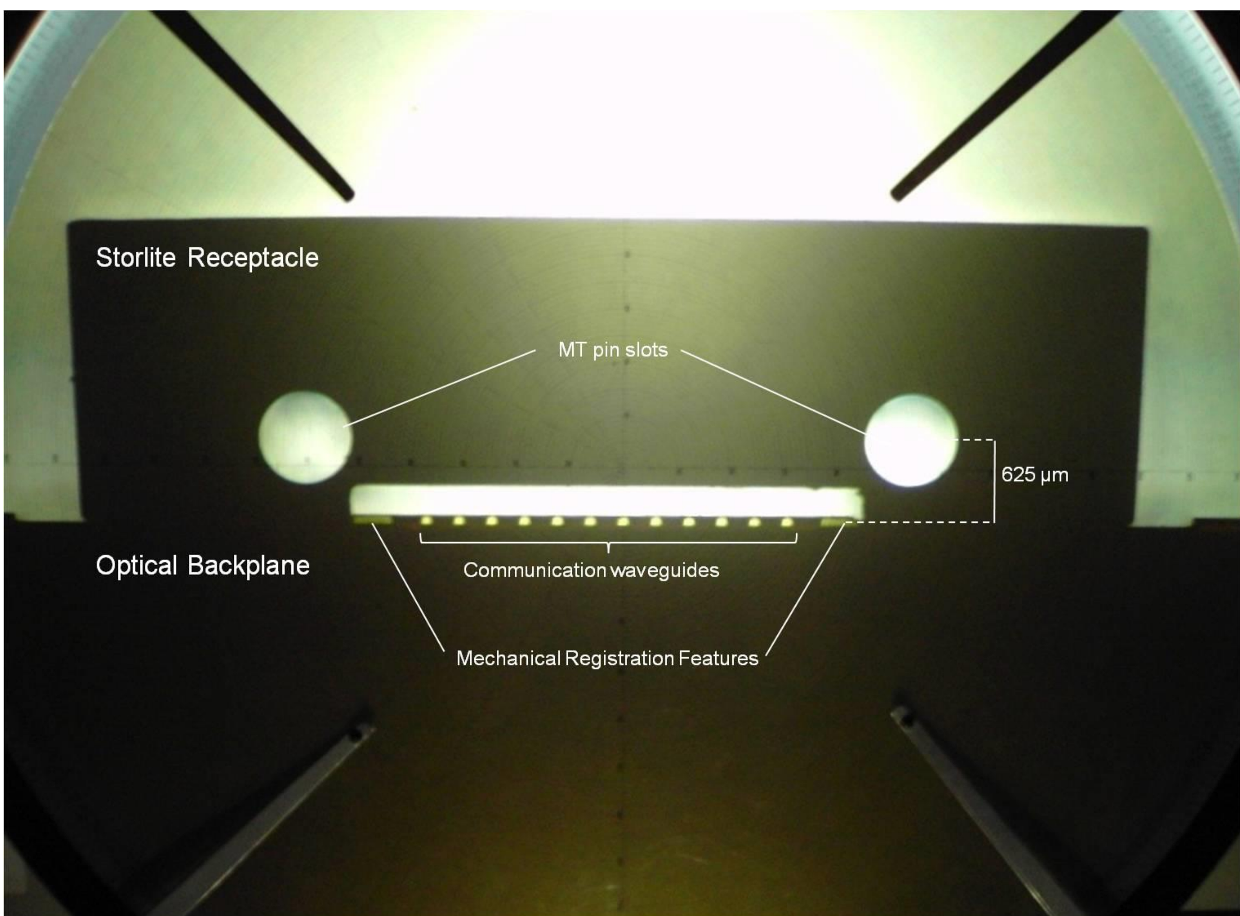


**Figure 3-33: Shadowgraph image of an optical waveguide group with a receptacle passively aligned over it using the mechanical registration features fabricated in the peripheries of the 12 communication waveguides**

In order to verify that the passive alignment has been successfully achieved prior to fastening the waveguide receptacle into place, a fibre fan-out jumper is used comprising a 6x12 MT ferrule on one end, of which the

lowest row only was populated with 12 individual multimode graded index fibres and each fibre terminated with an FC/PC fibre connector.

A first stage alignment verification process shown in Figure 3-34 required connecting one of the fibres to a visible 635 nm light source (Figure 3-34a) such that the corresponding channel in the MT interface was illuminated (Figure 3-34b). The MT ferrule was connected into the optical waveguide receptacle on its own (Figure 3-34c,d). The optical waveguide receptacle with the MT ferrule attached was passively aligned into the waveguide recess of the optical backplane. If the waveguide is visibly illuminated as shown in Figure 3-34e), this is an indication that at least partial alignment of the MT ferrule fibre and the waveguide has been achieved. A more stringent first stage validation of at least partial alignment of the entire MT ferrule and the waveguide interface would be to simultaneously illuminate the end channels (1 and 12) and to ensure that both waveguides are illuminated.

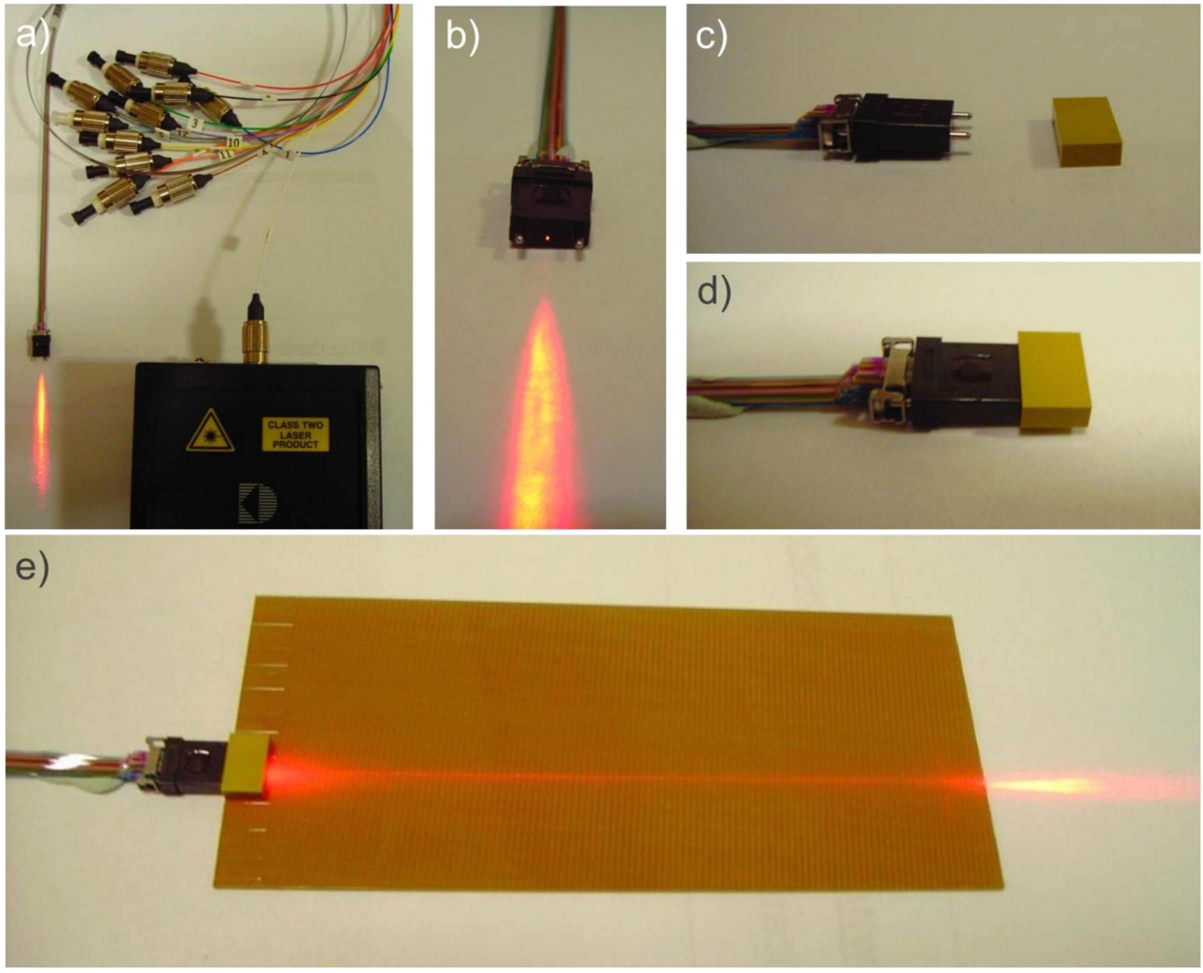


**Figure 3-34: First stage alignment verification using visible light: a) Class 2 635 nm laser source is connected to one fibre of a lowest row populated 6x12 MT to 12 LC fibre jumper, b) view of lowest row populated 6x12 MT interface with channel five illuminated (fifth channel from left), c) MT ferrule about to engage with optical waveguide receptacle, d) MT ferrule fully engaged with optical waveguide receptacle, e) optical waveguide receptacle passively aligned into waveguide registration recess with one waveguide channel illuminated**

This is only a first level alignment verification process. A second level alignment verification process would require, while the waveguide receptacle was held in place with the MT ferrule attached, each channel to be connected in sequence to a stable 850 nm source, and the light exiting the waveguide end facet on other side of the backplane measured directly. The level of light should be consistent with the expected performance of the waveguide. The expected performance of the OPCB waveguides would be determined prior to any assembly. The OPCB would be measured in a lab with the input launch source on an x-y-z translation stage used to align and butt-couple the source to the waveguide under test. A photodetector would be aligned to the other end of the

waveguide under test. The position of the launch source would be tuned until the photodetector measures the maximum amount of light through the waveguide. Ideally the moveable launch source is the same as the source that will be used in the actual system. In this case, the transceiver would need to be adapted such that it could be mounted on an x-y-z translation stage and powered. The MT pins in the Optical Interface section would need to be removed otherwise they would impede the lateral adjustment of the flat GRIN lens surface against the waveguide interface. In practise this can be cumbersome, so an optical test fibre with a similar core size to that of the VCSEL aperture can be used.

Once all these positional and measurement criteria are met then the unit can be fastened in place with epoxy (Figure 3-32b).

### *3.4.7.3. Misalignment tolerance of VCSEL and PD to waveguides*

A detailed study was carried out by Ioannis Papakonstantinou at UCL on the misalignment tolerances of VCSELs and PDs to waveguides on StorConnOpt2 board. Figure 3-35 shows 2 dimensional contour maps describing the changes in insertion loss as the lateral positions (parallel to the plane of the waveguide end facet) of a VCSEL is varied at the launch facet (Figure 3-35a) and the lateral position of a PD is changed at the receive facet (Figure 3-35b).

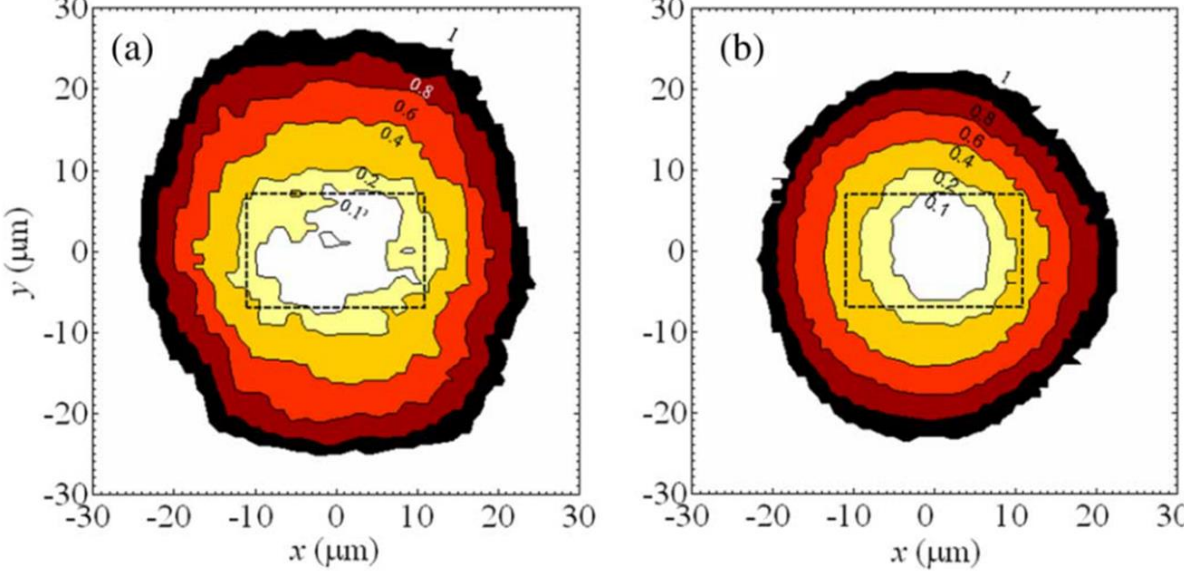


**Figure 3-35: Contour maps of relative insertion loss compared to the optimum coupling position for a) VCSEL misalignment at z = 0, b) PD misalignment** [2]

Furthermore, the reliability of repeated mating cycles of connectors to the waveguide receptacle was characterised and showed that during 75 repeated attachment and detachment cycles, the average insertion loss was measured to be 5.19 ± 0.16 dB. The results are reported in the joint journal paper [2].

#### *3.4.7.4. Waveguide end facet preparation*

The StorConnOpt2 OPCB was diced at each end to create the end facets of the waveguides. The surface quality of waveguide end facets is crucial for coupling efficiencies as end facet scattering is one of the key loss mechanisms in waveguides. In order to further reduce the surface roughness of the end facets, they were manually polished with optical cloths. Care was taken during polishing to avoid contamination of the soft polymer waveguide end facets by particulates released form the FR4 substrate during polishing. This was primarily achieved by wiping the sample in one direction along the optical polishing cloth leading with the optical layer and never wiping along the same section of the polishing cloth multiple times as contaminants from the FR4 layers will be lodged in the cloth after the first wipe. As a matter of course, the end facets were polished with different optical polishing cloths in sequence of decreasing coarseness.

Figure 3-36a shows a Scanning Electron Microscope image of a group of 12 waveguides on the Storlite optical waveguide backplane including the peripheral recesses in the upper cladding layer to enable mechanical access to mechanical registration features. Figure 3-36b shows a close-up of the waveguide interface. The degradation at the top upper cladding lip over the communications waveguides is clearly visible. Certain materials, such as, in this case, polyacrylates, will take on a brittle nature once cured, and this makes it easier for outer sections to crumble away when subjected to the unidirectional shearing forces consistent with the manual polishing process described above. This is a common issue with manual polishing of OPCB end facets. Furthermore this angled degradation of the upper cladding in the interface region is exacerbated when the polishing cloth is not moved exactly parallel to the plane of the waveguide end facet, but rather at a slight tilt.

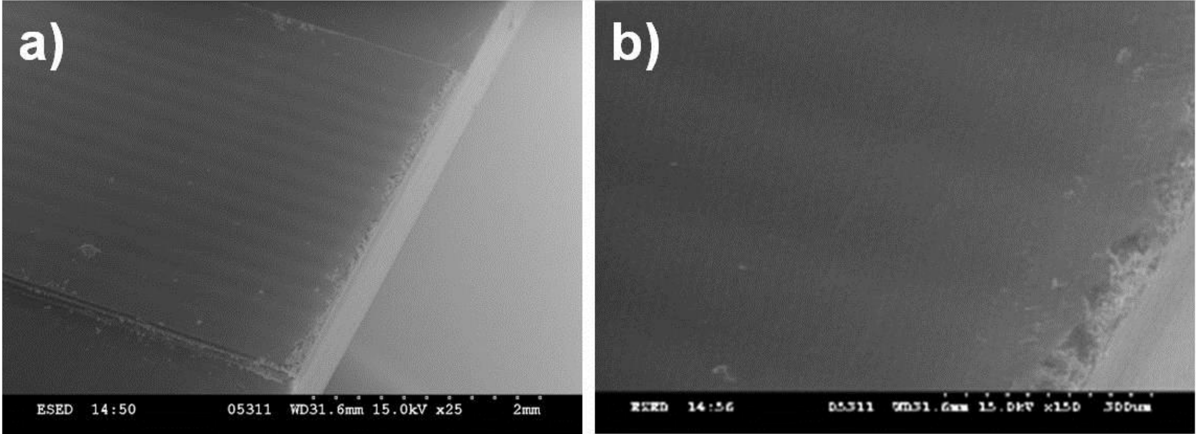


**Figure 3-36: SEM images of waveguide array interface which has been diced and polished: a) Full view of array of 12 communications waveguides and peripheral recesses in upper cladding layer to enable mechanical access to mechanical registration features, b) Close view of roughness and surface degradation over waveguide interface**

Figure 3-37a shows a photo of the optical backplane with waveguide receptacle aligned and assembled next to the constituent parts for the receptacle holder, which guides the optical connector housing into place during the insertion of the test daughtercard. The receptacle holder serves as a coarse alignment structure bringing the photonic interface of the transceiver into alignment over the waveguide receptacle. The waveguide receptacle is the secondary fine alignment structure, which receives the MT pins of the transceiver photonic interface and aligns the GRIN lens apertures accurately to the waveguide end facets.

Figure 3-37b shows a photo of optical waveguide interface with secondary alignment structure (receptacle holder) and primary alignment structure (waveguide receptacle) assembled.

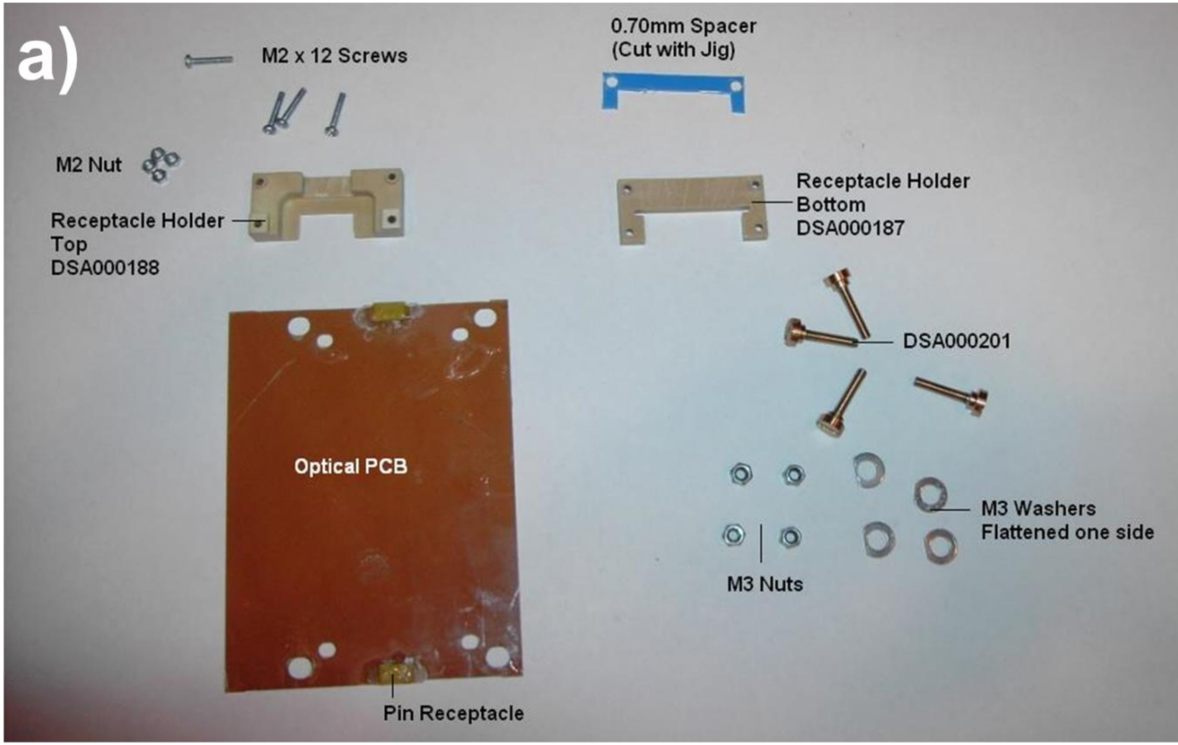


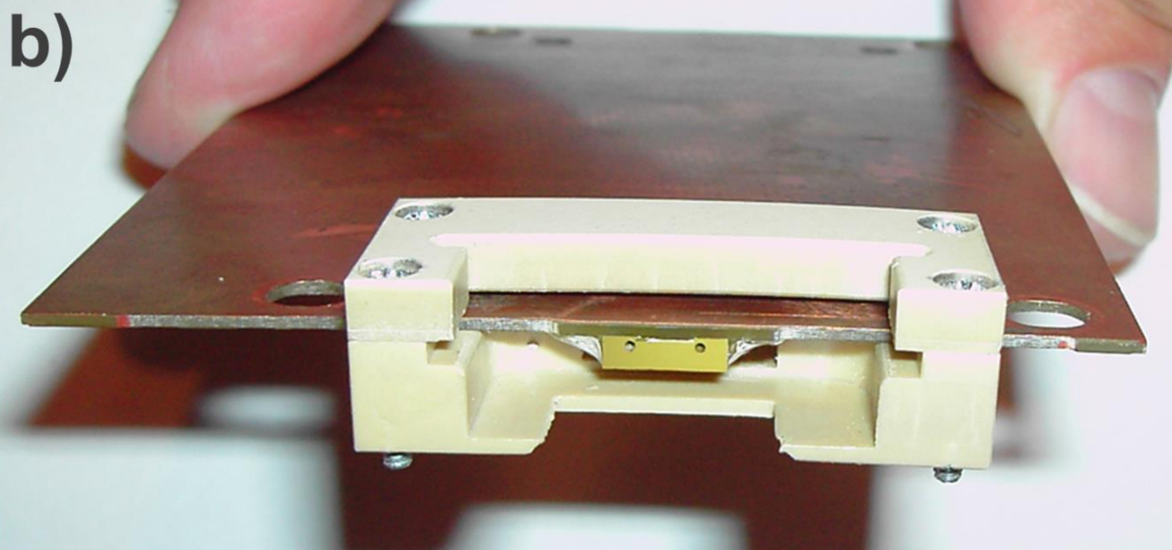


**Figure 3-37: optical backplane assembly: a) Exploded view of optical backplane assembly and parts including the receptacle holder, b) optical waveguide interface**

### 3.4.8 StorConnTest2 - Storlite 10 Gb/s test daughtercard

The purpose of the test daughtercards was to route high speed test data from an external 10 GbE LAN network analyser via the commercial XFPs on the front end to the StorConn2 transceiver circuit, which in turn is optically connected to the backplane. Each test daughtercard contained a GigArray® connector receptacle (Figure 3-38a) to accommodate the compliant GigArray® plug on the base of the StorConn2 transceiver and thus allow the fully assembled connector to be electronically attached to the line card. A microcontroller (Figure 3-38b) was included primarily to allow direct configuration of the StorConn2 transceiver parameters including VCSEL modulation and bias current, receiver squelch and enable all four transmit and receive channels. In addition the microcontroller was used to enable the four commercial XFPs (Figure 3-38c) mounted on the front end of the line card, which were required to convert external optical test data into high speed electronic test data and to be conveyed to the StorConn2 transceiver and vice versa. External communication with the test daughtercard was enabled through the RS232 connector and transceiver on the front end (Figure 3-38d).

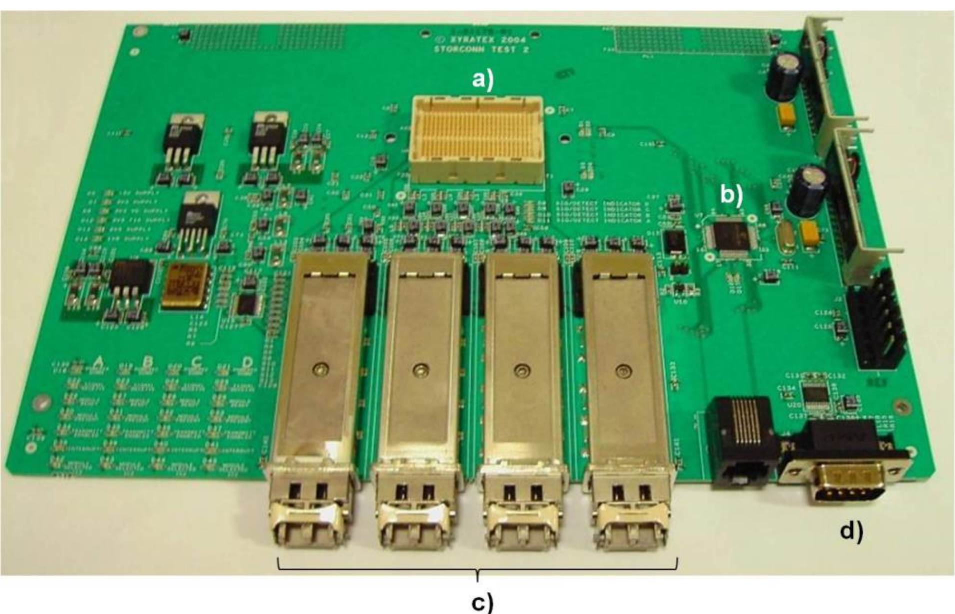


**Figure 3-38: Storlite test daughtercard: a) GigArray® connector receptacle, b) microcontroller, c) four commercial XFP transceivers, d) RS232 connector and transceiver**

#### *3.4.8.1. Functional overview*

Figure 3-39 provides a functional diagram of the StorConnTest2 physical layer relay card, which includes the following key features:

- The card houses four host ports to accommodate commercial XFP transceivers required for the transfer of 10.3 Gb/s test data between external protocol analyser devices and the adjoining StorConn2 circuit.
- PECL clock distribution to all XFP ports of 161.132 MHz to conform to XFP support of 10.3 Gb/s GbE test data
- Mezzanine connector to deliver all electrical and electronic signals including power, high and low speed and static control signals to the StorConn2 circuit.
- Microcontroller to regulate the functions of the StorConnTest2 XFPs and the StorConn2 transceiver.
- User interface to the microcontroller via a USART interface (RS232), which allows the user to configure the following parameters (see section 3.4.9):
    - *StorConnTest2 control:*
        - XFP power down
        - XFP transmit Enable
        - XFP I²C interface
    - *StorConn2 Control:*
        - VCSEL Driver TX enable
        - Optical Receiver RX Enable
        - Optical Receiver Squelch
        - TX and RX fault read-back
        - Advanced serial interface control to VCSEL Driver and Optical Receiver

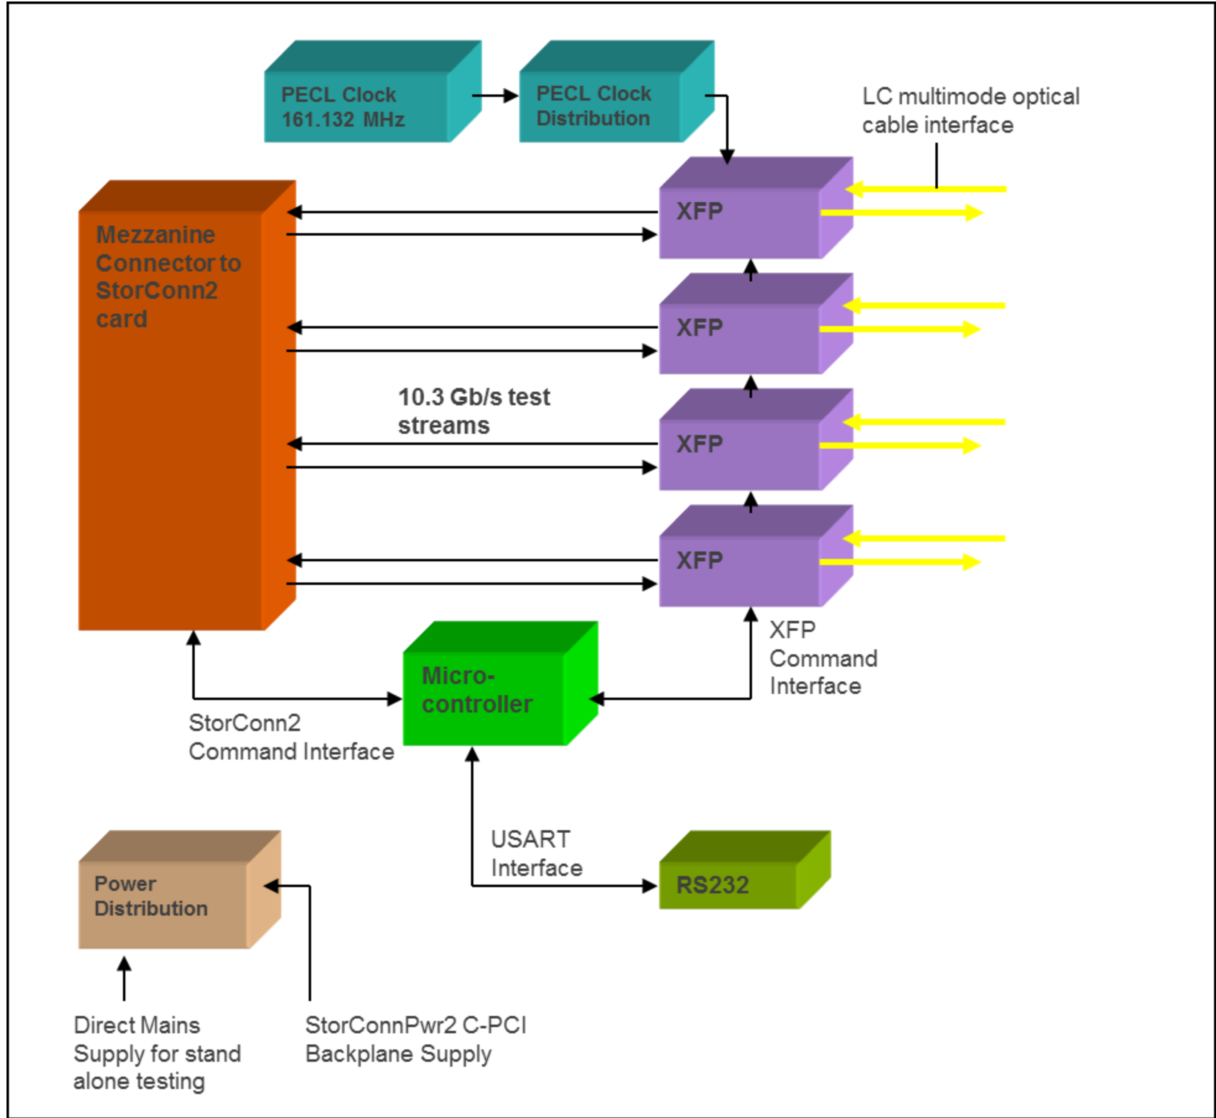


**Figure 3-39: StorConnTest2 functional diagram**

### *3.4.8.2. Circuit hardware design*

The StorConnTest2 circuit comprises a physical layer relay system between the four XFP ports and the high speed mezzanine connector to the adjoining StorConn2 device. This includes the controlling hardware for the StorConn2 and StorConnTest2 boards.

The challenge presented to the design of this circuit is the need to accommodate 10 Gb/s electronic data streams along differential copper traces over ample distances, while ensuring that the integrity of these high speed signals does not degrade beyond acceptable limits. This obstacle was tackled by adopting the appropriate PCB material, trace and via layout guidelines and signal conditioning hardware.

### *3.4.8.3. Raw card outline*

The StorConnTest2 raw card profile was chosen in accordance with Compact PCI standards for daughtercard form factors. Figure 3-40 shows the physical outline of the card and its principal dimensions.

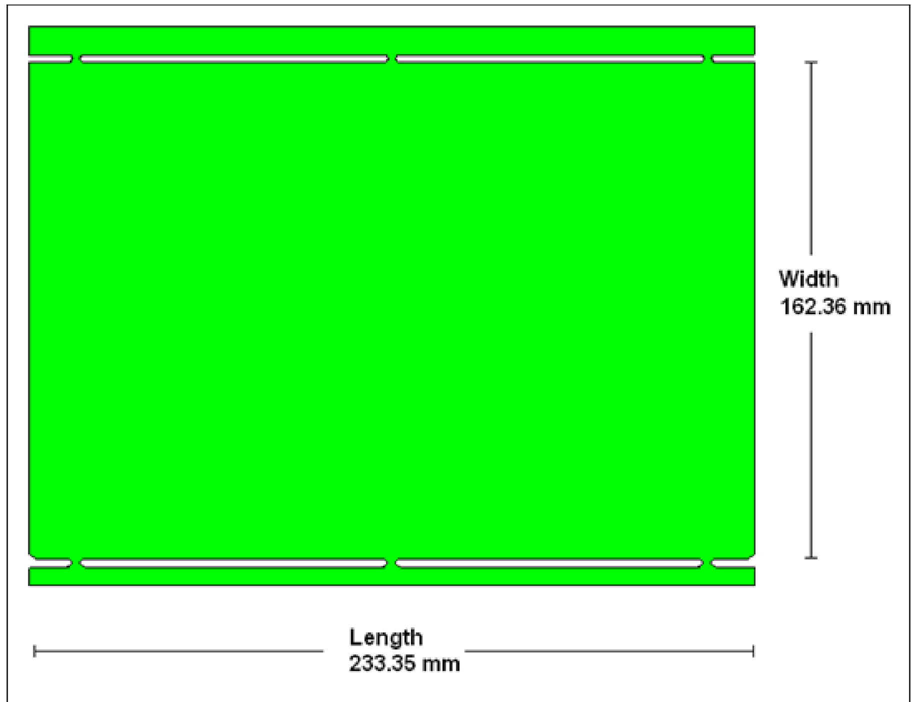


**Figure 3-40: StorConnTest2 card profile**

### *3.4.8.4. Raw card stack-up*

The circuit is housed on a standard six layer PCB, of which two layers, accommodating the high speed signal traces are composed of a Rogers dielectric material – RO4350B and the remaining layers accommodating low speed signals and power planes of a standard FR4 composite.

Controlled impedance tracks are located on the top layer (R_SIG00) and bottom layer (R_SIG01) subject to a differential impedance of 100 Ω ± 8%.

Figure 3-41 outlines the complete lay-up of this card. The layer names prefixed by “R_” denote the Rogers layers.

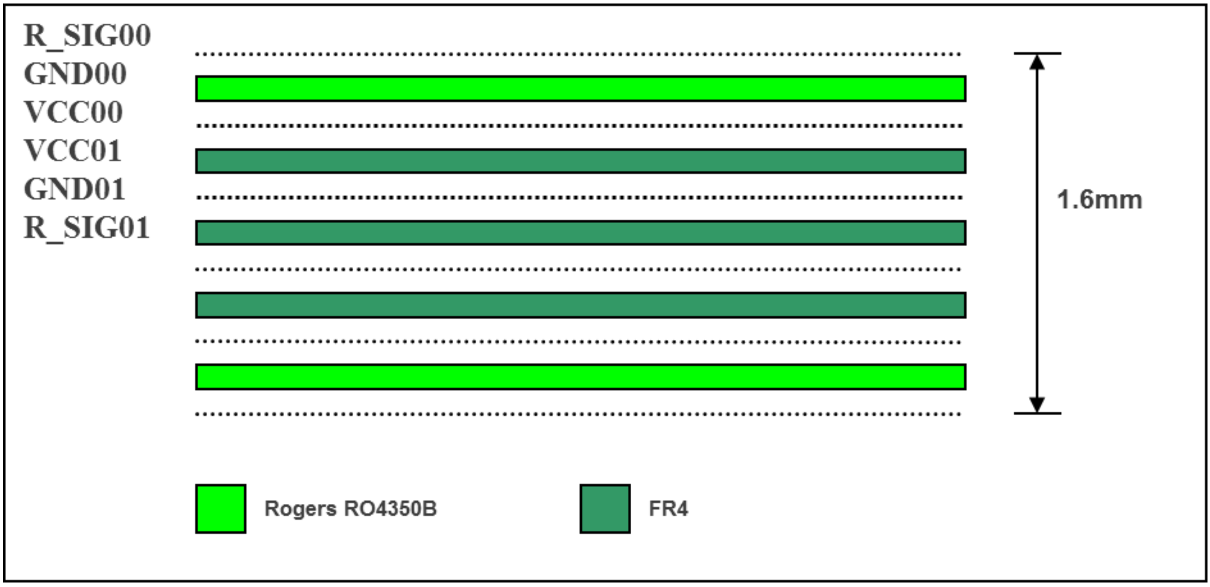


**Figure 3-41: StorConnTest2 raw card stack-up**

#### *3.4.8.5. StorConnTest2 component listing*

**Table 3-4** lists the main components of the StorConnTest2 card, while

**Figure 3-42**a shows their positions.

**Table 3-4: StorConnTest2 component listing**

| Component | Supplier | Part No. | Description |
|---|---|---|---|
| Mezzanine Connector | FCI Connect | 55737c | Copper array connector plug supporting differential signals up to and exceeding 10 Gb/s data rate and low speed control signals. |
| XFP transceiver | Picolight | PL-XXL-SC-S45 | 10 Gb/s optical transceiver with LC optical connector receptacles<br>850nm<br>multimode optical |
| XFP cages / connectors | Tyco | 788862-1 / 1489951-1 | Cages and connector bezels to accommodate XFPs on host board |
| Microcontroller | Microchip | PIC18F6620 | 64 pin microcontroller for user regulation of board functions |
| Signal Equaliser | Maxim | MAX3805 | 10.7 Gb/s Adaptive Receive Equaliser |
| PECL oscillator 161.132 MHz | | | Supplies clock signal for XFP CDR signal conditioning at 10.3 Gb/s |
| PECL differential clock driver | On Semiconductor | MC100EP14 | 1 : 5 PECL clock distribution buffer to relay single 161.132 MHz clock source to all four XFPs for CDR conditioning at 10.3 Gb/s |
| RJ-11 | Tyco | 520470-3 | Standard connector for programming and in-circuit debugging of microcontroller |
| DB9 | Farnell | | Standard serial connector for RS232 serial interface to microcontroller USART |
| Switching Regulator | Texas Instruments | PT6441N | 5V – 3V3<br>PIH 16-pin |

| | | | |
|---|---|---|---|
| | | PT6442N | 5V – 2V5<br>PIH 16-pin |
| Serial Interface | Maxim | MAX3232ECUP | RS232 Transceiver<br>TSSOP 20-pin |
| Linear Regulator | Micrel | MIC29502BT | 3V3 – 2V Low Voltage Dropout Regulator<br>SMD 6-pin |
| | | MIC39150-2.5BU | 3V3 – 2V5 Low Voltage Dropout Regulator<br>TO263-3 4-pin |
| 4 way right-angled connector | Molex | 53109-0410 | Auxiliary 4 way power connector to enable board to be powered outside the test enclosure |

#### *3.4.8.6. Power supply*

The card draws its supply from two alternative sources:

- The Molex power connector is used for stand-alone testing
- Either of the two C-PCI connectors will be used to supply power to the StorConnTest2 board when plugged into the electrical backplane – StorConnPwr2 – as part of the two card demonstrator assembly

#### *3.4.8.7. Connector positions and orientations*

- The C-PCI connectors are oriented such that upon rotation of the StorConnTest2 card by 180° around the axis of the adjoining StorConn2 card, the power connector orientation is identical. The purpose of this is to allow two StorConnTest2 daughtercards and more pertinently the two adjoining StorConn2 photonic interfaces to face each other when incorporated into the demonstrator assembly, without the need for separate circuits. This avoids the need for waveguide crossovers in the optical PCB, which will bind the two StorConn2 interfaces.
- The StorConn2 mezzanine connector is positioned centrally with respect to the C-PCI connectors in order to accommodate the requirement for rotational symmetry in the system.
- In addition, the StorConn2 mezzanine interface is orientated orthogonally to the axis of the StorConn2 card in order to help minimise the lengths of the high speed traces.

#### *3.4.8.8. High speed layout guidelines*

- The high speed differential copper signal traces are laid out across the top and bottom layers. These layers are composed of a Rogers dielectric material offering greater impedance control and a lower loss tangent at 10 Gb/s in comparison to FR4.
- The use of vias is unavoidable if crossovers on the high speed transmit and receive lines are to be prevented. In order to minimise unwanted reflections, the vias are direct and not stubbed.
- A ground via is located in the vicinity of each high speed via. The purpose of these "stitched" vias is to accommodate the return path of the signal at *all* points along the signal path.
- All aspects of the circuit layout have been adapted to minimise the lengths of the high speed traces.

#### *3.4.8.9. Microcontroller interfaces*

- The PIC microcontroller is programmed and debugged in circuit via a RJ-11 connector positioned on the front end of the card for ease of accessibility.
- User communication with the microcontroller by means of the embedded USART interface is maintained by a RS232 transceiver with a DB9 connector situated on the front end of the card for ease of accessibility.

#### *3.4.8.10. LED indicators*

- The card contains four LED banks to enable first-hand indication of the states of the XFPs' static control lines.
- LEDs are provided to indicate the status of every voltage rail on the card and the programming status of the microcontroller.

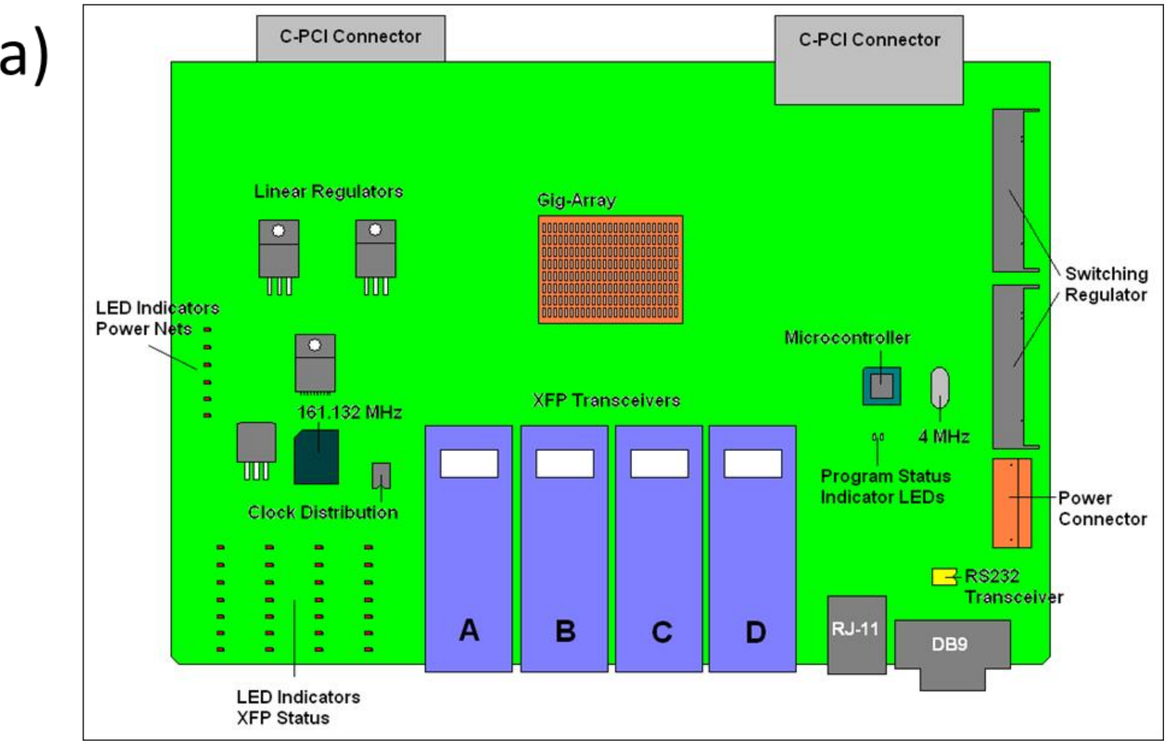


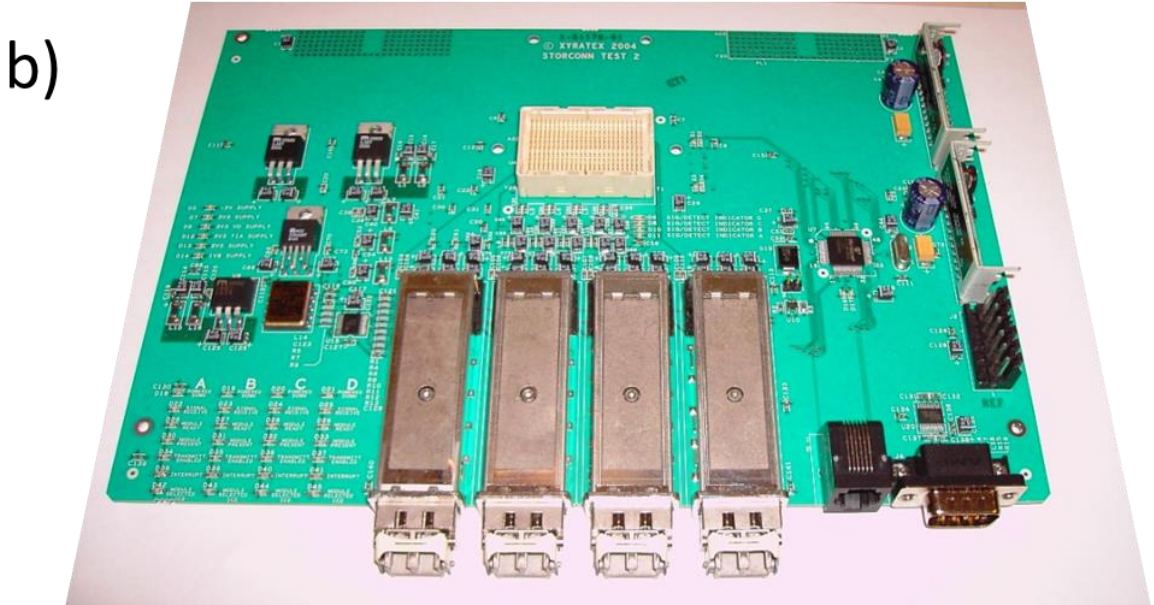


**Figure 3-42: a) Schematic showing StorConnTest2 component layout, b) photo of StorConnTest2**

### 3.4.9 Storlite platform firmware design

The StorConnTest2 microcontroller serves as the principal control mechanism on the unit, the main purpose of which is to both program and read-back data to and from the StorConn2 card and to regulate the XFP ports. The author designed the communication flow structure of the microcontroller and the firmware, which was written in Assembler code.

#### *3.4.9.1. Design structure*

The microcontroller flow design is partitioned into separate functional blocks in order to effectively accommodate the arbitration of user commands to different facilities and provide user feedback. This is shown in the flow diagram of Figure 3-43.

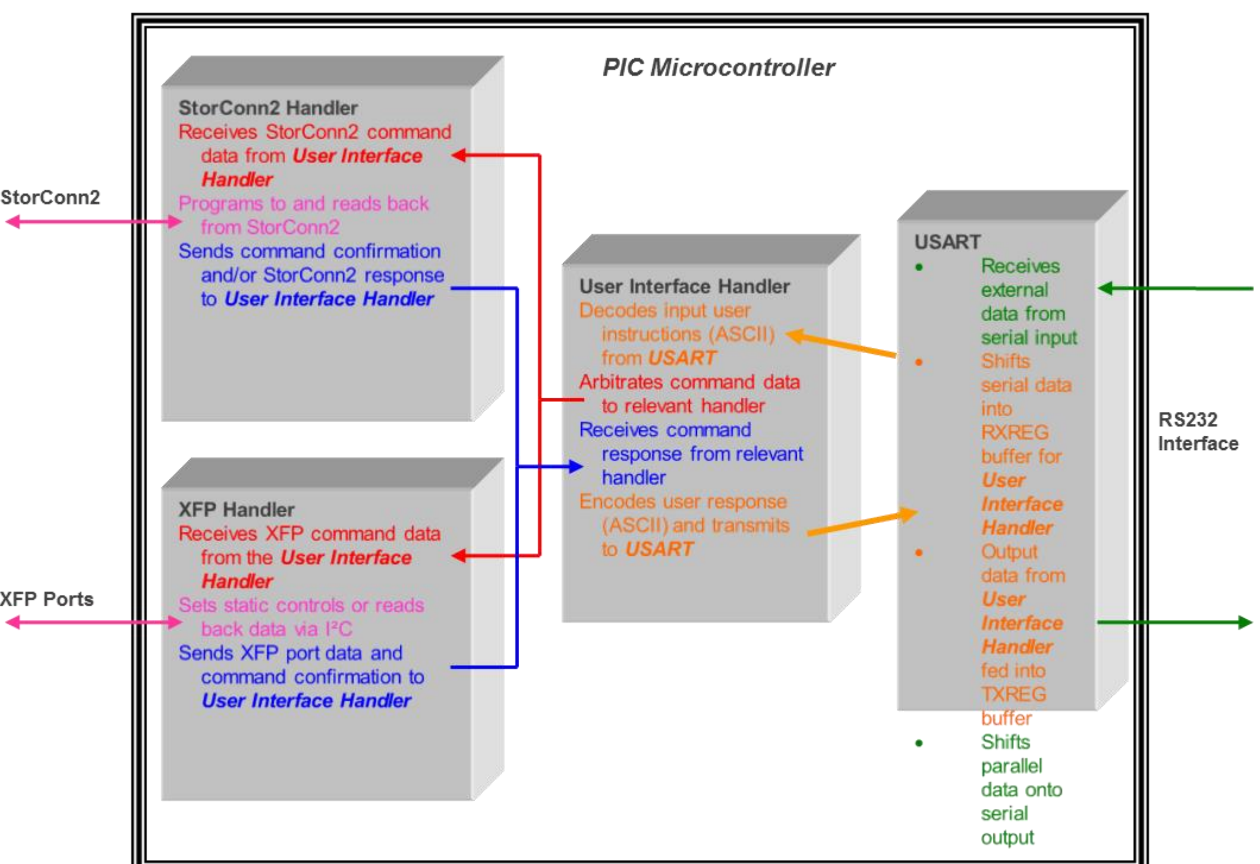


**Figure 3-43: StorConnTest2 microcontroller flow diagram**

#### *3.4.9.2. Command interface*

The command interface to the StorConnTest2 unit is regulated by the PIC microcontroller. The low level user interface has a rudimentary structure and is implemented by direct serial communication to the workstation via a null modem cable. An appropriate connection tool such as *HyperTerminal* allows user input to be transferred directly to the microcontroller USART via the serial connection and displays the data returned. The simple command structure underlies the development of the associated GUI described below to provide advanced user functionality.

#### *3.4.9.3. User command format*

Table 3-5 contains the complete command set for the StorConnTest2 unit.

The command format is structured in the following manner:

1) The command type is determined by the first letter.
2) The command body varies depending on which of the four commands is used.
3) The command is confirmed with “Enter”
4) If the command format is correct, the command will be executed and the appropriate response delivered, otherwise an error message will occur.

**Table 3-5: StorConnTest2 User Interface ASCII Command Format**

| ***Function*** | ***User Input*** | ***Response*** | ***Description*** |
|---|---|---|---|
| Static Control write | $\mathbf{wh_0h_1h_2h_3}$ (where h is a hex value) | **ok** | Set all static controls:<br>$\mathbf{h}_0$ = Enable StorConn2 Equalisers (3 to 0)<br>$\mathbf{h}_1$ = Enable XFP Tx Disable (D to A)<br>$\mathbf{h}_2$ = Power down XFPs (D to A)<br>$\mathbf{h}_3$ = $\mathbf{b_0\ b_1\ b_2\ b_3}$<br>$\mathbf{b_0}$ = Enable StorConn2 VCSEL driver<br>$\mathbf{b_1}$ = Enable StorConn2 optical receiver<br>$\mathbf{b_2}$ = Enable optical receiver squelch function<br>$\mathbf{b_3}$ = Assert microcontroller LEDs (diagnostic) |
| Static Control Read | **r** | $\mathbf{h_0h_1}$ | Reads all static control lines:<br>$\mathbf{h}_0$ = StorConn2 Equaliser Signal Detect (3 to 0)<br>$\mathbf{h}_1$ = $\mathbf{b_0\ b_1\ b_2\ b_3}$<br>$\mathbf{b_0}$ = StorConn2 optical receiver global signal detect<br>$\mathbf{b_1}$ = StorConn2 VCSEL Driver transmit fault<br>$\mathbf{b_2\ b_3}$ = Post-amble "00" (diagnostic) |
| XFP I²C Register Read | $\mathbf{yxh_0h_1}$ | $\mathbf{h_2h_3}$ | Reads back contents $\mathbf{h_2h_3}$ of channel **x** XFP at register address $\mathbf{h_0h_1}$ |
| VCSEL Driver temperature | **t** | $\mathbf{h_0h_1h_2h_3}$ | Reads back analogue to digital converted output of VCSEL driver temperature sensor |
| User Input Error | *Any input outside the command format* | **??** | |
| System Error | *Correctly formatted input* | ## | |

### *3.4.9.4. ASCII command exchange demonstration*

An example of the ASCII command use and format is shown below.

| **User Input** |
|---|
| **StorConnTest2 Output** |

**Table 3-6: StorConnTest2 Microcontroller ASCII Command Exchange**

| | |
|---|---|
| **Rdy** | *Upon reset the StorConnTest2 unit initialises user interface* |
| **wF00F** | • *Enable all StorConn2 Equalisers*<br>• *Enable transmit on all XFPs*<br>• *Power up all XFPs*<br>• *Enable StorConn2 VCSEL driver, optical receiver and optical receiver squelch function* |
| **ok** | *Operation confirmed* |
| **r** | *Read back static control lines* |
| **F8** | • *StorConn2 Equalisers all detecting signal from optical receiver*<br>• *Optical receiver verifies global signal detect*<br>• *No transmit faults on VCSEL driver* |
| **y100** | *Read back contents of register 00 on XFP A over the I²C interface* |
| **06** | *XFP reads back hex value "06" identifying itself as an XFP* |
| **ww** | *User input outside the specified format* |
| **??** | *User input error message* |
| **y100** | *Correct input* |
| **##** | *Error due to I²C interchange failure* |
| ***Carriage Return*** | *Carriage Return on its own* |
| ***No Response*** | *StorConnTest2 is forced into an idle state* |

### *3.4.9.5. Global User Interface (GUI) description*

The author designed the complete Global User Interface to control the Storlite demonstration platform using the Visual BASIC software language.

For the purposes of general user accessibility, the low level command interface is overlaid by the "**OptoPhy**" GUI, the main features of which are described below.

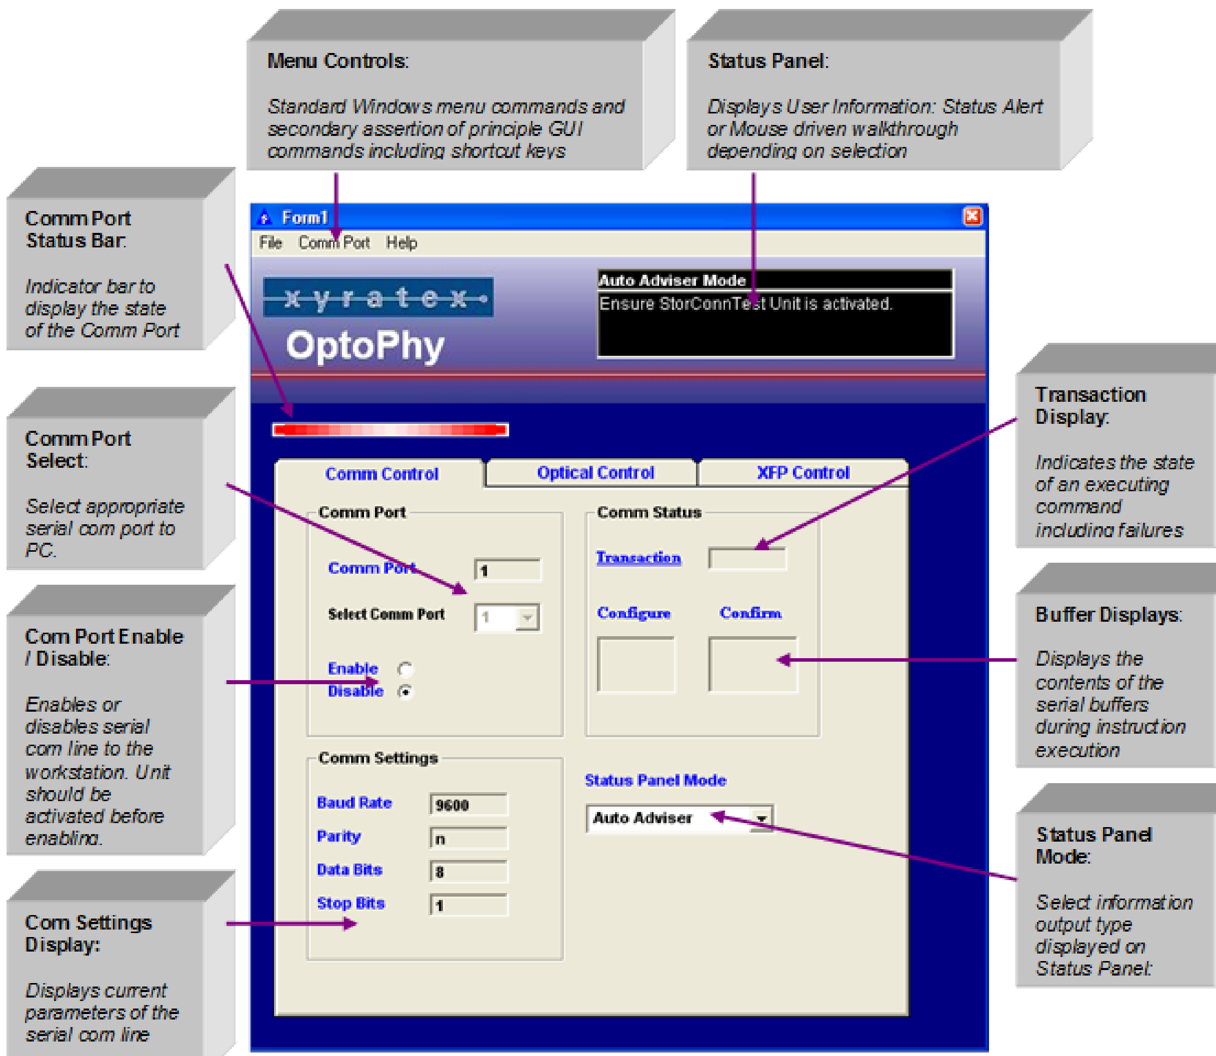


**Figure 3-44: Main screen / comm control tab**

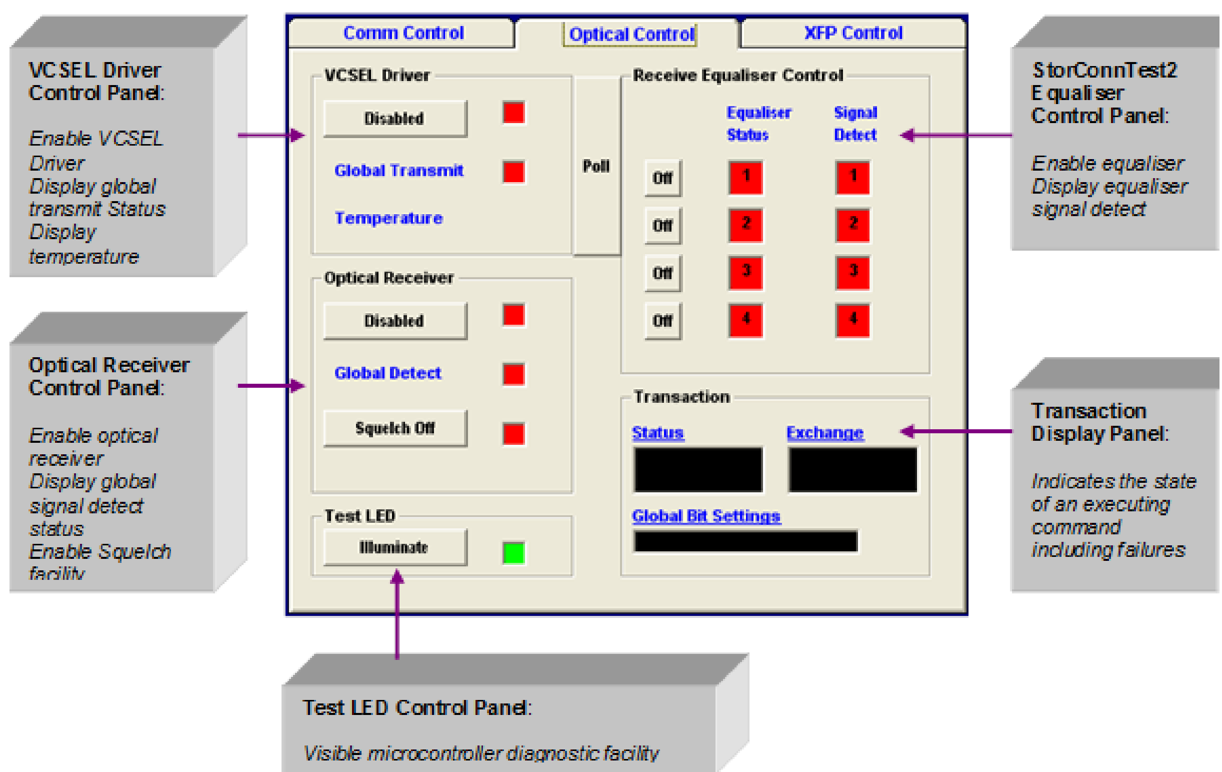


**Figure 3-45: Optical control tab**

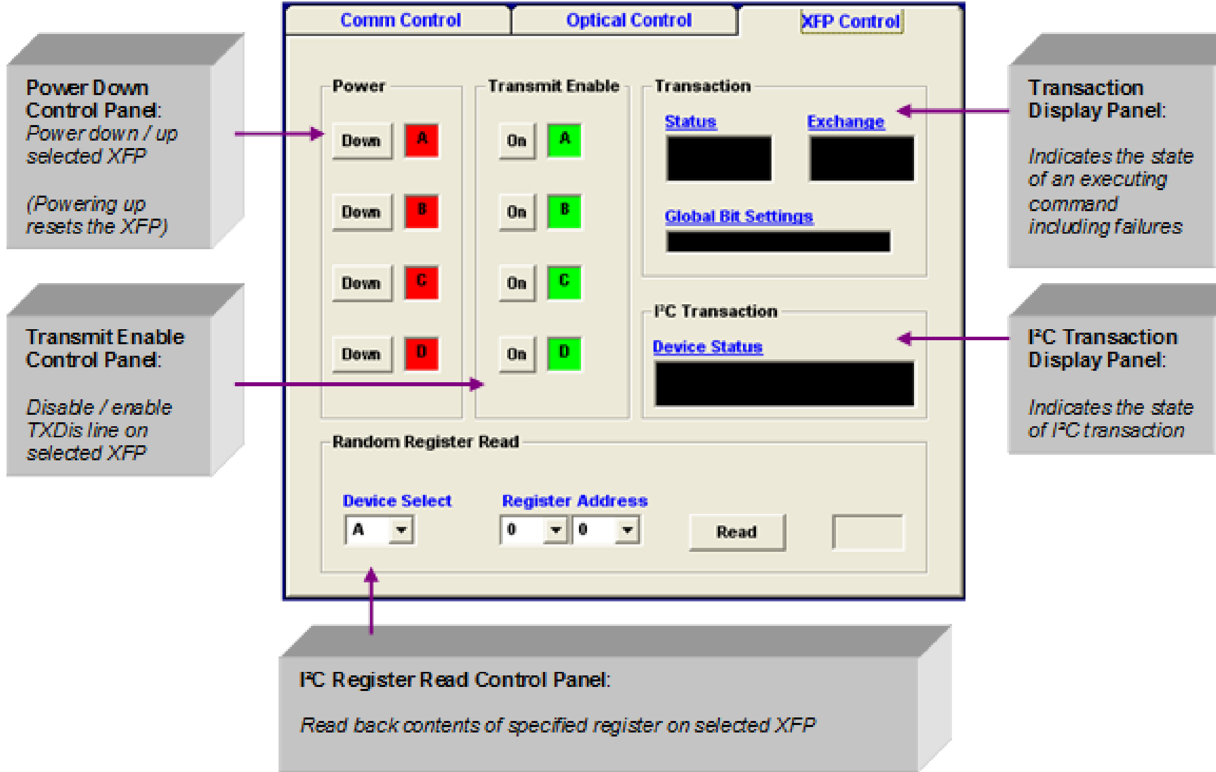


**Figure 3-46: XFP control tab**

## 3.4.10 StorConnPwr2 – passive electrical backplane

In addition to the optical backplane, *StorConnOpt2*, a separate passive electrical backplane was designed to supply power to the *StorConnTest2* daughtercards when they were in situ.

### *3.4.10.1. Physcial and component layout*

The *StorConnPwr2* card was populated with two C-PCI receptacles to accommodate the C_PCI power plugs on the *StorConnTest2* daughtercards and a five way header Power connector to accommodate an external power supply. Power filtering capacitors and indicator LEDs comprise the remaining components.

Table 3-7 shows the component listing.

**Table 3-7: StorConnPwr2 component listing**

| Component | Supplier | Part No. | Description |
|---|---|---|---|
| C-PCI Receptacle | ERNI | 923190 | Shielded electronic connector |
| Five way Header | Molex | | five pin power connector |

Figure 3-47 shows the component layout and physical outline of the *StorConnPwr2* card. Figure 3-48 shows the PCB layer stack-up of the *StorConnPwr2* card.

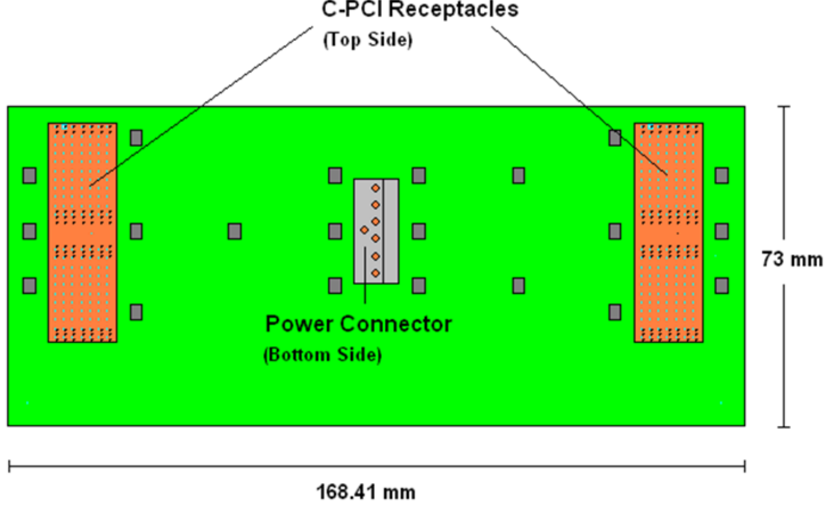


**Figure 3-47: StorConnPwr2 physical layout**

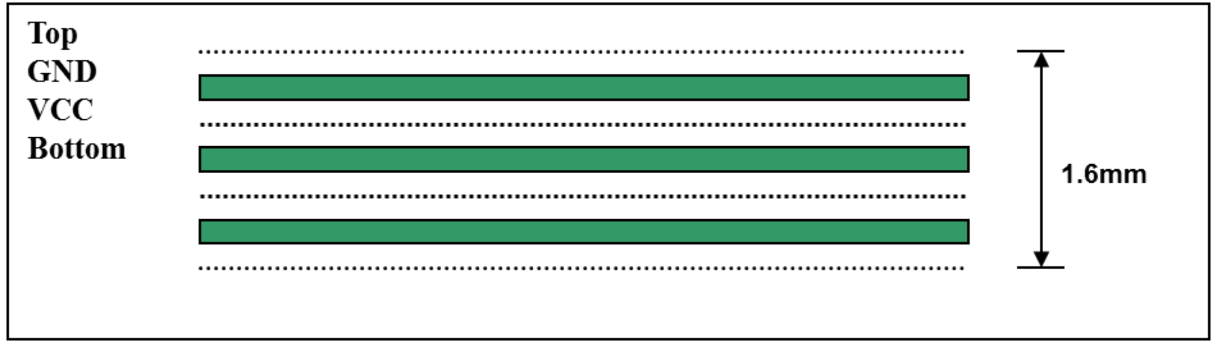


**Figure 3-48: StorConnPwr2 PCB layup**

### 3.4.11 Storlite demonstration platform

An assembly was constructed to demonstrate the proposed optical backplane connection system.

Figure 3-49 shows the iterative assembly of the Storlite demonstration platform starting from the aluminium enclosure (Figure 3-49a), the StorConnPwr2 backplane assembled (Figure 3-49b), one StorConTest2 daughtercard plugged into the StorConnPwr2 backplane (Figure 3-49c) and the completely populated enclosure (Figure 3-49d), which is elaborated in Figure 3-50.

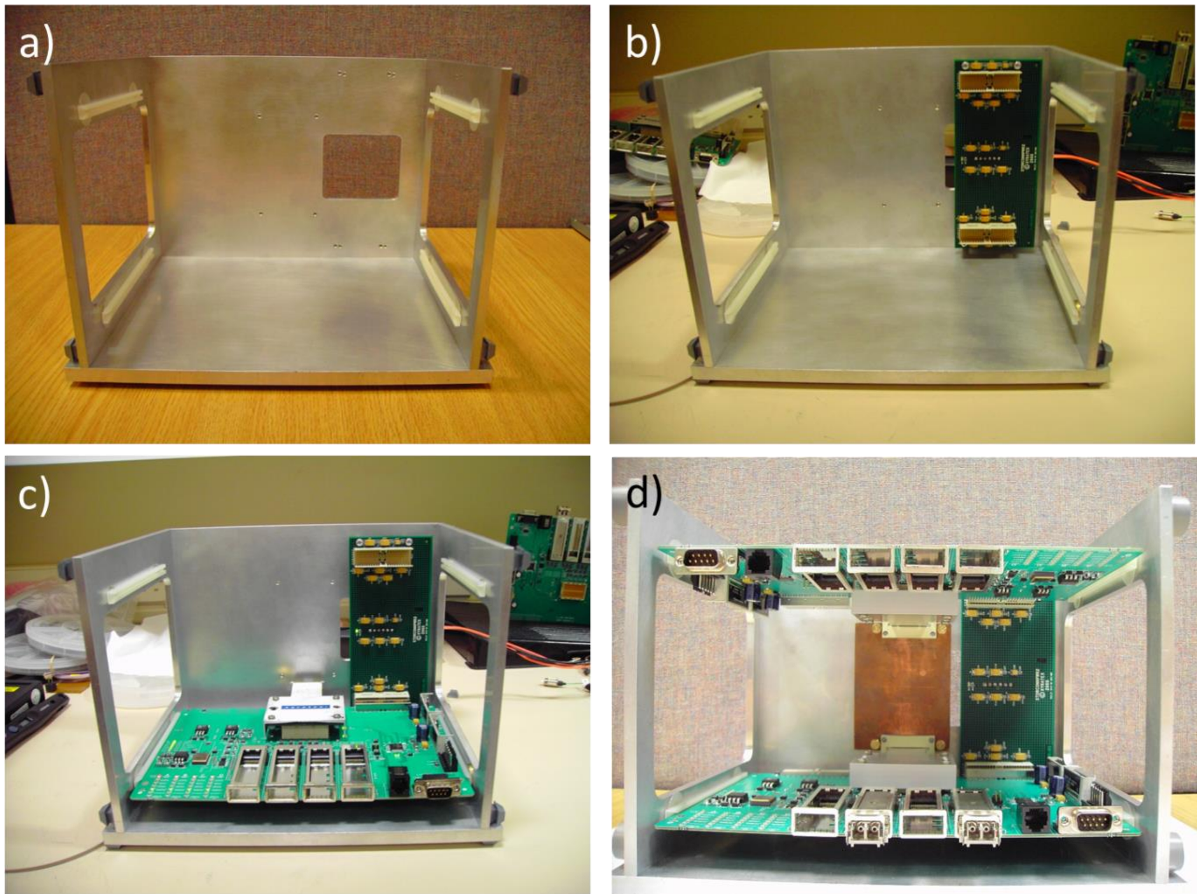


**Figure 3-49: Assembly process for Storlite demonstration enclosure: a) unpopulated enclosure, b) Enclosure populated with only StorConnPwr2, c) enclosure populated with StorConnPwr2 and StorConnTest2, d) fully populated enclosure**

The assembly was comprised of two test line-cards (Figure 3-50d), each housing the prototype Storlite active connector modules and the passive optical waveguide backplane (Figure 3-50a) described in section 3.4 with connector receptacles (Figure 3-50b) assembled at each end. In addition a separate electrical backplane (Figure 3-50c) was provided to supply power to both line cards.

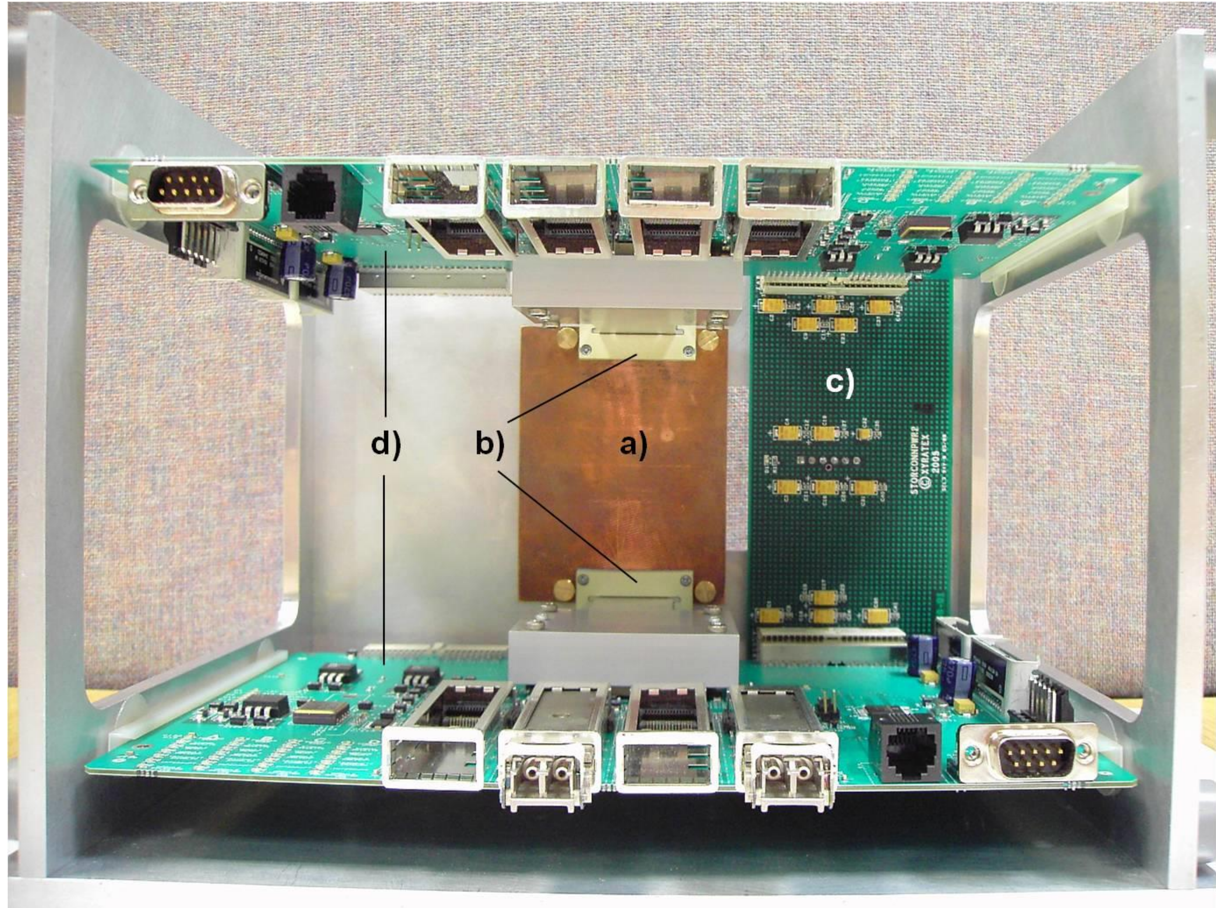


**Figure 3-50: Storlite demonstration platform a) Passive optical waveguide backplane (copper clad side shown), b) Pluggable optical backplane connectors, c) Electrical backplane, d) two test daughtercards**

The nature of the connection mechanism is such that the optical waveguides lie on the opposite side of the optical backplane to the side from which the connectors are plugged in. For this reason only the copper clad side of the backplane is visible when looking at the demonstration platform from the front, while the optical layer is visible when looking at the demonstration platform from the back as shown in Figure 3-51a. Figure 3-51b shows a close-up of a single waveguide on the Storlite backplane illuminated with visible light, while the backplane is in-situ in the demonstration chassis.

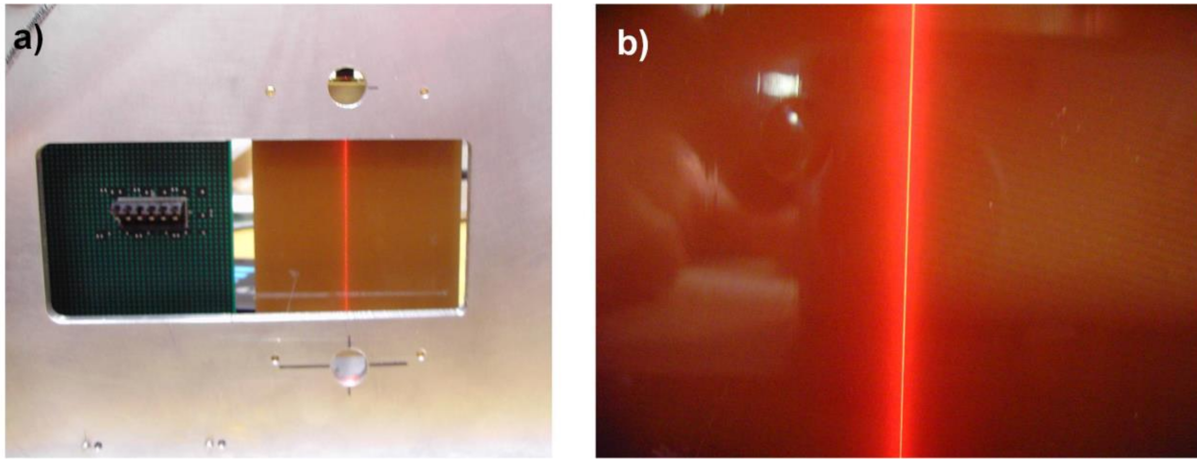


**Figure 3-51: a) View of Storlite demonstration system from the back, b) Single waveguide illuminated with 635 nm visible light**

## 3.4.12 Test and measurement

### *3.4.12.1. Xyratex 10 GbE LAN analyser*

In order to test the system for 10 GbE LAN applications, an in-house software tool developed by Xyratex was used to both generate a 10 GbE LAN test traffic stream and capture the received data and count the errors. Figure 3-52 shows a screen shot of the tool in operation after a 239 minute soak test with no errors recorded over the test period.

10GbE Test Card(12):0
ADDR BUFFER CONTENTS
0000 01234567 89ab1122 33445566 98badcfe 08070605 04030201 01020304 05060708
0020 efcdab89 67452301 10325476 98badcfe 08070605 04030201 01020304 05060708
DATA GOOD
Loop = 2879 Test Time = 239 Mins Time Left = 0 Mins FPGA Temp = 41 DegC
Elapsed Time = 239 Mins XFP Temp = 36 DegC
CAPTURE STARTED
CAPTURE STOPPED
Structure = StatisticsEther10GigType1 ID = 0x60280 Size = 0x280
STATS COUNTERS ACTUAL EXPECTED
Framer Received Registers
Frames Received OK : 4210237586
Octets Received OK : 1791648363024
Frames Received : 4210237586
Octets Received : 1791648363024
Unicast Frames Received OK : 0
Multicast Frames Received OK : 4210237586
Broadcast Frames Received OK : 0
Tagged Frames Received OK : 0
Pause MACControlFrame Received : 0

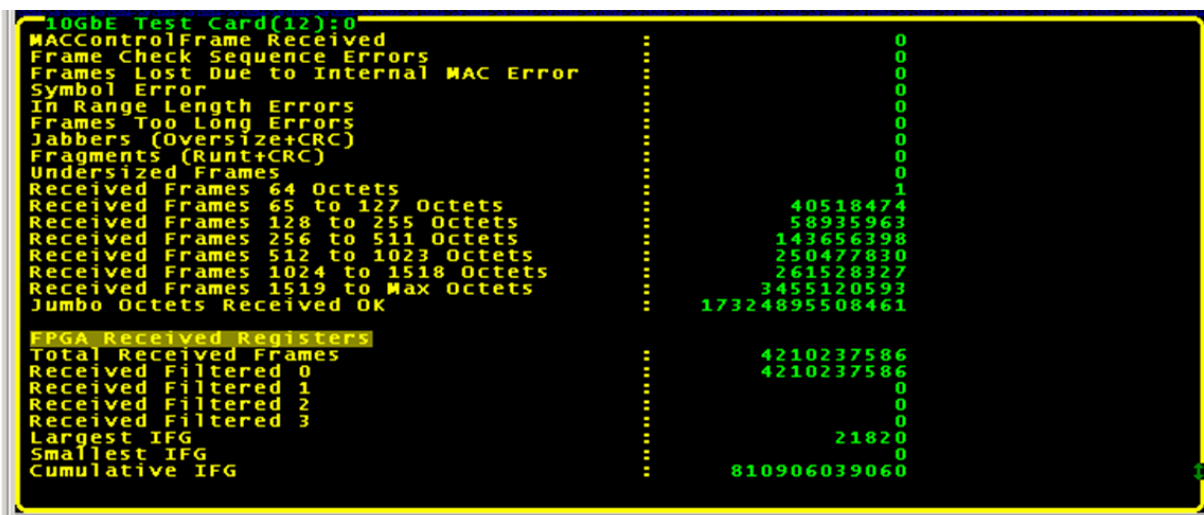


**Figure 3-52: Screen shots from Xyratex proprietary 10GbE LAN test traffic generator and error counter showing zero errors**

### *3.4.12.2. Bit error rate estimation from error counts*

In the absence of a bit error rate tester, the bit error rate of a signal can be estimated from the error count to a statistical confidence level of 99% by applying the Binomial Distribution Function as described by Redd [149].

(1 has been adapted from [149] to describe the number of error free bits (n) that need to be received in order to estimate the bit error rate (BER) of the signal with a given confidence level (CL).

CL can take on any value between 0 and 1.

$$n = -\frac{ln(1 - CL)}{BER} \qquad \textbf{(1)}$$

Given a required confidence level of 99% (0.99), the number of error free bits to be received is given by (2.

$$n = \frac{4.605}{BER} \qquad \textbf{(2)}$$

Equation (3 gives the minimum time t for which a signal of data rate (db/ds) can be measured and over which no errors should occur in order to estimate a bit error rate (BER) with a confidence level of 99%.

$$t = \frac{4.605}{BER(db/ds)} \qquad \textbf{(3)}$$

**Table 3-8: Test conditions to estimate BER on a 10 Gb/s signal from error counts with a 99% confidence level**

| **Bit error rate** | **Number of error free bits to be transferred** | **Required error free length of test time / hours** |
|---|---|---|
| $10^{-12}$ | $4.605 \times 10^{12}$ | 0.13 |
| $10^{-13}$ | $4.605 \times 10^{13}$ | 1.28 |
| $10^{-14}$ | $4.605 \times 10^{14}$ | 12.79 |
| $10^{-15}$ | $4.605 \times 10^{15}$ | 127.92 |

### *3.4.12.3. Parallel optical interface evaluation*

As mentioned, the purpose of the MT pins protruding through the GRIN lens ceramic support structure on the StorConn2 transceiver was to engage with a compliant receptacle on the optical waveguide backplane and draw the image points of the GRIN lens array into precise alignment with the waveguide end facets. Furthermore, this also allowed for a standard 6x12 MT ferrule terminated fibre jumper, of which the lowest row was populated to match the required offset, to be attached to the StorConn2 optical interface (Figure 3-53a), thus supporting stand-alone testability of the photonic interface itself using MT fibre-optic patch-cords (Figure 3-53b).

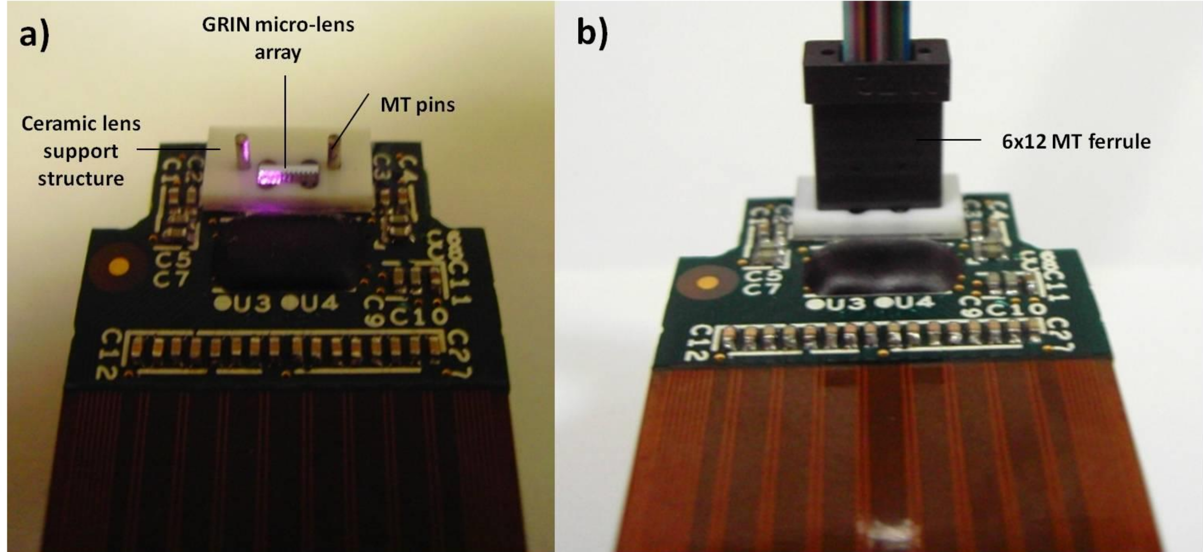


**Figure 3-53: a) Storlite transceiver opto-mechanical interface with all VCSELs activated, b) Storlite transceiver opto-mechanical interface with MT patchcord attached for stand-alone testing**

The StorConnTest2 test daughtercard was powered outside the chassis through the auxiliary power connector rather than via the CompactPCI connector. An external 10 GbE LAN test data pattern was generated by the Xyratex 10 GbE LAN analyser at a data rate of 10.3 Gb/s was conveyed to the transceiver transmit ports via the commercial XFPs on the front end of the test line card (Figure 3-54). The optical interface was mated to an MT patchcord with 62.5µm fibre fan-out, allowing direct characterisation of the transmitted output from each channel in terms of jitter and output power on a Tektronix CSA8000B communications signal analyser.

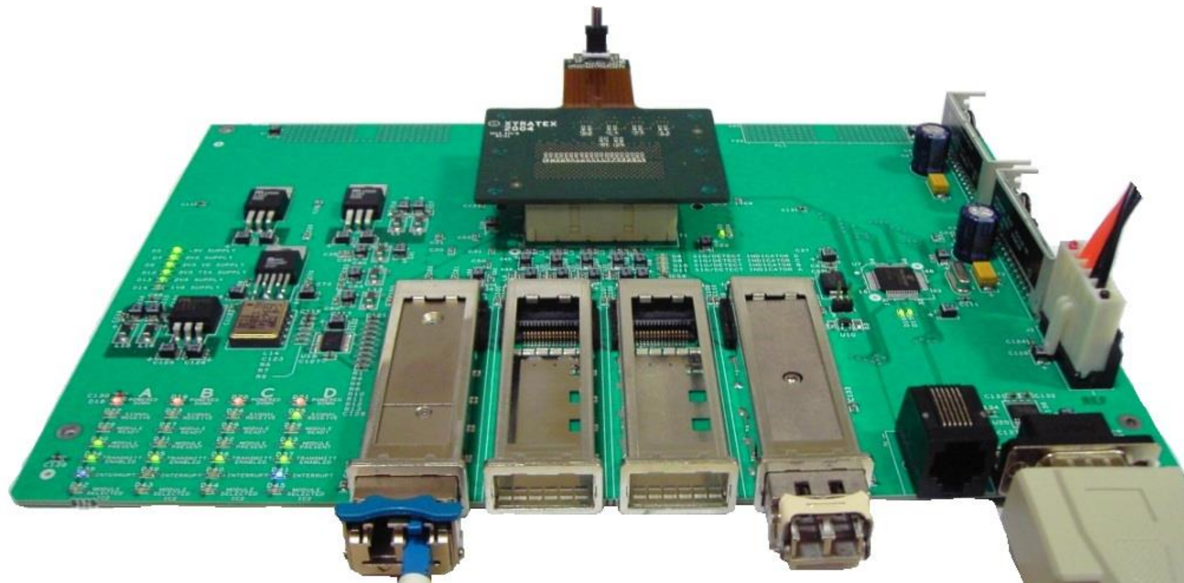


**Figure 3-54: Storlite test line card (powered) with transceiver attached (without connector housing) for stand-alone evaluation of the transceiver with an MT patchcord**

The VCSELs were driven with bias and peak-to-peak modulation currents of 11.91 mA and 9.8 mA respectively. The measurements were consistent showing an average jitter of 31.2 ps (0.32 UI, Unit Interval expressed as a fraction of the bit period), which is within the jitter thresholds for both transmitter input (0.61 UI) and receiver output (0.363 UI) as specified by the XFP MSA [150]. In addition the VCSELs exhibited an average optical power of 0.43 mW.

### *3.4.12.4. Eye diagram capture*

An eye diagram is a very useful means of evaluating the integrity of a digital signal. It is a superposition of multiple traces of a digital signal as would be recorded on an oscilloscope, thus over a sufficient amount of time all transitions (0-0, 0-1, 1-0 and 1-1) will be shown, giving rise to a pattern that resembles an eye.

A “clean” signal will exhibit little deviation from the ideal trigger point derived from the exact data rate of the signal, thus the superimposed signal traces will lie cleanly over each other, giving rise to an “open eye” with thin composite trace lines. A poor signal with high jitter, however, will deviate from the ideal trigger point, causing the superimposed lines to be staggered with respect to each other, and therefore the composite trace lines, the “sides of the eye”, will appear to become thicker. In the case of a very poor signal, the lines will be so thick that the central clearance will disappear entirely, giving rise to a “closed eye”. Analysis of an eye diagram can be used to extract a wide variety of signal characteristics including peak-to-peak jitter, extinction ratio, rise time and fall time. Each eye diagram shows the time per horizontal unit in the bottom right hand corner and the full voltage amplitude of the signal in the top right hand corner. The time per division is determined by the data rate. The full voltage varies depending on the strength of the received optical signal.

The signal eye diagrams of all four channels are recorded in Figure 3-55. Each eye diagram was measured with the three other sources off.

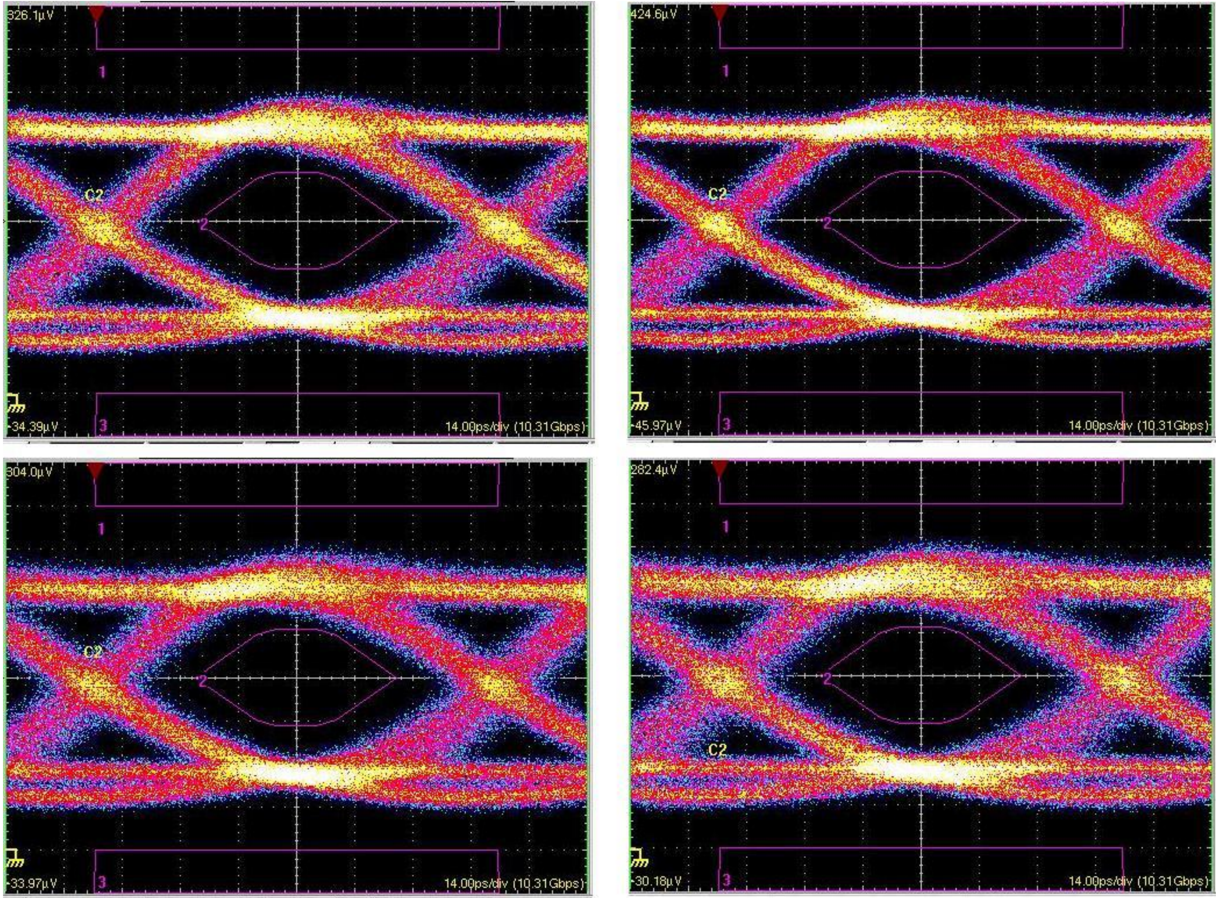


**Figure 3-55: Eye diagrams of all four VCSEL channels driven at 10.3 Gb/s with a bias current of 11.91 mA and a modulation current of 9.8 mA. Each VCSEL was measured with all other VCSELs off**

#### *3.4.12.5. Bit error rate performance estimate*

A loop-back test was carried out whereby the MT fibre patchcord attached to the Storlite transceiver optical interface was configured such that the four fibres coupled to the four VCSEL ports were looped back to the four fibres coupled to the four PIN photodiode ports. The performance of the four bidirectional optical links was characterised using the Xyratex 10 GbE LAN analyser, which generated a 10 GbE LAN test pattern and provided an error counter on the received stream. The 10 GbE LAN test data pattern was conveyed directly from the VCSEL under test to the PIN photodiode coupled to it. The signal was then reconverted into a high speed differential signal, passed to another XFP on the front end and the output of that sent back to the 10GbE LAN network analyser. A four hour soak test showed unimpaired traffic throughput in the loop-back configuration described. Based on the testing times indicated in Table 3-8, each link was measured for over two hours with no

errors recorded thus demonstrating a bit error rate of at least $10^{-13}$ at a confidence level of 99%. In accordance with 10 GbE LAN specifications, the BER limit of $10^{-12}$ was therefore satisfied throughout the test cycle.

#### *3.4.12.6. Optical waveguide signal integrity characterisation*

A group of 12 waveguides was directly characterised by attaching two MT patchcords with 62.5µm fibre fan-out to the compliant receptacles at each end of the waveguide group. Optical test data generated by the proprietary Xyratex 10 GbE LAN network analyser with a wavelength of 850 nm and a data rate 10.3 Gb/s was launched into the waveguide under test through the first connected MT patchcord and extracted from the waveguide by the second MT patchcord. The signal was then passed to the communications signal analyser. Each of the 12 waveguides was characterised in this fashion. The waveguide end facets had been diced but not polished.

The characterisation activity was repeated with the ingress and egress waveguide end facets untreated (Figure 3-56a) and treated (Figure 3-56b) with isopropanol, which served as an index damping fluid. The refractive index of isopropanol at 850 nm is 1.3776. As the refractive index of the core polymer is 1.56, this does not serve as an index matching fluid, however treating the interfaces with this fluid significantly reduces the index difference between the core and the material filling the gap between the waveguide end facet and the MT ferrule fibre. Therefore the end facet scattering losses are also reduced. It was decided to opt for the alcohol-based isopropanol rather than a refractive index matching oil as the former would be more easily removable, while the latter would have been difficult to completely remove making repeated measurements less reliable. Each eye diagram was measured with all other sources off. The horizontal time scale is 14 ps / horizontal division, which is consistent with the representation of a 10.3 Gb/s signal. The vertical voltage scale adjusts itself to the amplitude of the received signal and thus varies depending on the strength of the received optical signal.

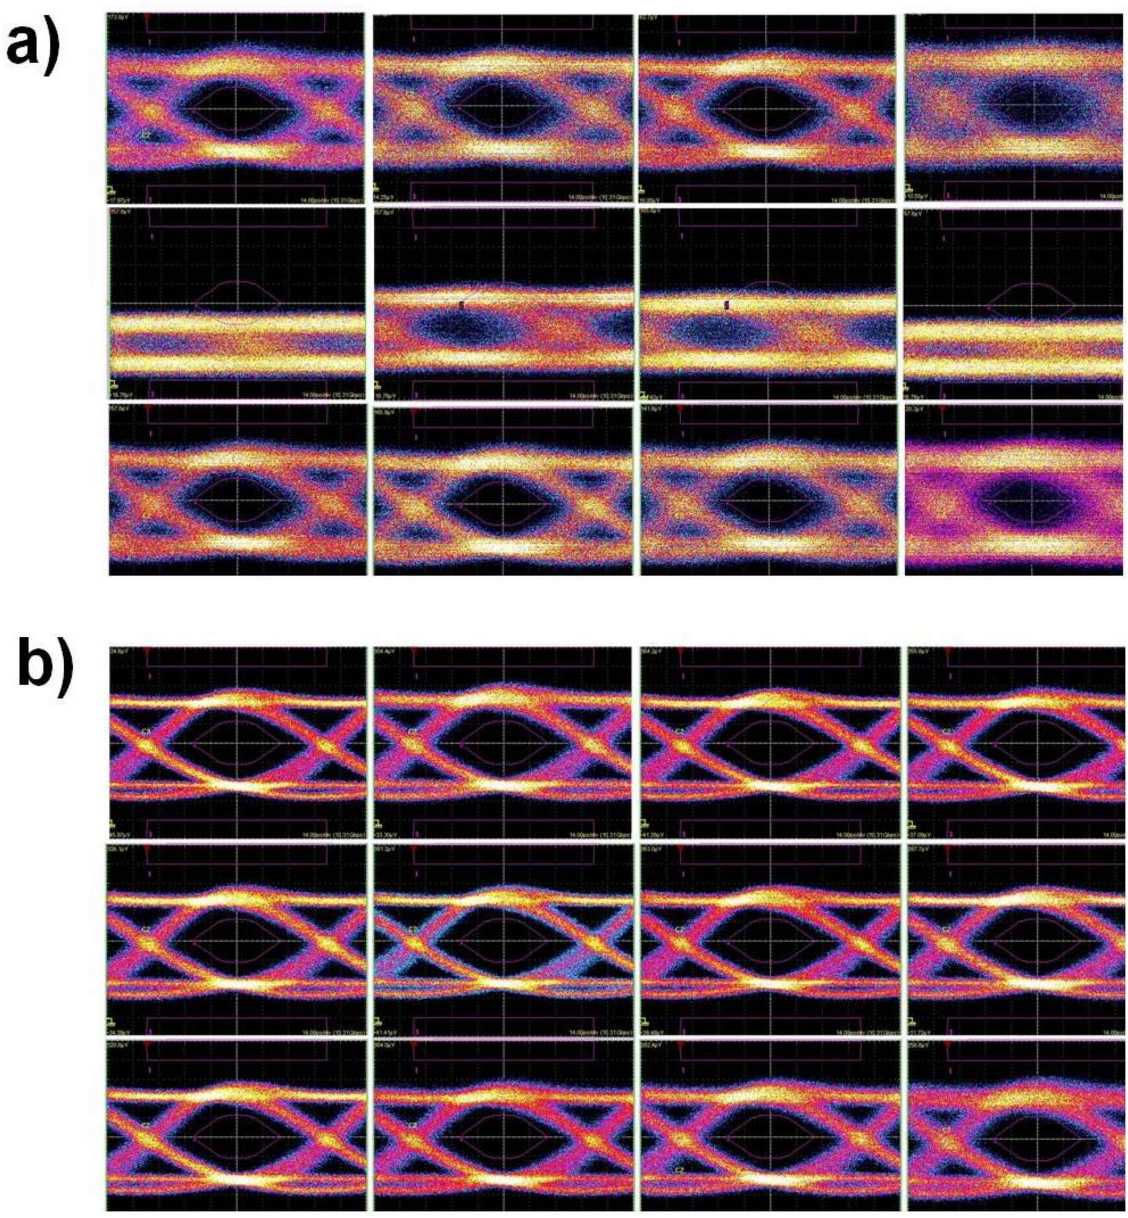


**Figure 3-56: Eye diagrams of 10.3 Gb/s optical signals after propagation through 12 waveguides under test. a) Eye diagrams captured without index damping fluid applied on waveguide end facets, b) Eye diagrams captured with index damping fluid applied on both ingress and egress waveguide end facets**

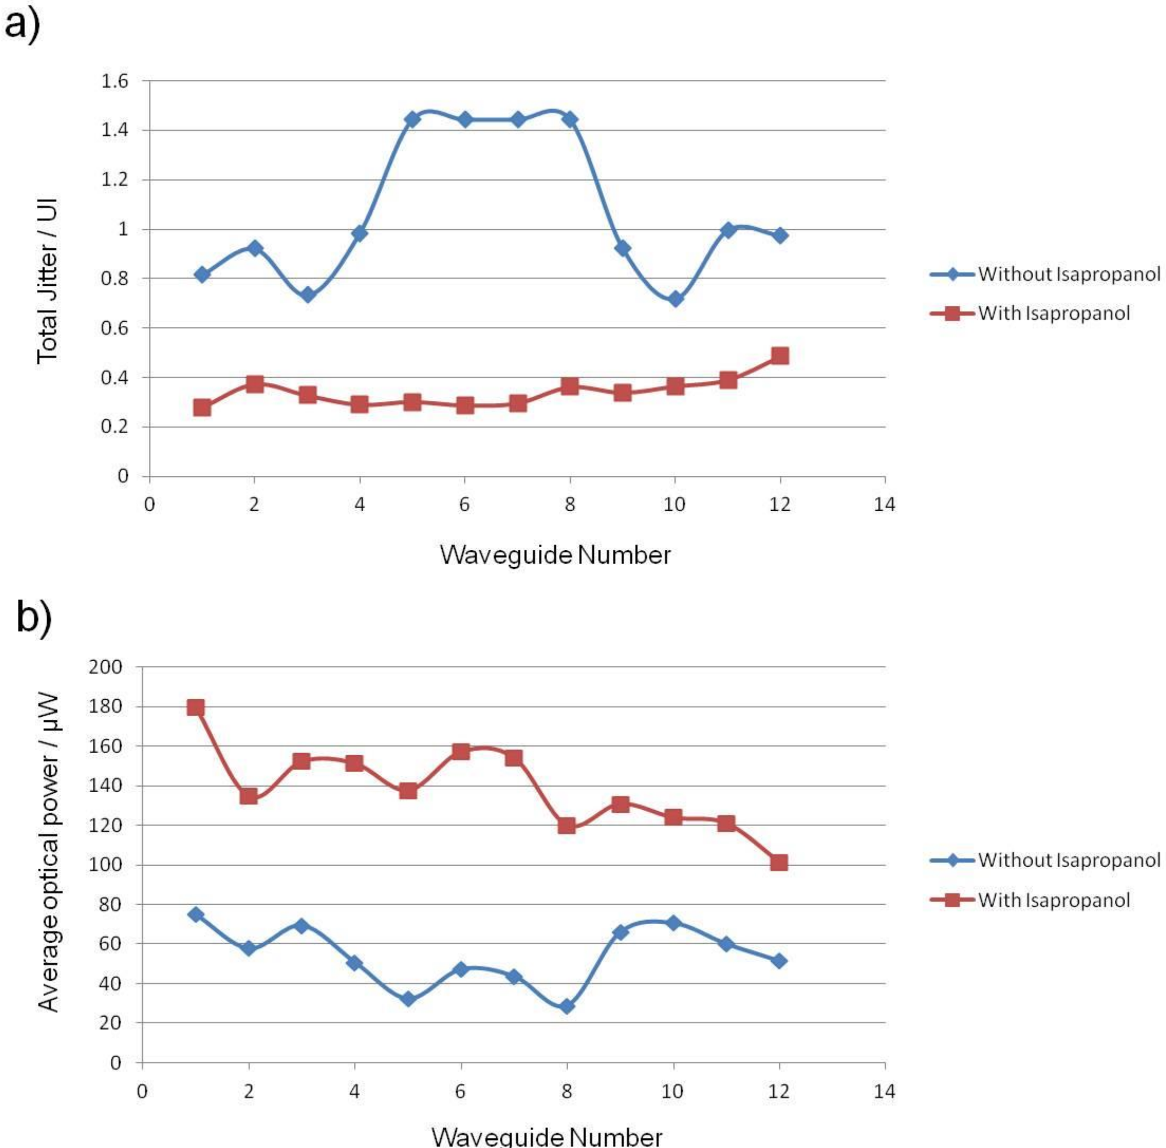


**Figure 3-57: a) Total jitter of 10.3 Gb/s signals measured across 12 waveguides without index damping fluid (blue graph) and with index damping fluid (red graph), b) Comparative average optical power measured at the communications signal analyser across 12 waveguides without index damping fluid (blue graph) and with index damping fluid (red graph)**

The application of index damping fluid on the ingress and egress waveguide end facets gave rise to average reduction in jitter of 65.82% and an average improvement in optical power coupling of 4.181 dB.

It is clear from the graphs that differences in end facet quality from waveguide to waveguide can have a substantial impact on signal integrity. As expected the application of index damping fluid mitigates the differences in end facet quality and in the case of Figure 3-57a transforms the strong variation in jitter to a more or less uniform profile. In Figure 3-57b, it can be seen that the optical power transmission profile across the 12

waveguides under test is more or less maintained, but simply shifted in magnitude denoting a larger attenuation on those channels, the end facets of which have not been treated with isopropanol compared to those with isopropanol treatment.

#### *3.4.12.7. Connector interface characterisation*

The Storlite transceiver connector was docked to the optical waveguide backplane and different VCSEL channels activated. Fig. 33a shows the egress waveguide interface with one VCSEL channel activated and one waveguide illuminated with 850 nm light. Fig. 33b shows the egress waveguide interface with two non-adjacent VCSEL channels activated and their corresponding waveguide channels illuminated.

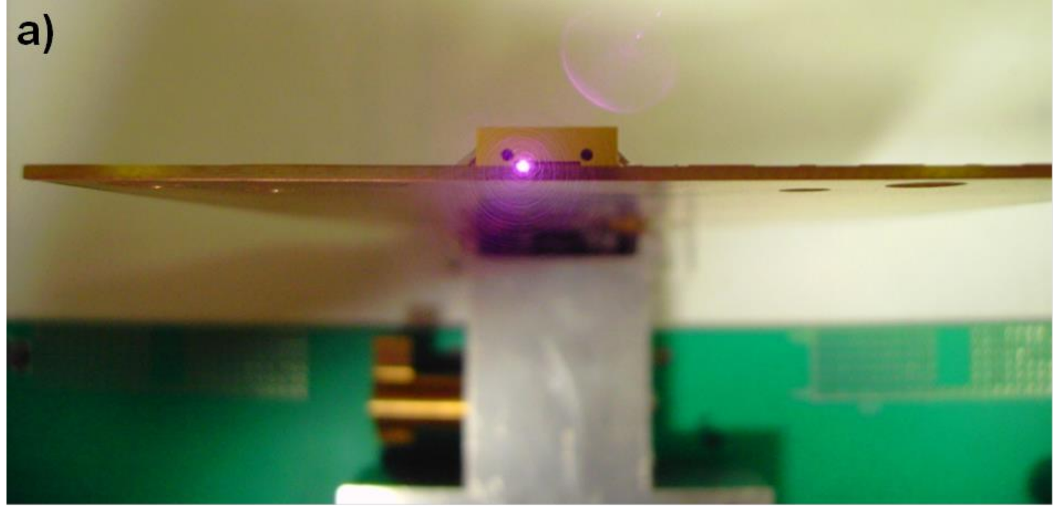


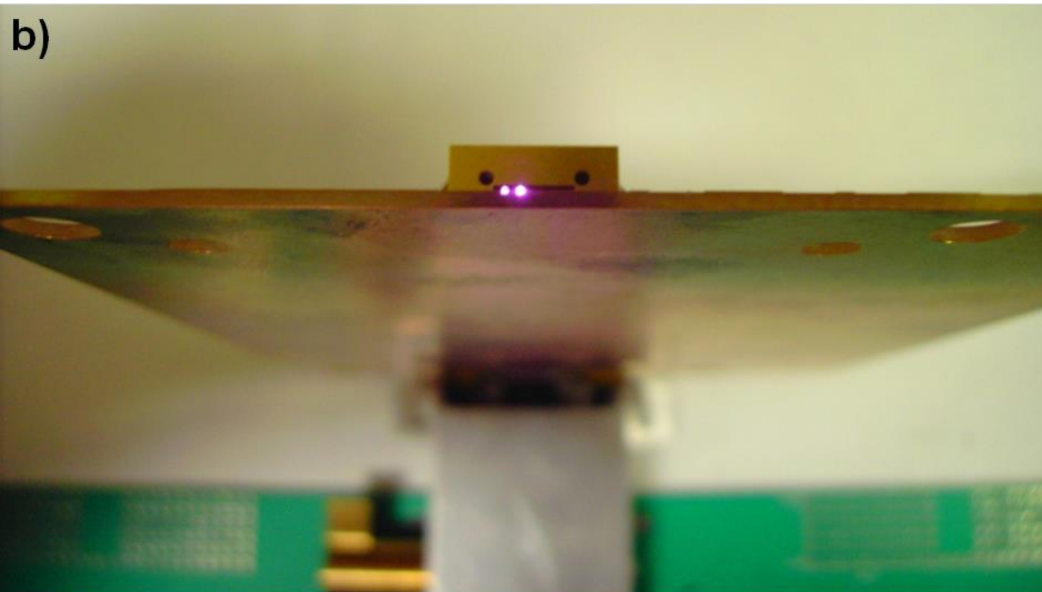


**Figure 3-58: View of exposed optical backplane waveguide egress interface with Storlite connector attached to ingress interface. a) VCSEL channel one (furthest left) activated and 850 nm light exiting single egress waveguide, b) VCSEL channels one and four activated and 850 nm light exiting two egress waveguides 750 µm apart**

The transceiver connector was docked to the optical PCB and an MT patchcord to the egress waveguides at the other end (Fig. 34). The extracted data-stream was characterised after passage through the butt-coupled connection and polymer waveguides.

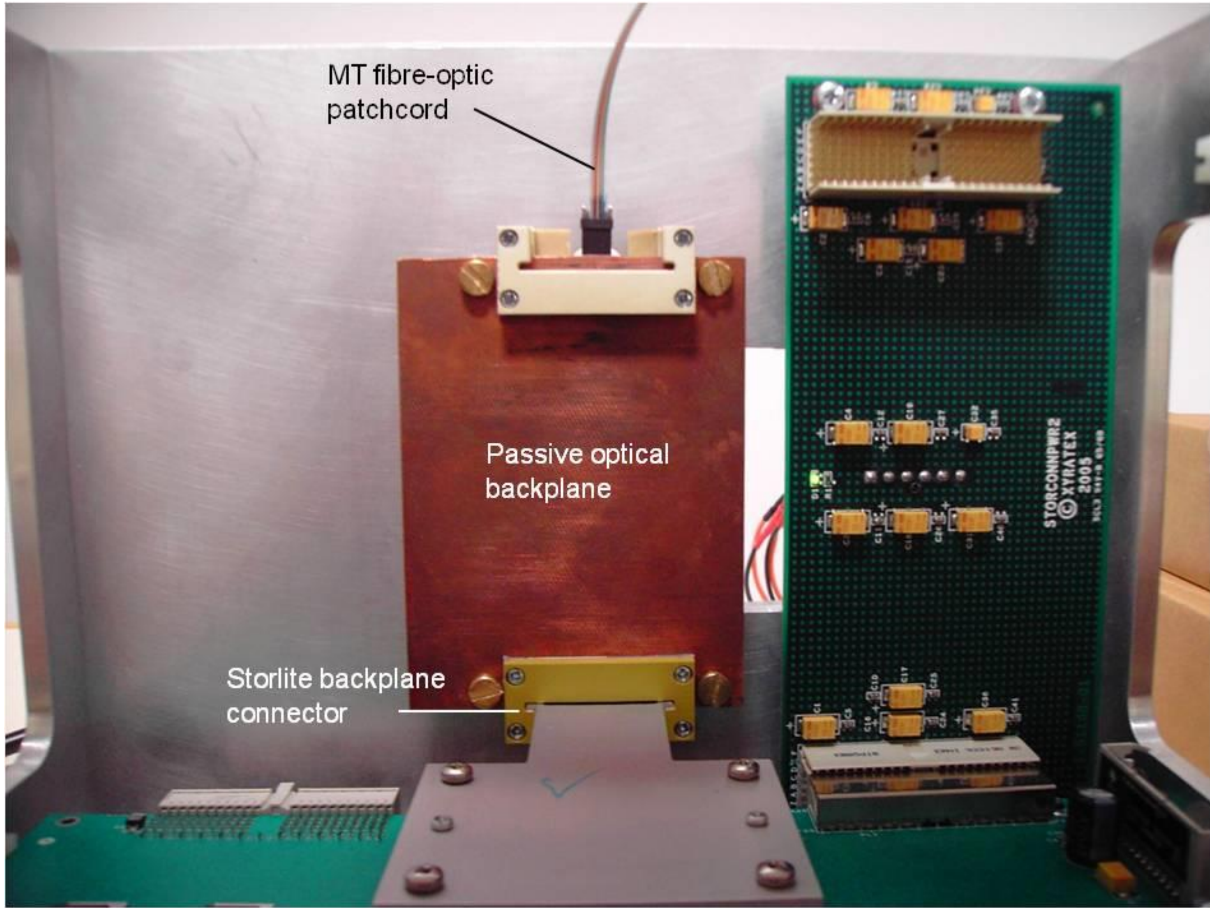


**Figure 3-59: Characterisation setup for Storlite optical interface connection to optical backplane**

The author carried out a comparative assessment on how the preparation of the optical PCB interface surface affects both interface scattering losses and high-speed signal integrity in terms of jitter. The author considered three waveguide end facet preparations: a) diced with isopropanol applied to the ingress and egress waveguide end facets (best case), b) diced and polished, c) diced only (worst case). The MT patchcord was attached directly to the optical transceiver interface and the output directly characterised to provide a reference against which these use cases could be evaluated. For each of the three end facet preparations, the total jitter (deterministic jitter + random jitter) and the optical insertion loss was measured on each of the four VCSEL and waveguide channels under test. The results in Table 3-9 show the average values of total jitter in UI and relative optical loss in dB compared to the reference signal. All measurements were made on a single waveguide, namely the second waveguide from the outer most waveguide.

**Table 3-9: Waveguide end facet characterisation**

| Case | Waveguide End Facet Preparation | Total Jitter / UI | Relative Loss / dB |
|---|---|---|---|
| One | No Waveguide (Reference) | 0.3389 | 0 |
| Two | Diced with Isopropanol applied to ingress and egress end facets | 0.3595 | 4.5 |
| Three | Diced and polished | 0.5593 | 6.9 |
| Four | Diced only | 0.8899 | 7.9 |

The total link margin is determined by the difference between the nominal emission power of the source and the minimum amount of optical power that the receiver can capture while still able to resolve the signal to an acceptable quality. The nominal emission power of the ULM VCSEL is 2 mW (3.01 dBm). An exact value for receiver sensitivity could not be provided by the manufacturer, therefore the value for receiver sensitivity used has been taken from commercial transceivers with similar components operable at the same wavelength. The receiver sensitivity of the PDs required to resolve a 10.3 Gb/s signal is therefore taken to be 0.15 mW (-8 dBm). The link margin is calculated as 3.01 dBm – (-8 dBm) = 11.01 dB. Therefore, based on the values in Table 3-9, the loss values on all waveguides fall within the link margin, therefore in principle each waveguide could sustain a 10.3 Gb/s optical signal.

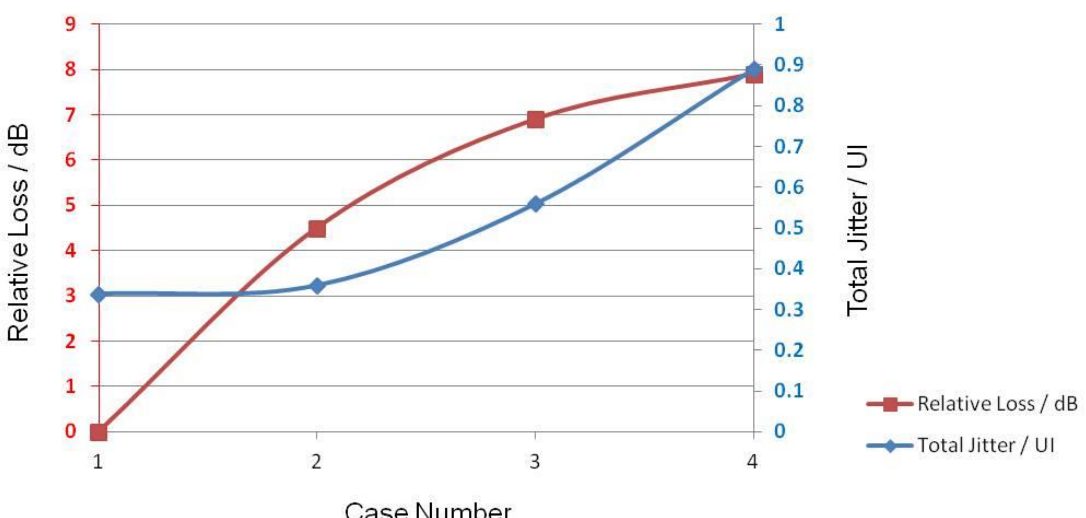


**Figure 3-60: Relative loss / dB and total jitter / UI for each use case listed in Table 3-9.**

The optical signal waveforms for the four cases outlined in Table 3-9 are shown in Figure 3-61.

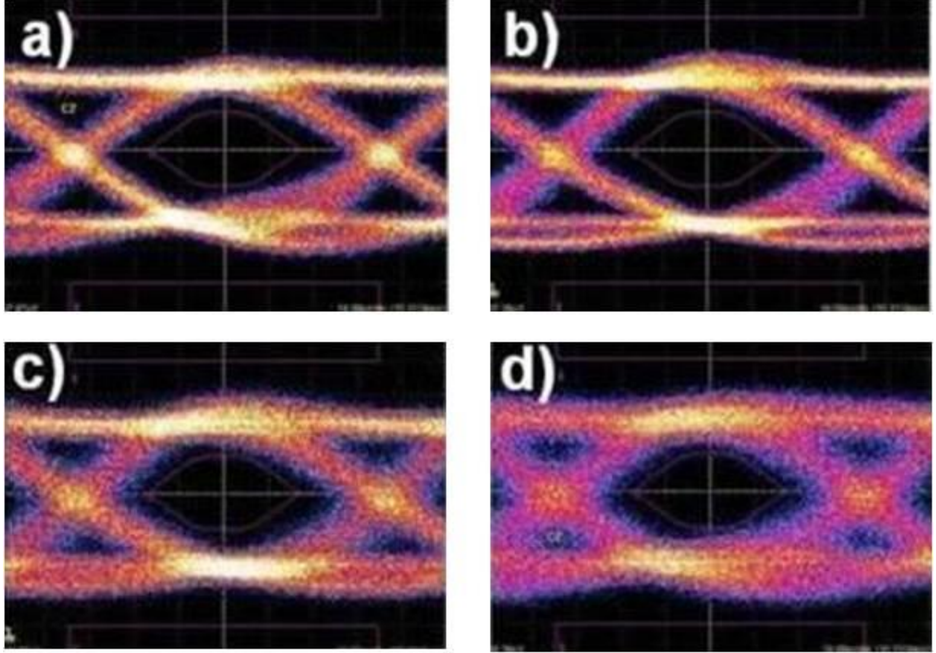


**Figure 3-61: Eye diagrams for four use cases: a) Storlite VCSEL is connected directly to communications signal analyser via MT fibre-optic patchcord (Reference), b) VCSEL connected to waveguide with isopropanol applied to ingress and egress end facets, c) VCSEL connected to waveguide with end facets diced and polished, d) VCSEL connected to waveguide with end facets only diced**

Figure 3-60 and the waveforms in Figure 3-61 show that both jitter and relative loss increase significantly as the surface quality of the waveguide end facets diminishes.

The impact of the surface quality of waveguide end facets on overall waveguide insertion loss will be treated in more detail in the next chapter with a novel and commercially viable solution proposed and characterised to reduce end facet scattering losses.

### *3.4.12.8. Full demonstrator loop-back test through optical backplane*

The demonstration platform was fully populated (Figure 3-62). The 10 GbE LAN test data pattern at a data rate of 10.3 Gb/s was generated by a proprietary Xyratex network analyser and conveyed along an optical fibre to a commercial XFP on the front end of the test line card. The XFP converted the optical signal into a differential electronic signal, which was routed to the dedicated Storlite transceiver VCSEL transmitter to which it was electrically wired on the line card PCB. The electronic test signal was then converted by the VCSEL driver and VCSEL on the first transceiver board into an optical signal, which was imaged into the waveguide ingress facet through the GRIN micro-lens in the transceiver interface. The optical signal propagated along the embedded

multimode waveguide to the GRIN micro-lens of the optical interface of the second transceiver attached to the other backplane receptacle and was imaged into the active area of the PIN photodiode. The resulting modulating electrical current was then converted by the TIA/LA back into a differential electronic signal. The electronic signal was routed along the second test card to an XFP on the front-end. The XFP converted the signal into an optical signal, which was conveyed by a multimode 62.5 µm graded index fibre to the communications signal analyser.

Repeated four hour soak test cycles consistently showed unimpaired traffic throughput in the loop-back configuration described. Based on the testing times indicated in Table 3-8, each link was measured for over two hours with no errors recorded thus demonstrating a bit error rate of at least $10^{-13}$ at a confidence level of 99%. In accordance with 10 GbE LAN specifications, the BER limit of $10^{-12}$ was satisfied throughout the test cycle.

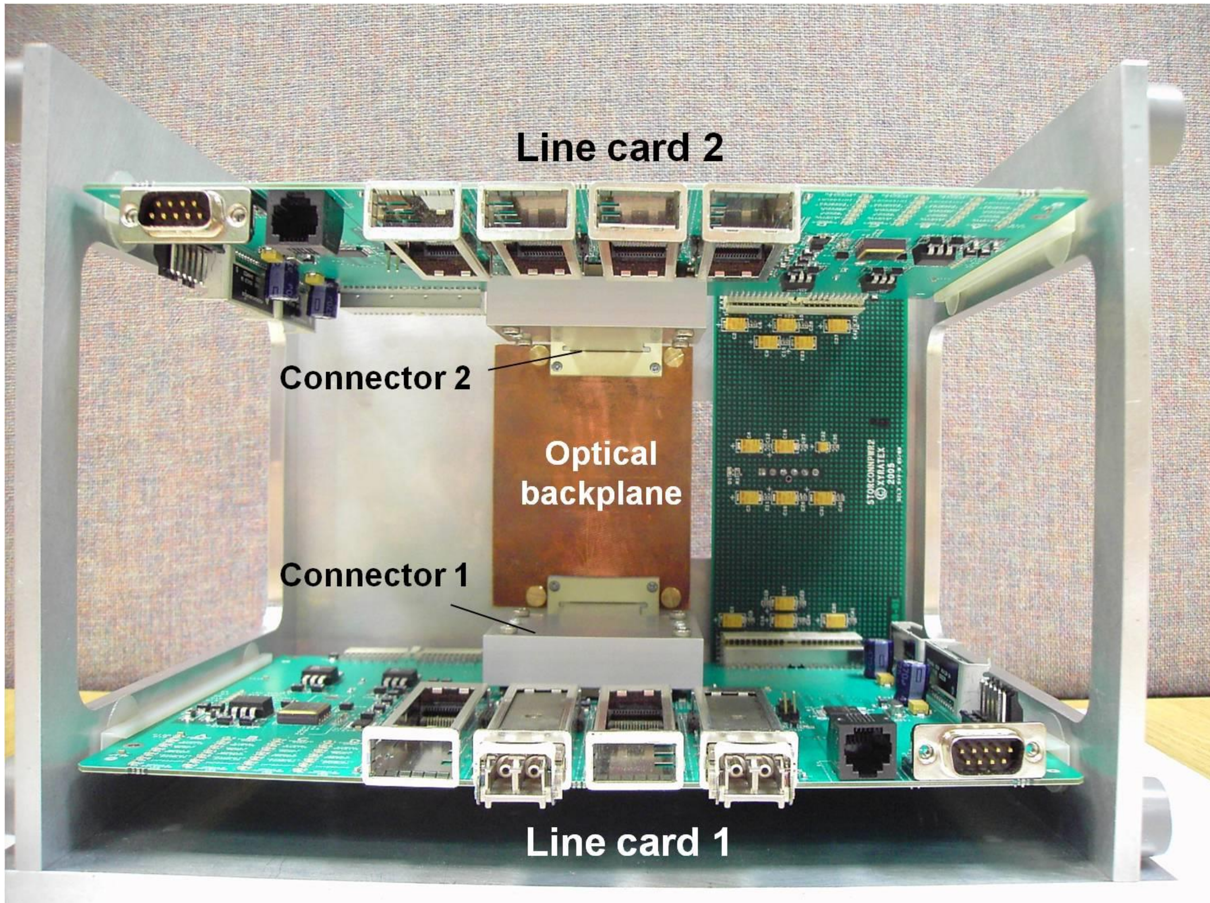


**Figure 3-62: Storlite demonstration platform populated with two test line cards with proprietary pluggable connectors plugged into a passive optical backplane and a passive electrical backplane.**

## 3.5 Summary

In this chapter, the invention, design and development of a first generation active pluggable optical connector has been described in detail, which allowed peripheral daughtercard devices to plug into and unplug from a passive optical backplane with embedded multimode polymer waveguides. The first generation optical connector required a two stage connection process, by which the active connector, mounted on a daughtercard was first plugged into the optical backplane, and then a user required to directly turn a cam on the connector housing to initiate the second, final stage of optical coupling, which brought the GRIN micro-lens and the waveguide interfaces into physical contact. Of course this would not be an appropriate process for a commercial connector in the long term as it would require users to reach into the enclosures to turn a cam each time a daughtercard was to be plugged in or removed. The efficiency of the module operating at 10.3 Gb/s was calculated to be 21.6 pJ/bit.

A technique was invented and successfully deployed to passively align and assemble optical receptacles to the multimode waveguide interfaces on an OPCB backplane with high precision. In order to avoid drilling notches either side of each waveguide interface close to the registration stubs, it was necessary to use an offset receptacle design, whereby the datum between the mechanical alignment slots is offset from the channel interface. This offset could however give rise to additional misalignment issues due to the pivoting effect between the datum and the channel interface. The greatest weakness of the first generation receptacle design, is that the waveguide end facets are exposed to the elements. In a forced air environment, this increases the risk of dust contamination and accumulation on the polymer acrylate end facets, which, after all, share many of the same physical properties as many adhesives.

A complete demonstration platform was designed and developed to test two daughtercards connecting to a separate passive optical backplane and power backplane. In order to focus on the validation of the active optical connector and the passive alignment techniques, the design of the demonstrator was simplified, with just the two daughtercards facing each other. This way the simplest optical waveguide design, based on straight groups of 12 waveguides could be used. This however would not be consistent with a typical backplane communication system, in which multiple daughtercards would be typically oriented in parallel rather than facing each other.

In the next chapter, the design and development of a second generation of active OPCB connector, enhanced backplane waveguide receptacle and an advanced demonstration system will be detailed, which addresses all the concerns raised above.

# 4 SECOND GENERATION PLUGGABLE ACTIVE OPTICAL CIRCUIT BOARD CONNECTOR FOR POLYMER WAVEGUIDE BASED ELECTRO-OPTICAL CIRCUIT BOARDS

## *4.1 Introduction*

### 4.1.1 FirstLight project summary

This chapter details the activities and achievements of the FirstLight project, which advances four key technology enablers namely i) low-cost, high precision techniques for passive alignment and assembly of optical interface components onto polymer waveguides, ii) complete pluggable in-plane connector solutions for

polymeric OPCBs, iii) design of compact optical waveguide layouts appropriate for data system midplane form factors and iv) fabrication of OPCBs with embedded polymer waveguides

As the backplane and its peripheral daughtercards are connected in a mutually orthogonal way, the author developed a more advanced in-plane pluggable connector technology and connection scheme whereby the optical interfaces of optical transceiver modules housed on the mating edge of the peripheral cards can be butt-coupled to optical channels embedded on the backplane (Figure 4-2a). This builds on the connection methodology demonstrated in the Storlite project [3], however utilizing an expanded beam connection allowing reduced susceptibility to dust and protecting the waveguide interfaces. According to this scheme the optical axis of the peripheral transceiver module is again collinear with the OPCB embedded optical channels, thus eliminating the need for right-angled mirrors and minimising the number of boundaries incurring optical loss.

In order to evaluate the viability of these technologies in a data centre environment, the author designed a second generation optical backplane connection demonstration platform, the FirstLight demonstration platform, which brings these technologies together into a high bandwidth density data communications enclosure.

### 4.1.2 FirstLight system design

Table 4-1 contains the formal designations for each circuit board type designed for the FirstLight system.

The author carried out the full design, development and characterisation of the FirstLight active pluggable optical connector including the StorConn3 quad parallel optical transceiver, and the StorConnTest3 10 GbE LAN physical layer relay card in the FirstLight platform.

The author jointly designed the FirstLight electro-optical backplane (StorConnOpt3) with Kai Wang of University College London. Specifically, the author designed the electrical circuit board layers and electrical and optical component layout and developed the waveguide requirements specification including the optical channel mapping. Kai Wang co-designed and characterised the complex optical waveguide layout, which will be detailed in section 4.3.

**Table 4-1: Storlite card designations**

| Card designation | Description |
|---|---|
| StorConn3 | Quad parallel optical transceiver circuit |
| StorConnTest3 | 10 GbE LAN physical layer relay card |
| StorConnOpt3 | Electro-optical Compact-PCI backplane |

The FirstLight platform (Figure 4-2a) comprised four test daughtercards (StorConnTest3), which were plugged electrically and optically into an electro-optical backplane (StorConnOpt3), which incorporated conventional copper layers for electrical power distribution and electronic signal interconnect, and an optical interconnect layer comprising polymer optical waveguides.

Each StorConnTest3 daughtercard supports one active pluggable optical connector, which incorporates the StorConn3 quad 10 Gb/s parallel optical transceiver circuit.

The author developed an improved iteration of the pluggable active optical connector, which incorporated a high speed parallel optical transceiver and a passive alignment mechanism to ensure accurate dynamic optical engagement between the transceiver interfaces and the embedded polymer optical waveguides in the OPCB.

As with Storlite, the polymer waveguide layer included self-alignment features to enable passive alignment and assembly of proprietary optical connector receptacle devices onto the polymer waveguides. A complex optical interconnect pattern was designed to meet exacting specifications to demonstrate how polymer waveguides would perform when subjected to the routing constraints expected within a conventional midplane form factor.

Figure 4-1a shows the fully assembled FirstLight demonstration platform with the motherboard and all four test daughtercards inserted. Figure 4-1b shows the fully populated FirstLight enclosure powered up.

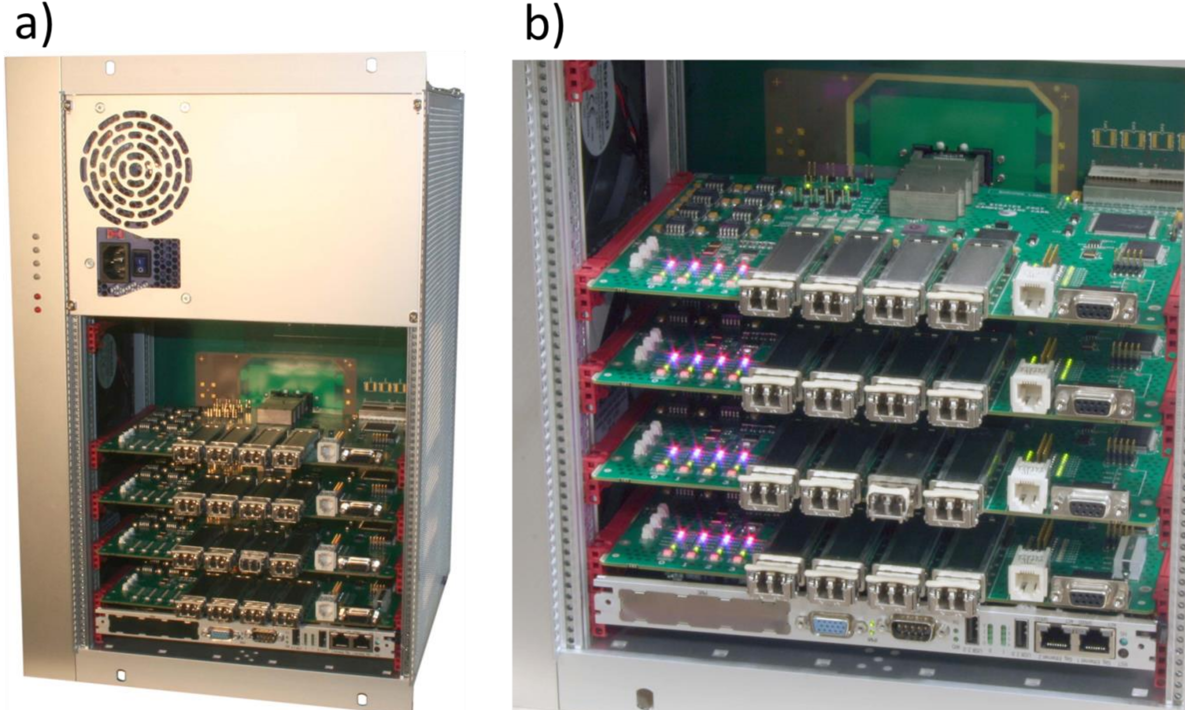


**Figure 4-1: a) FirstLight demonstration enclosure fully populated, b) FirstLight demonstration enclosure powered**

## *4.2 FirstLight pluggable active electro-optical backplane connector*

The author developed a prototype active pluggable connector (Figure 4-2b) to allow optical connection between the peripheral line cards and the optical layer embedded in the backplane. The connector comprised a parallel optical transceiver (StorConn3), connector housing and a pluggable engagement mechanism. The connector housing and engagement mechanism were designed in conjunction with US connector company Samtec.

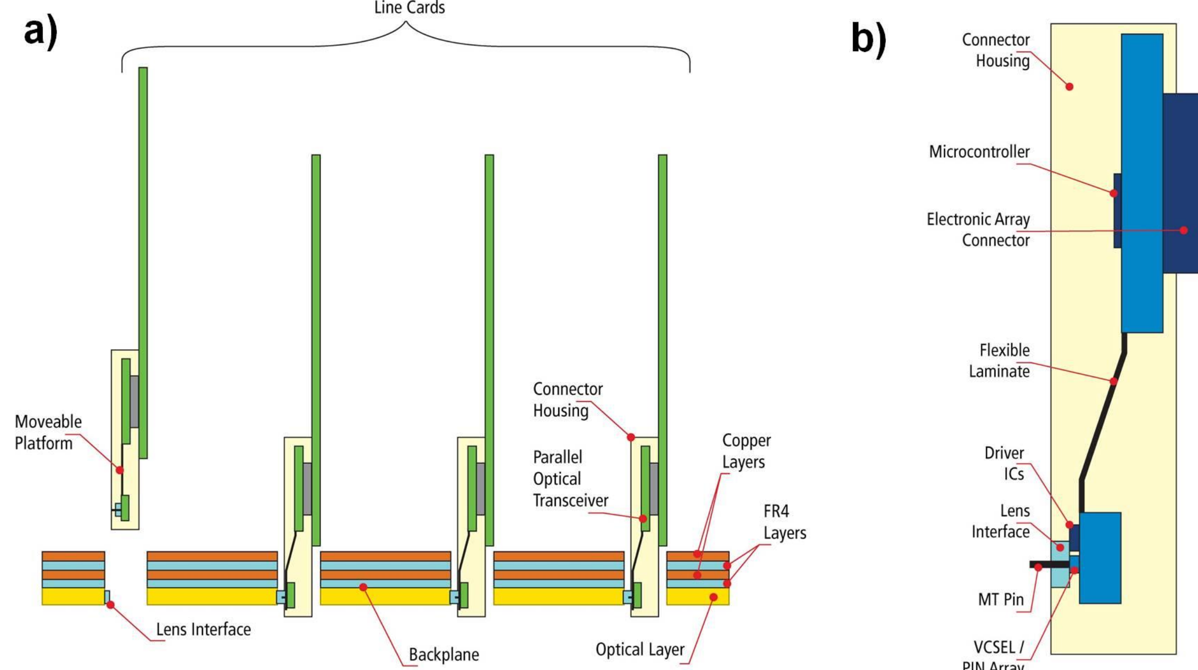


**Figure 4-2: FirstLight pluggable optical PCB connector concept. Pluggable active connector modules are inserted into the top of an optical backplane and engage with an optical layer on the bottom side of the backplane: (a) Electro-optical backplane connection scheme; active pluggable connectors housed on the edge of peripheral line cards engage with the embedded optical layer in the backplane PCB.**

**(b) Active optical connector comprising a parallel optical transceiver, connector housing and engagement mechanism.**

### 4.2.1 StorConn3 – FirstLight quad parallel optical transceiver

A quad parallel optical transceiver circuit was designed and developed to better implement the proprietary connection technique and address some of the weaknesses of the preceding StorConn2 circuit including: 1) the need to physically turn the cam on the connector to engage the optical interface after the daughtercard has been fully inserted, 2) the susceptibility of the exposed optical waveguide interface and optical connector interface to dust and 3) the lateral rigidity of the flexible bridge section i.e. a large amount of force is required to move the optical interface laterally, in any direction in the plane of the transceiver circuit.

Figure 4-3 shows a functional diagram of the StorConn3 transceiver circuit. The functional principle of the StorConn3 circuit is similar to that of the preceding StorConn2 transceiver, in that it will allow the photonic interface to float freely within a given vertical range relative to the test daughtercard on which the transceiver is mounted. As with StorConn2, the StorConn3 base section contains a high speed electronic array connector, which conveys electronic signals to and from the transceiver circuit from and to the StorConnTest3 test daughtercard PCB, on which the transceiver is mounted. The high speed differential signals are conveyed from the base section across a flexi-rigid bridge to the optical interface section. The optical interface comprises a four channel 850 nm VCSEL transmitter array and corresponding VCSEL driver array, and a four channel 850 nm PIN photodiode receiver array with corresponding TIA array. Each of the four differential electronic signals output from the TIA pass through a 10.7 Gb/s adaptive receive equaliser after which it is conveyed via the high speed electronic array connector to the StorConnTest3 test daughtercard.

StorConn3 has the following differences compared to its predecessor:

i. The optical interface is designed to be compliant with a 1x12 MT ferrule interface, in which the optical channels are located on the datum connecting the two MT pins protruding through the board, rather than offset.

ii. The lens used is a geometric microlens array rather than a GRIN lens array and forms one half of a dual lens expanded beam optical interface, which both renders the connection less susceptible to dust and protects the OPCB waveguide interface.

iii. The flexible bridge contains slots between the differential pairs in order to improve the mechanical flexibility in both the vertical direction (normal to the plane of the transceiver circuit) and the lateral direction (in the plane of the transceiver circuit)

iv. A microcontroller is on the transceiver itself rather than the test daughtercard

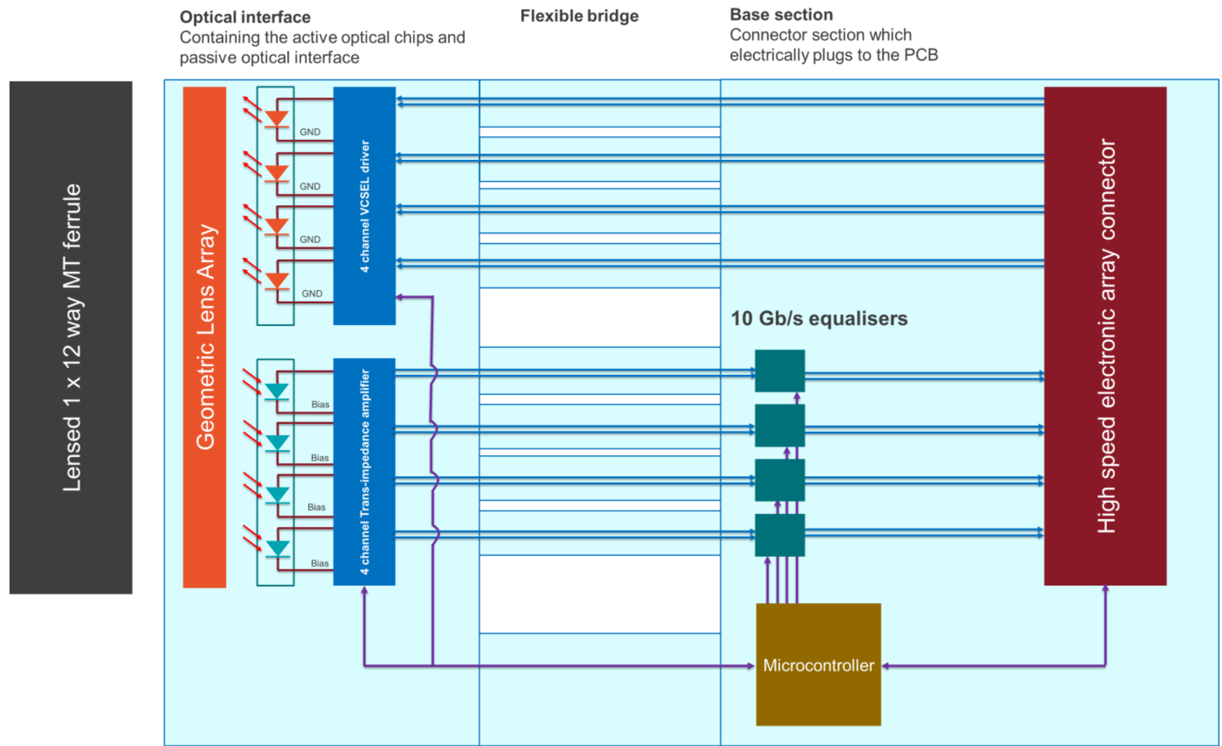


**Figure 4-3: Functional diagram of StorConn3**

Figure 4-4 provides an annotated view of the StorConn3 quad transceiver circuit mounted on a partially flexible and rigid substrate. The transceiver circuit comprised three sections: a base section (Figure 4-4e), a flexible bridge section (Figure 4-4d) and a moveable optical platform (Figure 4-4c). The circuit was constructed on a flexible laminate substrate, which was reinforced with rigid FR4 layers in the base section and optical platform leaving the intermediary bridge section flexible. The base section allowed for the electrical connection of the transceiver to the peripheral line card by means of a high speed electronic array connector provided by Samtec capable of supporting data rates up to 11 Gb/s.

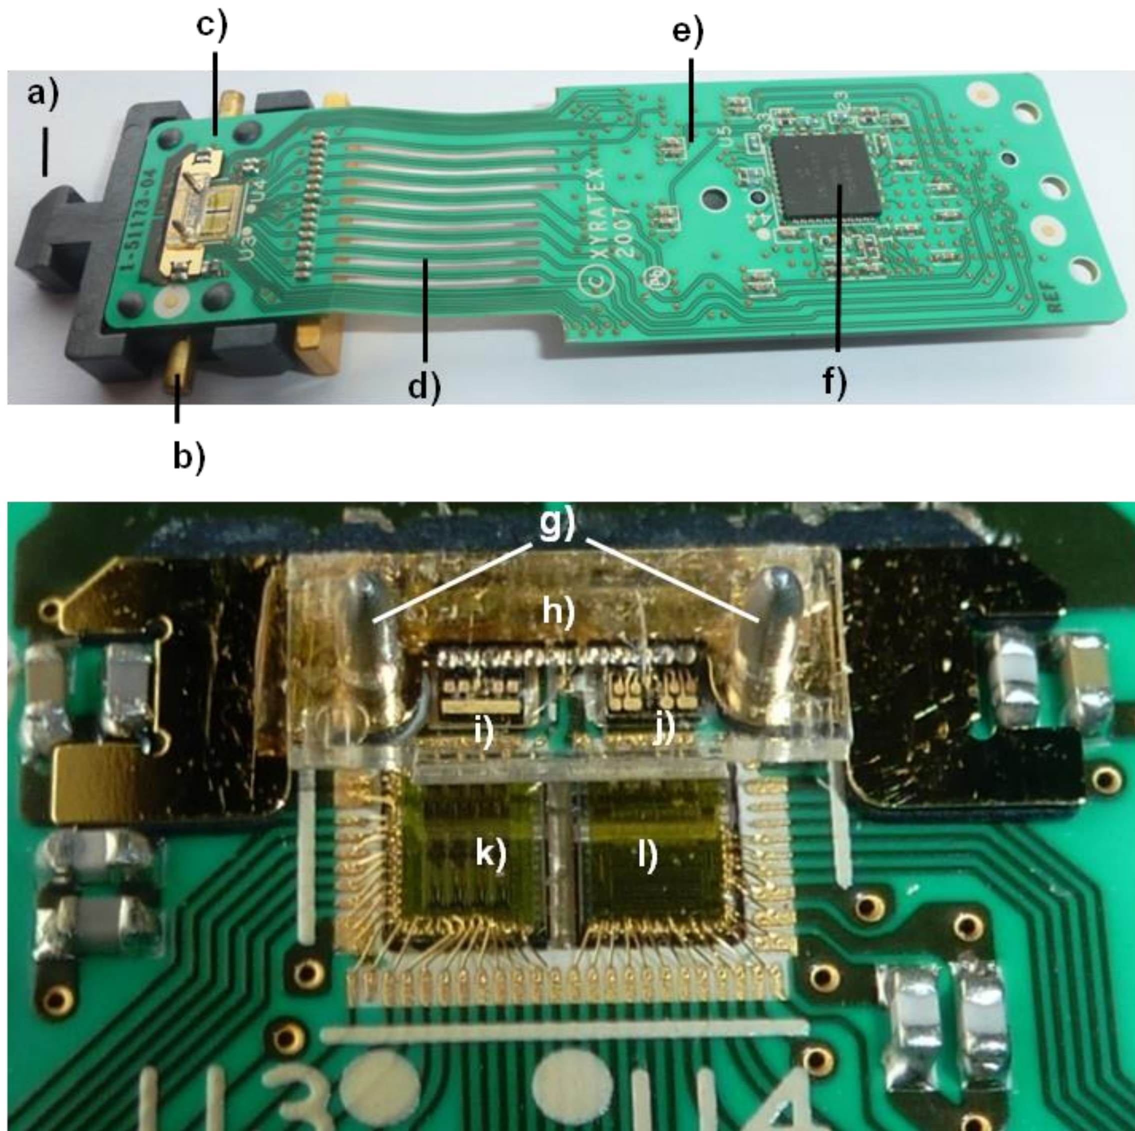


**Figure 4-4: Parallel optical transceiver circuit and optical interface a) Front ramped plug, b) Optical platform guide pins, c) Optical interface, d) Flexible bridge, e) Rigid base section, f) Microcontroller, g) MT compatible alignment pins, h) Micro-lens array, i) Quad VCSEL array, j) Quad PIN photodiode array, k) Quad VCSEL Driver array, l) Quad transimpedance amplifier/limiting amplifier TIA/LA array**

A microcontroller (Figure 4-4f) provided a two-wire control interface through which various parameters of the transceiver could be externally controlled or monitored including laser bias and modulation currents, receiver squelch, signal detect and IC temperature read back. The design allowed for multiple transceivers to operate as slave devices on a single two-wire communications bus.

The moveable optical platform contained a quad VCSEL array (Figure 4-4i), a VCSEL driver array (Figure 4-4k), a PIN photodiode array (Figure 4-4j) and a quad transimpedance amplifier/limiting amplifier (TIA/LA) (Figure 4-4l). The uncooled VCSELs emitted at a nominal wavelength of 850 nm with a full beam divergence of 30º giving rise to an N.A. of 0.26, which is 78% of the N.A of the backplane waveguides. The PIN photodiodes were responsive to the same wavelength and had a circular receive aperture of 70 μm diameter, which was chosen to be large to maximise misalignment tolerances and to reduce modal or speckle noise. The PIN photodiodes, which were held under reverse bias, had a nominal responsivity of 0.62 A/W and a bandwidth of 7.8 GHz. Zarlink provided the quad VCSEL arrays, VCSEL driver arrays, TIA arrays and PIN photodiode arrays.

Two MT pins were assembled into the optical interface platform. The VCSEL and photodiode arrays were pre-attached to a lead frame, which included two MT compliant guide slots. These were used to guide the lead frame over the pins and thus accurately align the VCSEL and photodiode arrays to the protruding MT pins.

An MT compliant, twelve channel microlens array (Figure 4-4h) with two MT guide holes was aligned along the MT pins protruding from the optical platform, over the active emitting and receiving apertures of the VCSEL and photodiode arrays and attached to the lead frame with UV curable adhesive. The VCSEL and photodiode arrays were spaced apart on the lead frame such that the VCSEL emitting apertures were aligned to the four micro-lenses on the left hand side and the photodiode receiving apertures were aligned to the four microlenses on the right hand side (as seen from above with the base section on the bottom and optical interface on the top), thus the central four lenses remained unused. In future, it would be easy to incorporate higher channel numbers into the same form factor such as 12 or even 24 channel arrays.

Each channel was capable of sustaining a data-rate of 10.3 Gb/s giving rise to an aggregate bandwidth of 82.4 Gb/s. The optical connection interface comprised a collimating 1×12 micro-lens array and a pair of mechanical registration pins, designed to be compliant with MT style parallel optical interfaces. This allowed for stand-alone testing of both the transceiver and the backplane waveguides with a lensed MT terminated fibre-optic patch-cord. The flexible bridge section allowed the optical platform to float relative to the peripheral device, thus ensuring that when coupled to the backplane, the optical connection remained relatively impervious to transient movements and vibrations in the system (Figure 4-4d).

The circuit was designed to be mounted into a connector housing (Figure 4-16), which included grooves to enable the required movement of the optical interface during the pluggable connection process described below. The transceiver included guide pins on the sides of the optical platform (Figure 4-4b), which engage with slots in the connector housing to support the connection process described later. The transceiver also included a ramped plug (Figure 4-4a) on the front of the optical platform, which engages with the primary receptacle in the backplane in the first part of the connection process.

### 4.2.2 StorConn3 circuit hardware design

The StorConn3 raw card consists of a Kapton Polyimide based flexible PCB laminated onto a six layer FR4 reinforced rigid PCB in the optical head and base sections. The bridge section of the Kapton Polyimide PCB is not laminated onto an FR4 reinforced rigid PCB and is therefore mechanically flexible. In order to maximise the mechanical flexibility of the flexible bridge, slots were cut in the substrate between the high speed differential pairs.

Figure 4-5 shows the physical shape of the card and the principal dimensions of its outline.

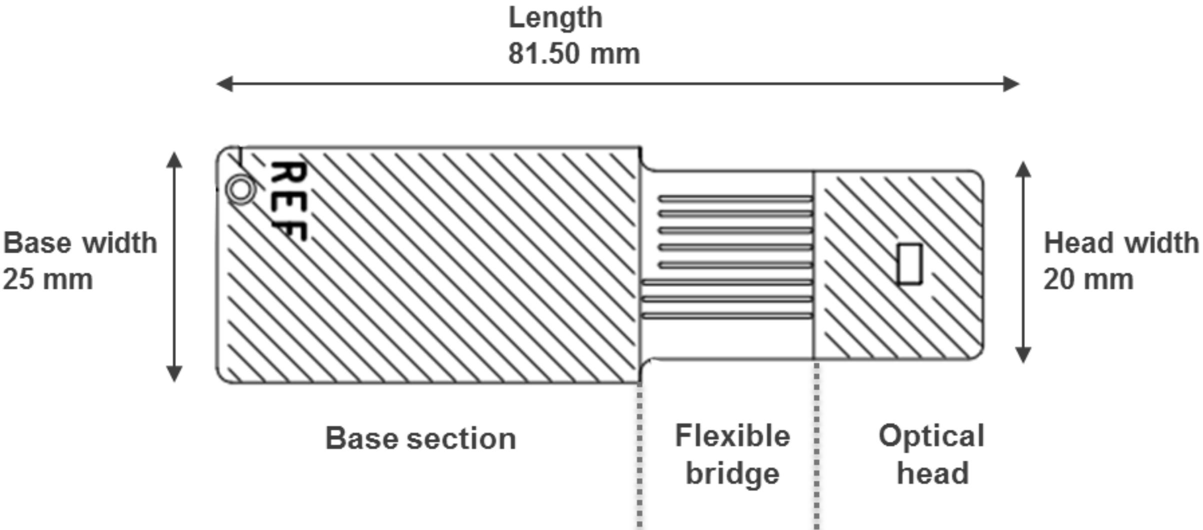


**Figure 4-5: StorConn3 circuit board form factor**

#### *4.2.2.1. PCB layer stack-up*

Figure 4-6 outlines the complete layer stack-up of this circuit board.

StorConn3 has eight electrical layers, of which four are signal layers (layers one, three and six), two are ground layers (layers two and seven) and two are power layers (layers four and five).

The high speed differential tracks are located on the top layer (layer one) and bottom layer (layer eight) and designed with a controlled differential impedance of 100 Ω ± 8%. As these differential tracks are on the surface layers and referenced to an adjacent ground layer, they are referred to as microstrip type.

Layers three and six are used for low speed signals such as the two-wire interface signals between the microcontroller and the VCSEL driver and TIA arrays.

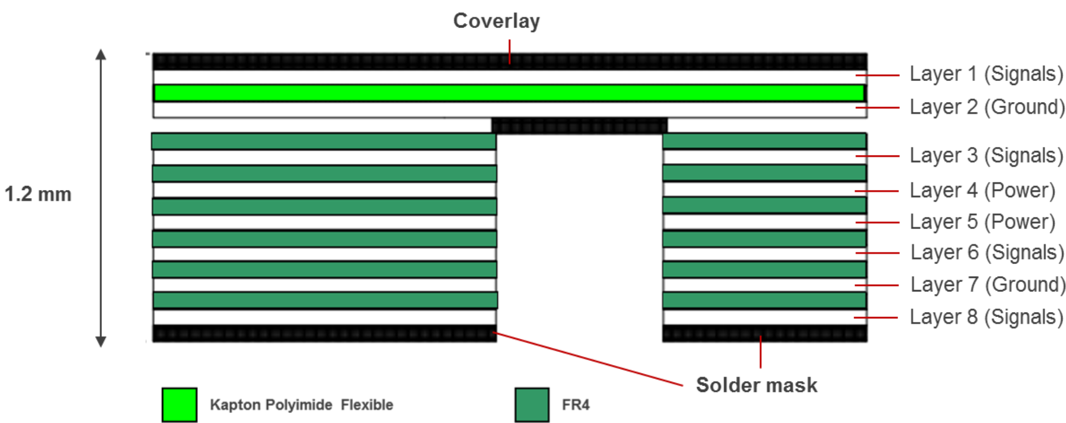


**Figure 4-6: StorConn3 PCB layer stack-up**

### 4.2.3 Dual lens coupling solution

As misalignment tolerance is important to minimise the connection cost, an expanded beam connector design was devised. The micro-lens array attached to the transceiver formed one half of a dual lens arrangement, and the second micro-lens array was part of the secondary receptacle, assembled over the waveguide interface on the backplane. When coupled, the transceiver lens array in combination with the backplane lens array served to image the VCSEL output into the backplane waveguide and the waveguide output into the photodiode aperture. Simulations were carried out by Omron using Zemax ray tracing software to select the optimum configuration of coupled micro-lenses required to image the output of the transceiver VCSEL onto the launch facet of the waveguide (Figure 4-7a) and the output of the waveguide onto the receiving aperture of the photodiode (Figure 4-7b). It was determined that the best way to achieve this with minimum insertion loss was to deploy a focusing micro lens array P1L12A-F1 over the transceiver interface and a collimating micro lens array P1L12A-C1 over the waveguide interface as shown in Figure 4-8. A total insertion loss of 0.72 dB was calculated over the dual lens interface between the VCSEL and the waveguide, while a total insertion loss of 1.11 dB was calculated between the waveguide and the photodiode.

It should be noted that the launch profiles from both waveguide and VCSEL were again assumed in the simulation to have a Gaussian profile. The output of a waveguide could be assumed to be Gaussian, if the waveguide geometry (length, bends) is such that it will promote sufficient mode mixing, however VCSEL source profiles are not completely Gaussian in nature, therefore this simulation could be further optimised in future. In the simulation the rectangular step-index waveguide was approximated by a cylindrical MMF with a core diameter of 75 µm and an N.A of 0.27, in order to accommodate the limitations of the modelling capability available. This is designed to produce a far field pattern, which is closely similar to that of the waveguide, however, although this is a free space coupling arrangement, the relative distances between the waveguide and coupling interfaces are not large enough to fulfil the far field condition. The far field condition requires that in order for a distance to be safely considered in the far field domain, the distance would have to be substantially greater than $D^2/\lambda$, where D is the diameter of the aperture and $\lambda$ is the wavelength. Therefore in future the simulation model would have to take the real shape and size of the aperture into account in order to provide more accurate results.

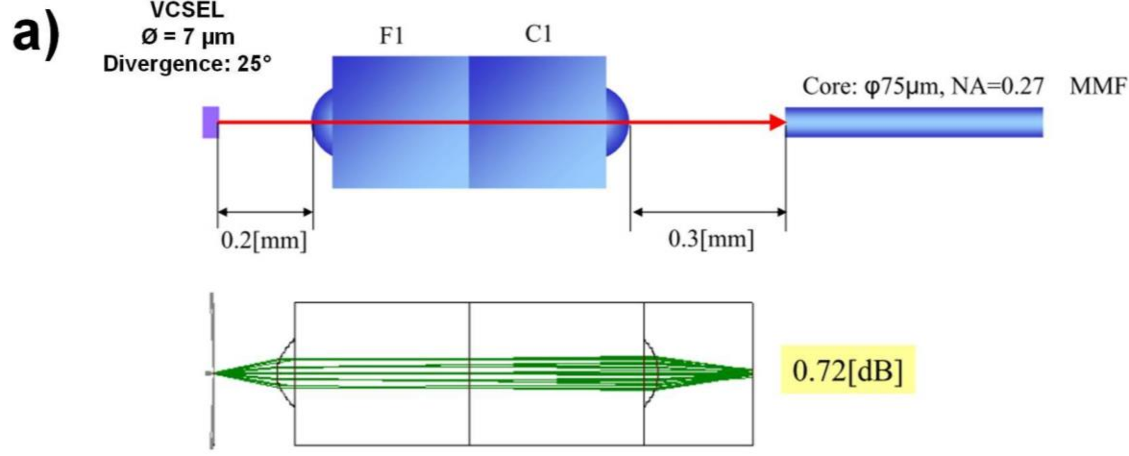


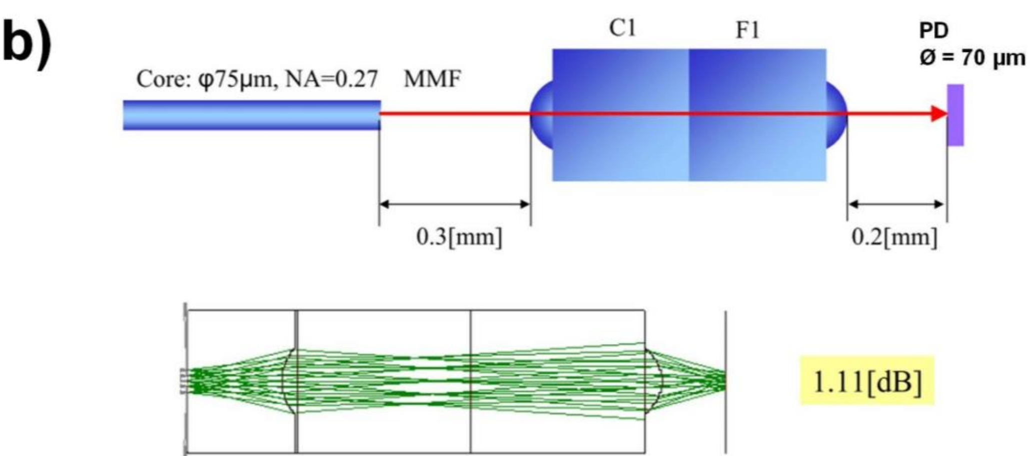


**Figure 4-7: a) Model and simulation results of optimum dual lens arrangement to image VCSEL to multimode ingress waveguide end facet predicting 0.72 dB loss in this transition, b) Model and simulation results for optimum dual lens arrangement to image egress waveguide end facet onto PD active area predicting 1.11 dB in this direction. (Source: Omron)**

The free space distance between the VCSELs, photodiodes, waveguides and their respective lens arrays was chosen to ensure that, at the point of interface between the two lenses, the optical beam was expanded to a width many times that of the source width, whether the source was the VCSEL or the waveguide. An expanded beam width of 105 µm was predicted at the optical mating interface between the VCSEL and waveguide, while a 195 µm beam width was predicted at the optical mating interface between the waveguide and the photodiode. One crucial benefit of this arrangement was to make the connector far less susceptible to contamination as any stray contaminants that settle on the lens interface would block a smaller proportion of the expanded beam than they

would a beam of similar size to the sources.

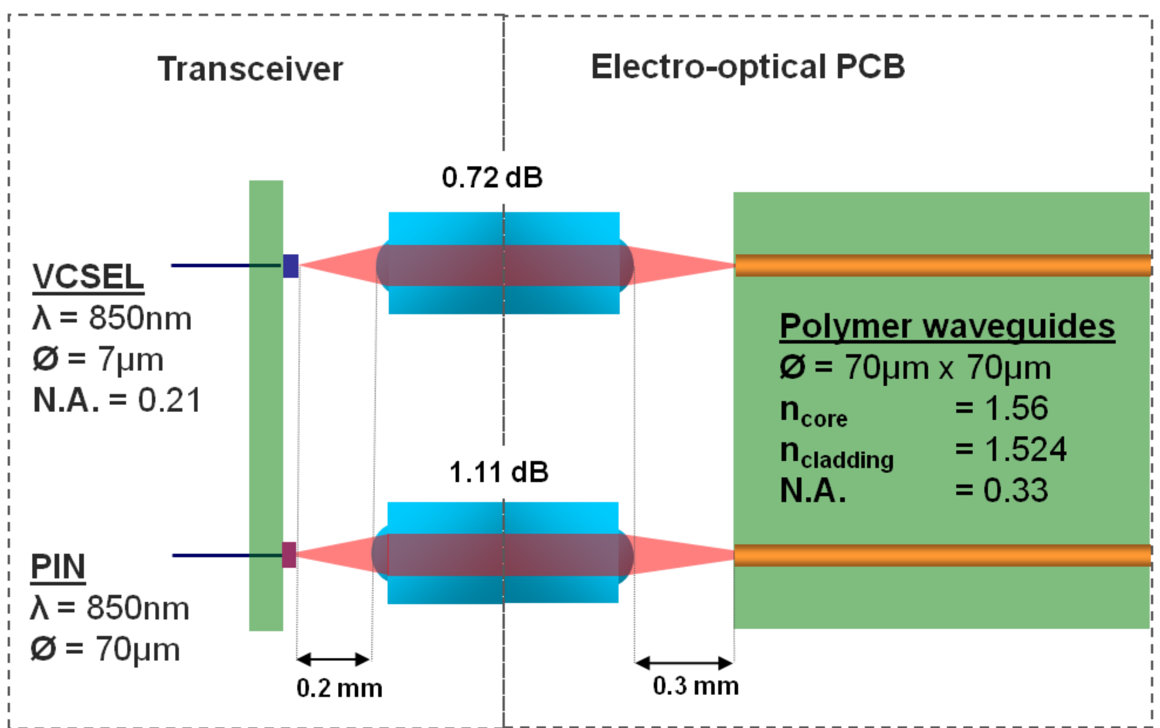


**Figure 4-8: Dual lens coupling interface between transceiver and OPCB**

### 4.2.4 StorConn3 component layout

Table 4-2 lists the main components of the StorConn2 card. Figure 4-9 shows the component layout.

**Table 4-2: StorConn3 component list and description**

| Item | Supplier | Part No. | Description |
|---|---|---|---|
| **Optical Head** | | | |
| VCSEL Driver | Primarion (Acquired by Infineon Technologies in April 2008) | PX6514 (obsolete) | 10 Gb/s data rate per channel (40 Gb/s aggregate data rate) |
| TIA / LA | | PX6524 (obsolete) | 10Gb/s optical receive system, Transimpedance amplifier & Limiting amplifier per channel |
| Passive heatsink | Xyratex inventory | custom | Attached to PCB onto copper slugs forming thermal contact with VCSEL Driver / Optical Receiver die |
| VCSEL array (chip) | Zarlink Semiconductor | ZL60126 | 10 Gb/s Vertical Cavity Surface Emitting Laser array (1 x 4) |
| PIN array (chip) | Zarlink Semiconductor | ZL60131 | 10 Gb/s PIN photodiode array (1 x 4) |
| Lead frame | Zarlink Semiconductor | 113876 | Lead frame with two MT pin slots, which holds the VCSEL and PIN array at a precise distance from each other such as to be compliant with 1x12 MT ferrule interfaces |
| MT pin holder and ramped plug | Samtec | Custom design | Custom part comprised of PC-ABS thermoplastic to fit under optical head including MT pin holding slots, ramped plug and heat staking pillars |
| Heatsink | Bromfield Precision | Custom design | Custom heatsink, which supports |

| | | | |
|---|---|---|---|
| | Engineering Ltd | | copper slugs under the VCSEL driver and TIA array ICs for more efficient heat transfer from ICs to heatsink |
| Microlens array | Omron | P1L | 1 x 12 geometric microlens array with MT pin slots, compliant with 1x12 MT ferrule interfaces |
| **Flexible Bridge** | | | |
| No active or passive components | | | |
| **Base** | | | |
| Signal equalisers | Maxim | MAX3805 | 10.7 Gb/s Adaptive Receive Equalisers to improve signal integrity of 10 Gb/s differential signals received from TIA/LA |
| Mezzanine connector | Samtec | SEAF-15-05.0-S-10-2-A | Copper array connector receptacle supporting differential signals up to and exceeding 10 Gb/s data rate and low speed control signals. |
| Microcontroller | Microchip | PIC18LF4331 | Attached to PCB onto copper slugs forming thermal contact with VCSEL Driver / Optical Receiver die |

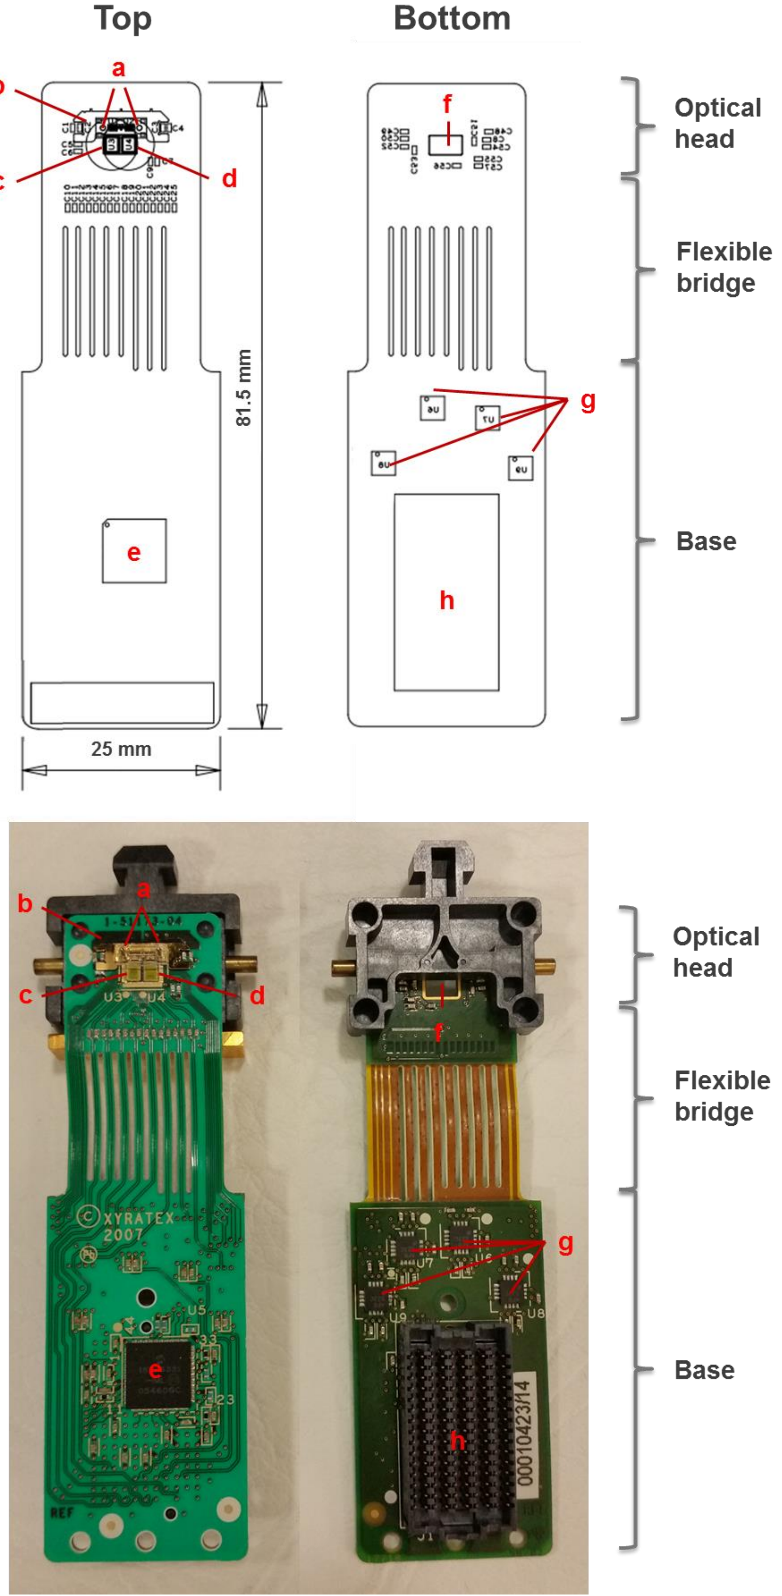


**Figure 4-9: StorConn3 component layout: a) MT compatible alignment pins, b) lead frame holding the VCSEL array and PIN photodiode array, c) quad VCSEL driver array IC, d) transimpedance amplifier/limiting amplifier TIA/LA array IC, e) microcontroller, f) rectangular slot in PCB for copper slug, g) adaptive signal equalisers, h) electrical array connector**

### 4.2.5 StorConn3 card assembly process

The complete assembly of the StorConn3 card requires, in addition to the conventional process of card population with standard passive and active components, the direct wire-bonded assembly of the VCSEL drive and TIA/LA array chips in raw die form, the integration of novel thermal dissipation structures and high-precision passive alignment processes for both chips and opto-mechanical structures.

The assembly process and sequence is outlined below.

#### *4.2.5.1. Mezzanine connector placement*

The Samtec **SE ARRAY® (SEAF)** high speed mezzanine connector is designed to provide high speed differential (10 Gb/s) and single ended electrical connection between two parallel boards. The connector utilises Ball Grid Array (BGA) for solder attachment to the PCB.

#### *4.2.5.2. Equaliser chip assembly*

The Maxim **MAX3805** 10.7 Gb/s adaptive signal equalisers, also used on the StorConn2 transceiver circuit, were packaged in a thin “quad-flat no-leads” (QFN) surface mount package with a thermal pad on the base and had to be assembled onto the board by means of a solder reflow process.

Subsequently, all passive components including resistors and capacitors were solder-attached to the card in the standard manner.

#### *4.2.5.3. MT pin holding unit assembly*

Figure 4-10 shows a custom MT pin holding unit designed by Samtec to fit under the optical head section, which included:

- MT pin holding slots populated with steel MT compliant pins of 0.7 mm diameter
- a ramped plug, which forms part of the connection mechanism
- four heat staking pillars, which allow the part to be permanently fastened to the PCB
- cylindircal optical platform brass guide pins

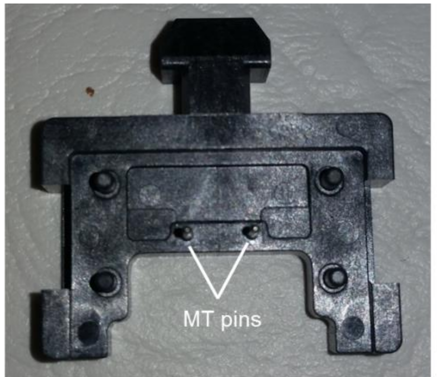

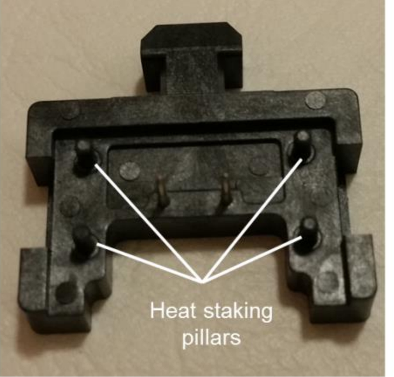

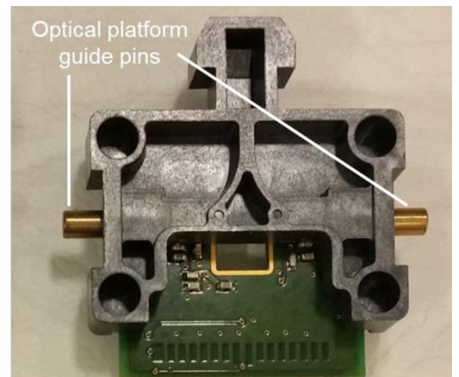


**Figure 4-10: MT pin holder and ramped plug part**

The MT pin holding unit is aligned to the bottom of the optical head section, such that its four heat staking pillars and MT pins are slotted through compliant holes in the PCB.

The MT pin holder is then fastened to the PCB through a thermoplastic staking process, whereby the four protruding heat staking studs are then softened and deformed through temporary application of heat in order to create an interference fit between the studs and the PCB upon cooling.

### *4.2.5.4. Custom heatsink and thermal dissipation structures*

The author designed a custom heatsink, which was optimally shaped to maximise heatsink volume in the confined space under the optical interface. The heatsink, shown in Figure 4-11, includes a rectangular clearance (black), which will span both the VCSEL driver and optical receiver die, which each dissipate 0.55W. The clearance will allow a rectangular copper slug to be inserted, which will protrude through the compliant rectangular clearance in the PCB (Figure 4-9f) and form direct thermal contact with the base of the die and transfer the heat effectively into the heatsink volume.

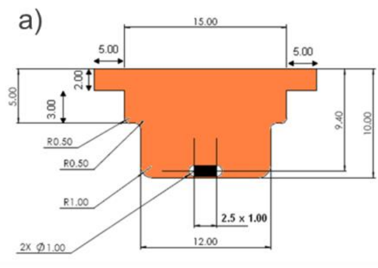

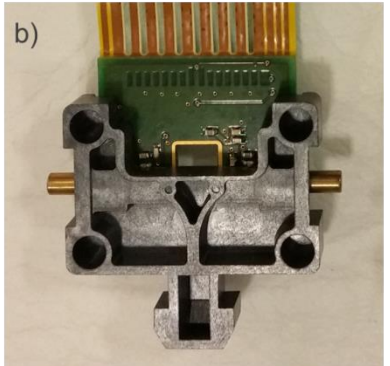

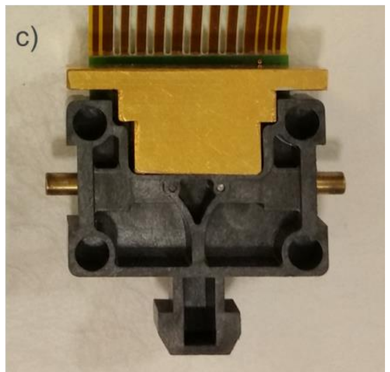


**Figure 4-11: Custom heatsink: a) heatsink design, b) optical head section prior to heatsink attachment, c) optical head section with heatsink attached**

#### *4.2.5.5. Optical driver die assembly*

The Primarion **PX6514** VCSEL driver and **PX6524** optical receiver die (also used on the StorConn2 transceiver card) were positioned over the single rectangular thermal dissipation slug and attached to the top surface of the slug with a thermally conductive epoxy, thus establishing a thermal channel to the underlying custom heatsink. The die pads were then wire-bonded to the compliant pads on the PCB.

#### *4.2.5.6. VCSEL array and photodiode array lead frame assembly*

The Zarlink **ZL60126** quad VCSEL array and Zarlink **ZL60131** quad PIN photodiode array die were pre-attached to a lead frame, which held the VCSEL array and PIN array at a precise location relative to each other and to two MT compliant pin slots (Figure 4-12).

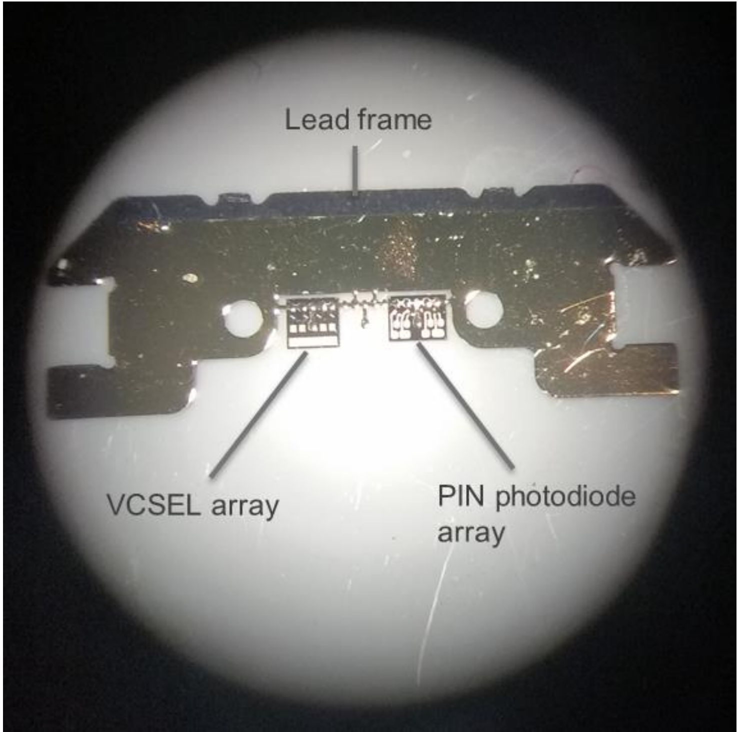


**Figure 4-12: Zarlink lead frame holding VCSEL array and PIN photodiode array**

Figure 4-13 shows how the lead frame is aligned onto the PCB using the MT pins protruding through the PCB from the MT pin holding unit under the PCB. Once the lead frame is fully descended, the bases of the VCSEL and PIN photodiode die will be in physical contact with their respective pads (Figure 4-13a).

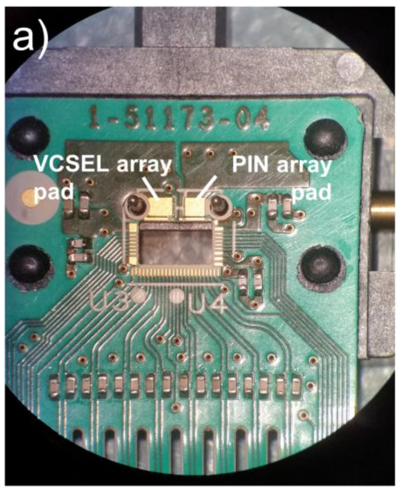

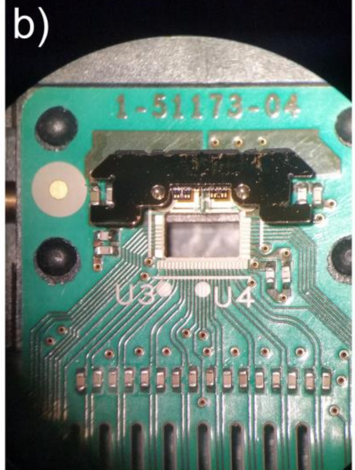

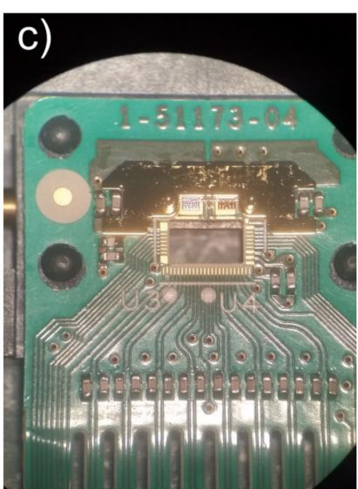


**Figure 4-13: Lead frame alignment: a) optical head with two MT pins protruding through the PCB, b) lead frame slots aligned over the MT pins, c) lead frame fully descended into place. During actual assembly the VCSEL driver and PIN photodiode array ICs will already have been secured by this time**

The bases of the VCSEL array and PIN array are secured in place with a thermally and electrically conductive solder paste. The cathode and anode pads on the VCSEL and PIN arrays were wire bonded directly to the recipient pads on the VCSEL driver IC and TIA/LA IC respectively, thus substantially reducing attenuation and signal degradation effects compared to the approach with StorConn2 where the VCSEL and PIN arrays were bonded to the PCB first and then to the VCSEL driver and TIA/LA array ICs.

#### *4.2.5.7. Geometric lens assembly*

Figure 4-14 shows dimensioned drawings of the Omron **P1L** geometric microlens array, with 12 microlenses and 2 MT compliant pin slots, which was designed to fit over a 1x12 MT ferrule interface to form a lensed, expander beam assembly. The entire part was comprised of a proprietary polycarbonate material, which was injection moulded. On the bottom surface there are four mounting feet and 12 protruding half convex microlenses. The mounting feet are the same thickness as the microlens protrusions and prevent damage to the microlens surfaces when the microlens assembly is placed on a flat surface. The other surface is completely flat and is designed to be the mating surface, which makes physical contact with the flat surface of another lensed MT ferrule.

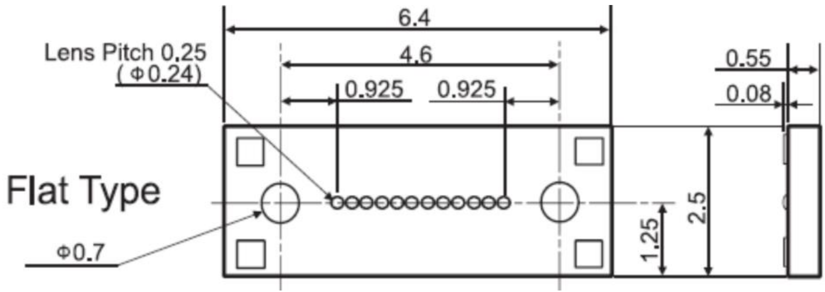


**Figure 4-14: Dimensioned drawings of Omron P1L lens array (Source: Omron)**

As with the lead frame, passive alignment of the microlens array is achieved by easing the MT compliant pin slots in the microlens assembly onto the protruding MT pins and gently manoeuvring it down into position over the lead frame.

Figure 4-15 shows the optical head with the Omron microlens assembled and the MT pins protruding through them

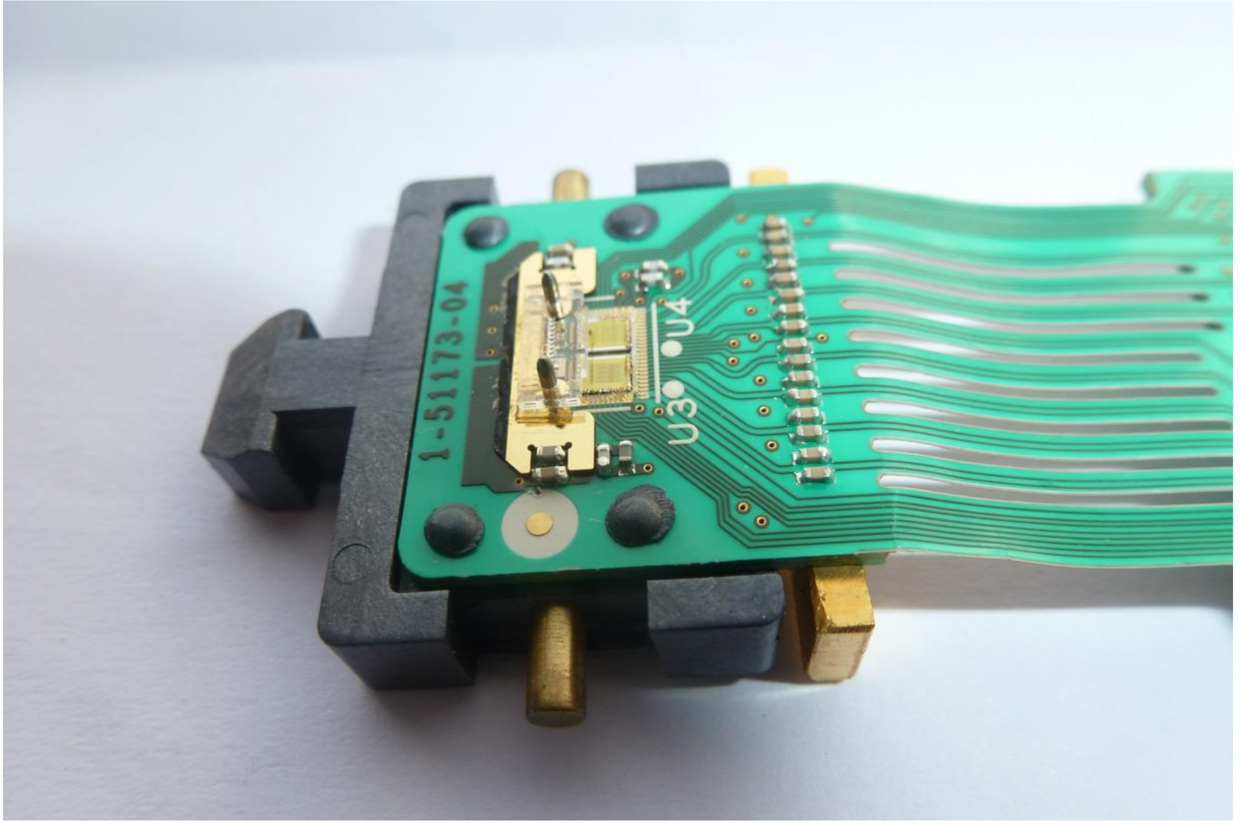


**Figure 4-15: Optical interface section with Omron microlens assembled**

### 4.2.6 Pluggable connector mechanism

As shown in Figure 4-16 the transceiver circuit (Figure 4-16a) was assembled into a connector housing (Figure 4-16b) wherein the two lateral guiding pins, which form part of the transceiver optical platform were slotted into compliant grooves in the housing (Figure 4-16c).

This enabled the controlled movement of the optical platform relative to the housing as required during the two-stage pluggable connection process, which is described as follows: As the peripheral line card is first inserted into the backplane, the ramped plug at the front end of the transceiver is funneled into the larger primary backplane receptacle and the transceiver lens array moved into position under the backplane lens array housed in the secondary receptacle (Figure 4-17a).

As the connector is then pushed further into the larger backplane receptacle, the lateral guiding pins on the optical platform are guided along the grooves in the connector housing, which are angled such as to move the transceiver lens array towards the backplane lens array. The MT pins on the optical platform then engage with the MT compliant slots in the secondary receptacle aligning both lens arrays to each other with a high degree of precision (Figure 4-17b). When the peripheral line card is extracted, the connection process is reversed.

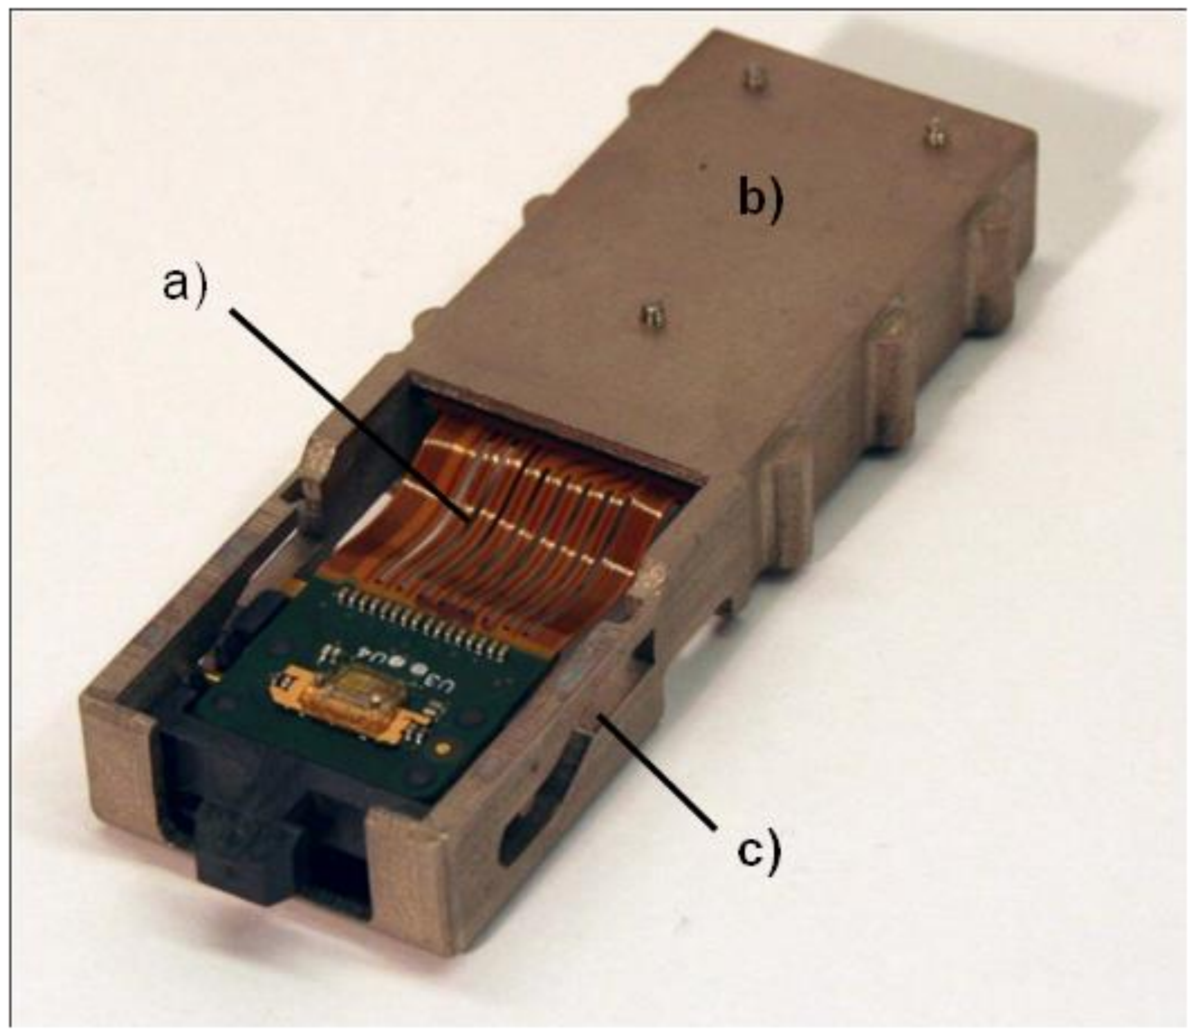


**Figure 4-16: a) Optical transceiver circuit mounted on flexi-rigid substrate, b) Connector module housing, c) Grooves to enable required movement of optical interface during mating process.**

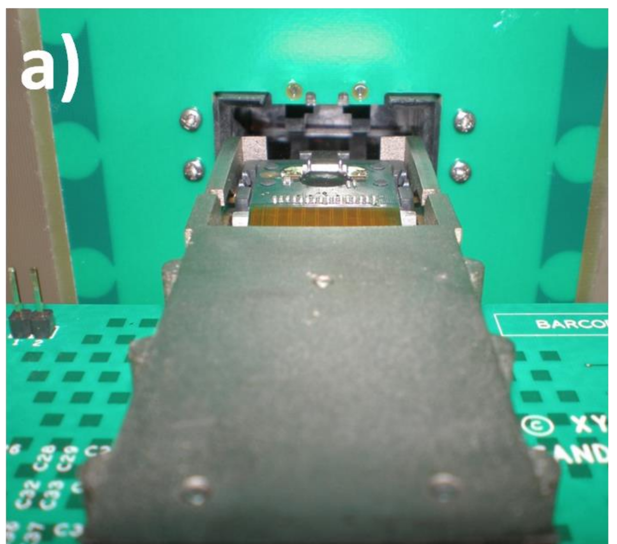


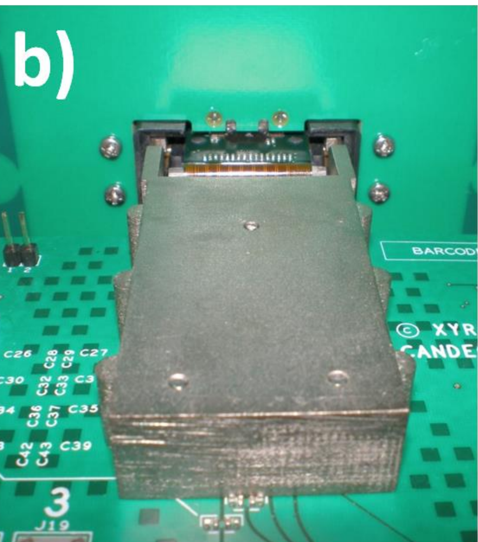


**Figure 4-17: a) Optical connector during first stage of coarse engagement with OPCB connector receptacle, b) Optical connector fully plugged into OPCB connector receptacle**

## *4.3 StorConnOpt3 - electro-optical backplane with embedded polymeric waveguides*

### 4.3.1 Polymer optical waveguide layer

A passive OPCB backplane was designed, which included ten electrical layers devoted to power distribution and low-speed signal communication and one optical polymer interconnect layer to convey high speed (10.3 Gb/s) serial data between peripheral active optical connectors plugged into the backplane. The optical interconnect layer was fabricated from an acrylate/polyurethane polymer, which was proprietary to IBM Research, exhibiting a propagation loss of 0.03 - 0.04 dB/cm at a wavelength of 850 nm. The optical layer stack comprised a core layer, sandwiched between a lower and an upper cladding, whereby the polymer in the guiding core layer had a higher refractive index than that in the bounding cladding layers (Figure 4-18). The refractive index of the core material was $n_{core}$ = 1.5600, while that of the cladding was $n_{cladding}$ = 1.5240, giving rise to step-index multimode waveguides with a core - cladding index difference of $\Delta n$ = 2.3% and a numerical aperture (N.A.) of 0.33 as measured at a wavelength of 850 nm. Rectangular channels were patterned using a vectorial Laser Direct Imaging (LDI) writing process to define waveguides in the core layer with a cross-section of 70 µm x 70 µm, which was suitable to meet the launch and capture tolerances on the optical transmit and receive elements [151].

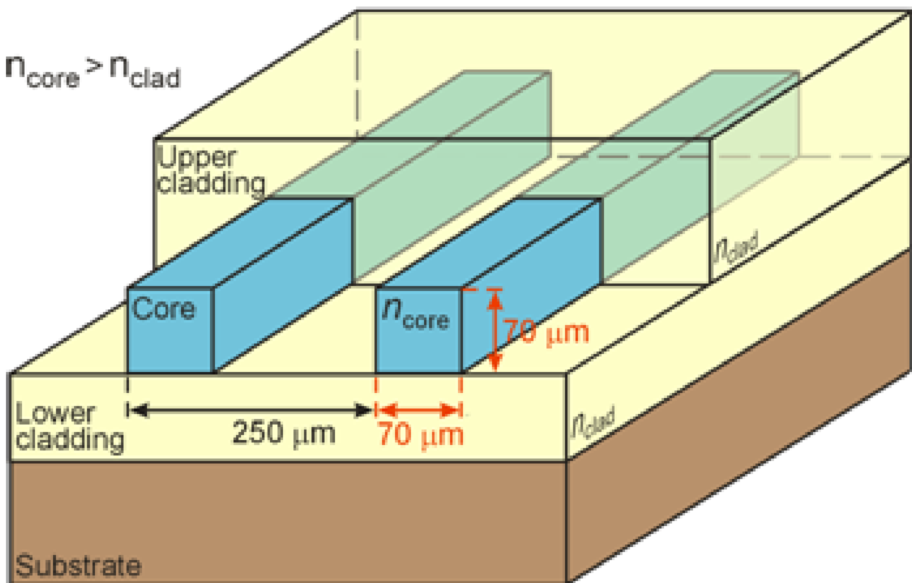


**Figure 4-18: Schematic illustrating waveguide embedded in lower and upper claddings (Source: IBM Research – Zürich)**

### 4.3.2 Optical interconnect design

A complex polymer waveguide interconnect layout was designed to form the optical layer of a 262 mm x 240 mm electro-optical PCB and was guided by two key requirements:

1) All midboard optical engagement interfaces were oriented in the same direction to allow line cards to be pluggable in arbitrary slots and interchangeable

2) The design needed to demonstrate the routing compactness and manoeuvrability, which is typically required on a high density data storage system backplane.

Therefore waveguide structures needed to be carefully designed according to a set of basic optical waveguide layout design rules [151] in order to minimise the optical loss in each waveguide segment and ensure aggregate (total) insertion loss for each waveguide falls within the receiver sensitivity threshold to allow bit error free signal transmission.

The optical interconnect layer design, shown in Figure 4-19, was jointly designed by the author, Kai Wang at UCL and IBM Research. It was defined by a complex routing pattern (Figure 4-19a), which included four quasi-rectangular optical engagement apertures, multiple non-orthogonal crossings and both negatively (Figure 4-19b) and positively (Figure 4-19c) cascading 90° bends. The engagement apertures were interconnected by a point-to-point topology, whereby each aperture is connected by one bidirectional link (comprising two waveguides) to every other aperture, resulting in a total of 12 waveguides on the board. The sizes of the engagement apertures were determined by the form factor of the pluggable connector prototype, which will be described in the next section. It should be noted that the optical layout was not designed to minimise link losses, but to highlight different stressors.

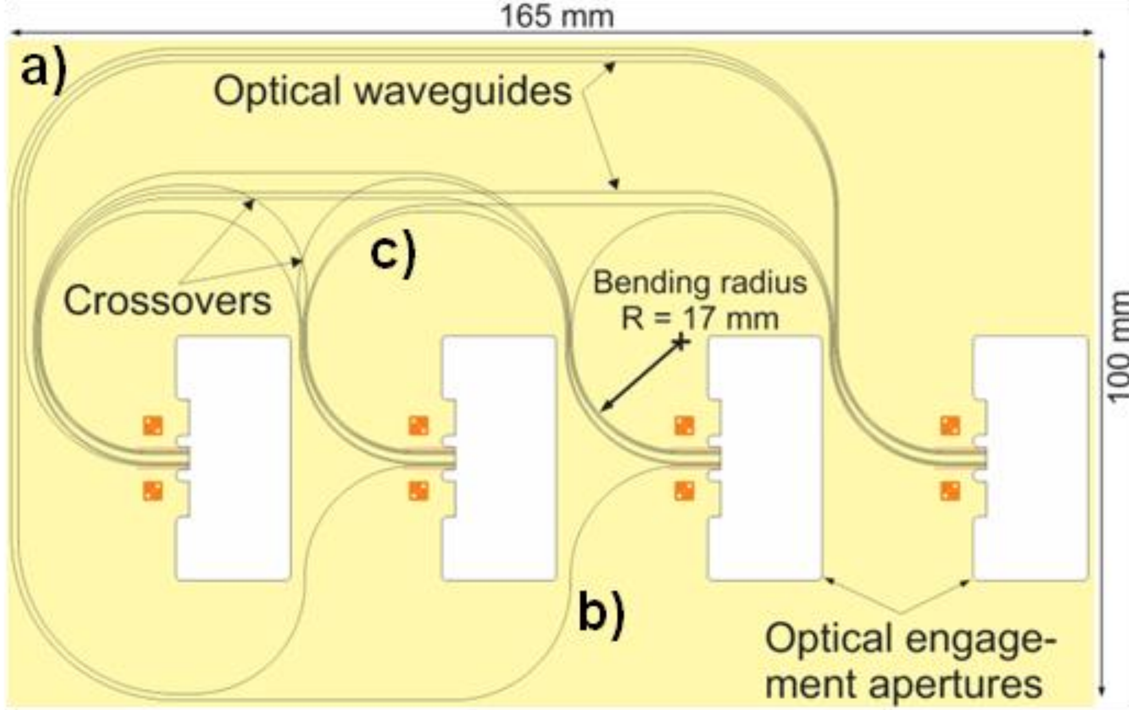


**Figure 4-19: a) Waveguide routing pattern with four optical engagement apertures for the FirstLight demonstrator, b) Negative cascading waveguide, c) Positive cascading waveguide**

A negative cascade is defined as one that occurs when one 90º bend is followed by another 90º bend, which curves in the opposite direction to the first bend, giving rise to an inflection point in the waveguide and a net waveguide angle change of 0º. A positive cascade occurs when one 90º bend is followed by another 90º bend, curving in the same direction as the first and thus giving rise to a net waveguide angle change of 180º. All, but two of the waveguides had four cascaded 90º bends comprising a negative cascade followed by a positive cascade. To minimise bend losses, a radius of curvature of 17 mm [152] was applied on all bends as this was the maximum permitted by the routing constraints on the board. A number of waveguides intersect in one or more positions to accommodate space restrictions. The crossing angles chosen range from 130º to 160º and the measured optical losses are 0.03 dB to 0.08 dB per crossing (Figure 4-20).

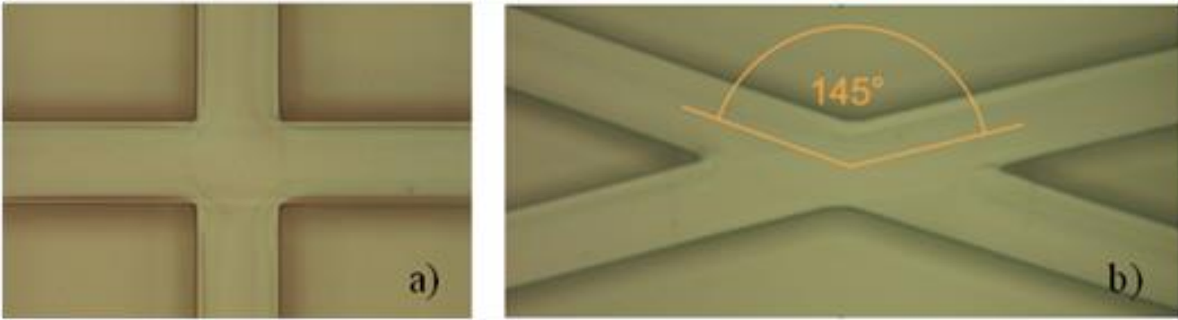


**Figure 4-20: Top view micrographs showing two waveguides intersecting with (a) a crossing angle of 90$^o$ and (b) a crossing angle of 145$^o$ (Source: IBM Research)**

The design of the optical layout was guided by optical insertion loss predictions, based on previous experimental measurements [152] carried out on photolithographically fabricated polymer acrylate test waveguide [120] samples. Optical bend loss over 90° waveguide bends, with radii higher than 13 mm has been previously measured to be at a consistent low value [151], therefore the minimum bend radii were chosen to be higher than 13 mm in order to take polymer and waveguide fabrication into account. The power budget of each optical waveguide link depends on the optical transmitter output power and on the sensitivity of the receiver used in the system. A receiver sensitivity of -11.1 dBm [153], matching the photodiodes and receiver circuitry designed for the active connector, was selected as the threshold for receiving error free signals at 10.3 Gb/s for the system design. The output power of the optical transmitters deployed in the system was -1.48 dBm, which limited the maximum tolerable insertion loss in each optical link to a power budget of 9.62 dB.

Though the relatively large minimum bend radii required at this stage would place significant routing constraints on future optical printed circuit board layouts, these could be effectively mitigated by refinement of manufacturing techniques or novel structuring of the waveguide to reduce bend loss as demonstrated by Xyratex [154].

### 4.3.3 Electro-optical PCB fabrication process

The electro-optical backplane fabricated by Varioprint and IBM Research-Zürich was built up of ten copper layers and one polymeric layer. Preliminary tests on the electrical layers showed significant thickness variation on the complete stack-up of the PCB. Experimental characterisation of surface height variation of the optical substrate on the PCB was carried out by IBM Research-Zürich through surface roughness scans at various locations including an area containing copper fiducials, which served as alignment reference features for high

precision positioning of the waveguides onto the PCB (Figure 4-21a). Scan 1 was taken on the surface near the edge of the optical substrate with no copper features underneath and yielded a measured surface height variation of ±12 µm across the PCB. Scan 2 was taken on the surface along the centre line of the substrate, which traversed some copper features and yielded a measured surface height variation of up to ±50 µm across the PCB (Figure 4-21b).

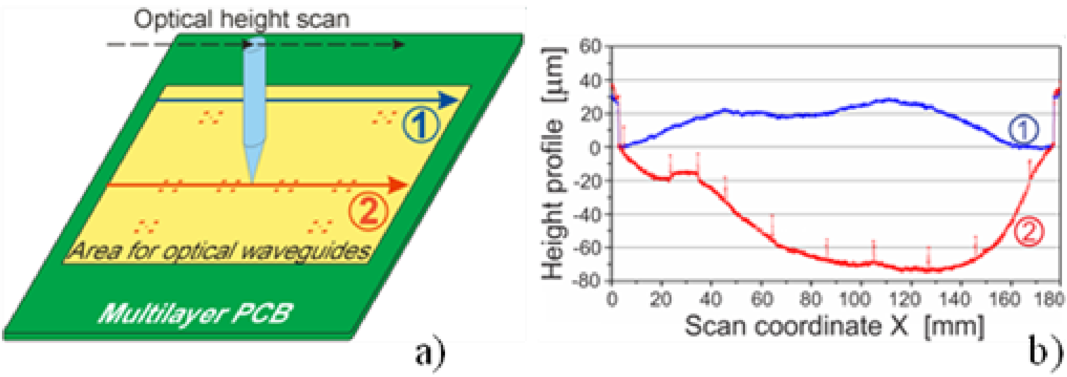


**Figure 4-21: a) Schematic of optical surface scan performed on an OPCB containing up to ten laminated copper layers. b) Experimental results of optical surface scan showing height variations up to ±50 µm. (Source: IBM Research – Zürich)**

If the optical layer were to have been deposited directly onto the electrical PCB surface by using the “doctor blade” method shown in Figure 4-22, the substrate surface variation would only be reduced by the factor of 0.6. This is due to the fact that the thickness of the polymer layer deposited by the doctor-blading process is 0.60 of the blade gap (the distance from the blade tip to the substrate). If the optical lower cladding layer were deposited over those features, the variations of lower cladding thickness would theoretically be reduced to ±30 µm (±50 µm × 0.6) while the core thickness variation would be reduced to ±18 µm (±30 µm × 0.6) accordingly.

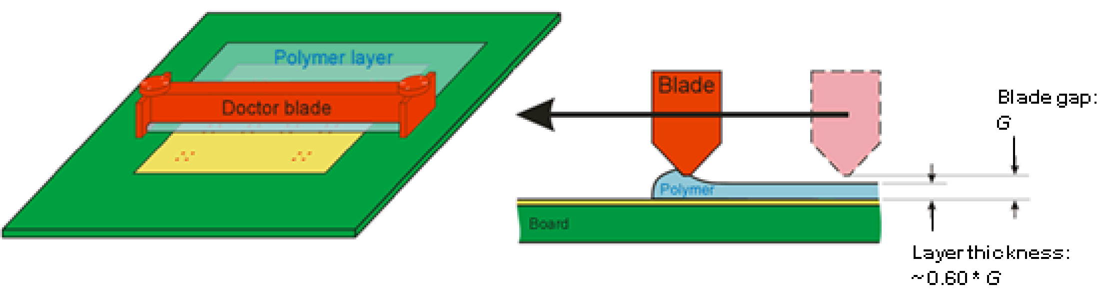


**Figure 4-22: Schematic depiction of polymer layer deposition using doctor blade method. (Source: IBM Research – Zürich)**

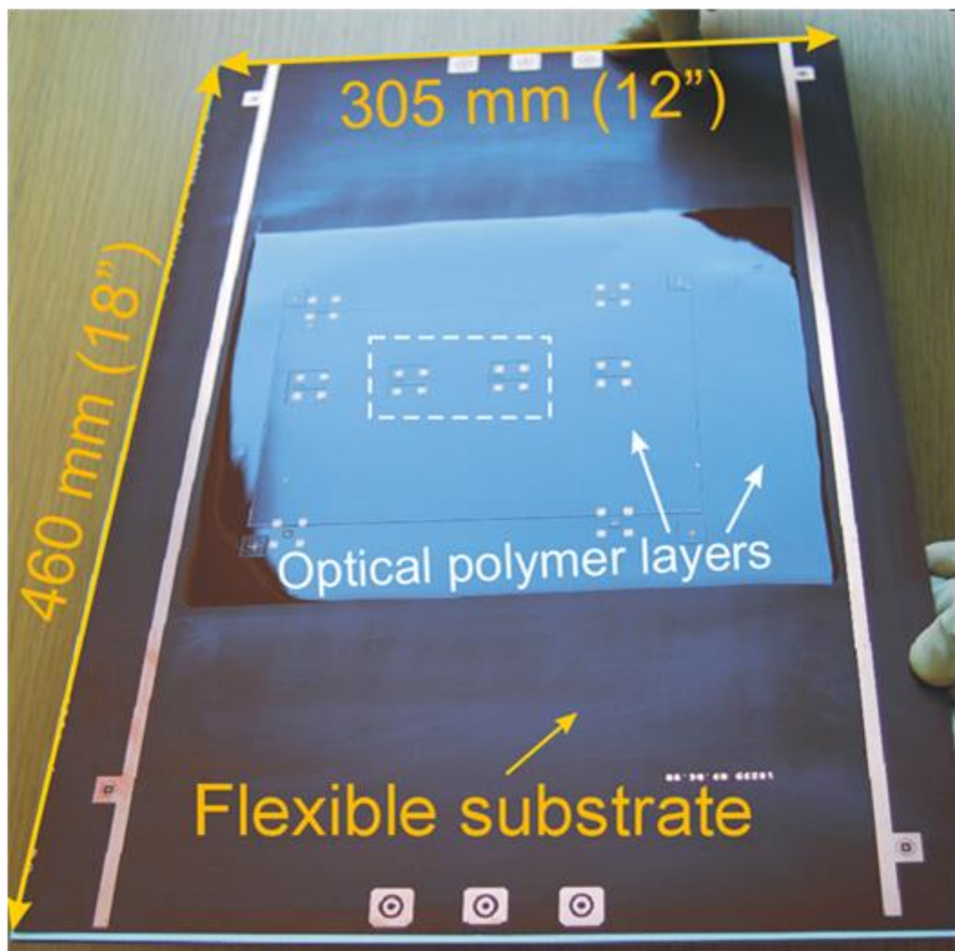


**Figure 4-23: Separate flexible polyimide substrate with optical polymer layers deposited and patterned. (Source: IBM Research – Zürich)**

It was consequently decided by IBM Research-Zürich that the optical layer should be fabricated on a separate flat flexible substrate which would later be laminated to the electrical PCB (Figure 4-23) in order to minimise the thickness variation of the optical waveguide features.

A proprietary liquid cladding polymer was deposited onto a 100 μm thick polyimide substrate with a doctor-blading process applied to control the thickness. The photosensitive polymer was uniformly cured with collimated ultraviolet light (365 nm) from a Mercury/Xenon (Hg/Xe) lamp to polymerise and cure the 100 μm thick lower cladding layer. A higher refractive index liquid core polymer was then deposited onto the lower cladding layer and doctor-bladed to a thickness of 70 μm. The core features were patterned using the laser direct imaging (LDI) technique whereby the beam of a GaN ultraviolet laser diode operating in continuous wave mode at 372 nm wavelength was moved across the substrate to selectively cure those parts of the core layer, which would form the waveguides (see Figure 4-24a). By means of a subsequent wet-chemical development process

step, the non-exposed parts were then removed. As the optical layers were fabricated on the comparatively smoother polyimide substrate, the waveguide core thickness variation was reduced to within ± 2 µm (Figure 4-24b).

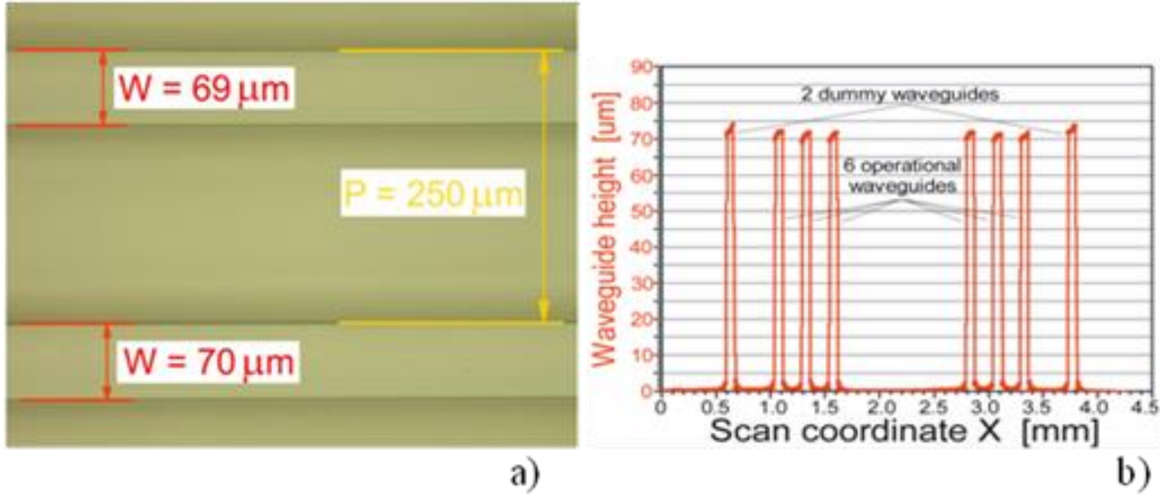


**Figure 4-24: a) Waveguides measured by microscopic top view fulfil width specification of 70 µm within ± 2 µm. b) Waveguide heights experimentally determined by optical surface scanning satisfy target height of 70 µm within ± 2 µm. (Source: IBM Research – Zürich)**

The electrical backplane was fabricated separately (Figure 4-25). Mechanical slots were milled into the polyimide backed optical layer and compliant pins assembled onto the electrical PCB in order to align both substrates together prior to a cold lamination process. Finally the optical engagement apertures were milled out within the outline of the backplane. The backplane was 262 mm long, 240 mm high and 4 mm thick (Figure 4-26).

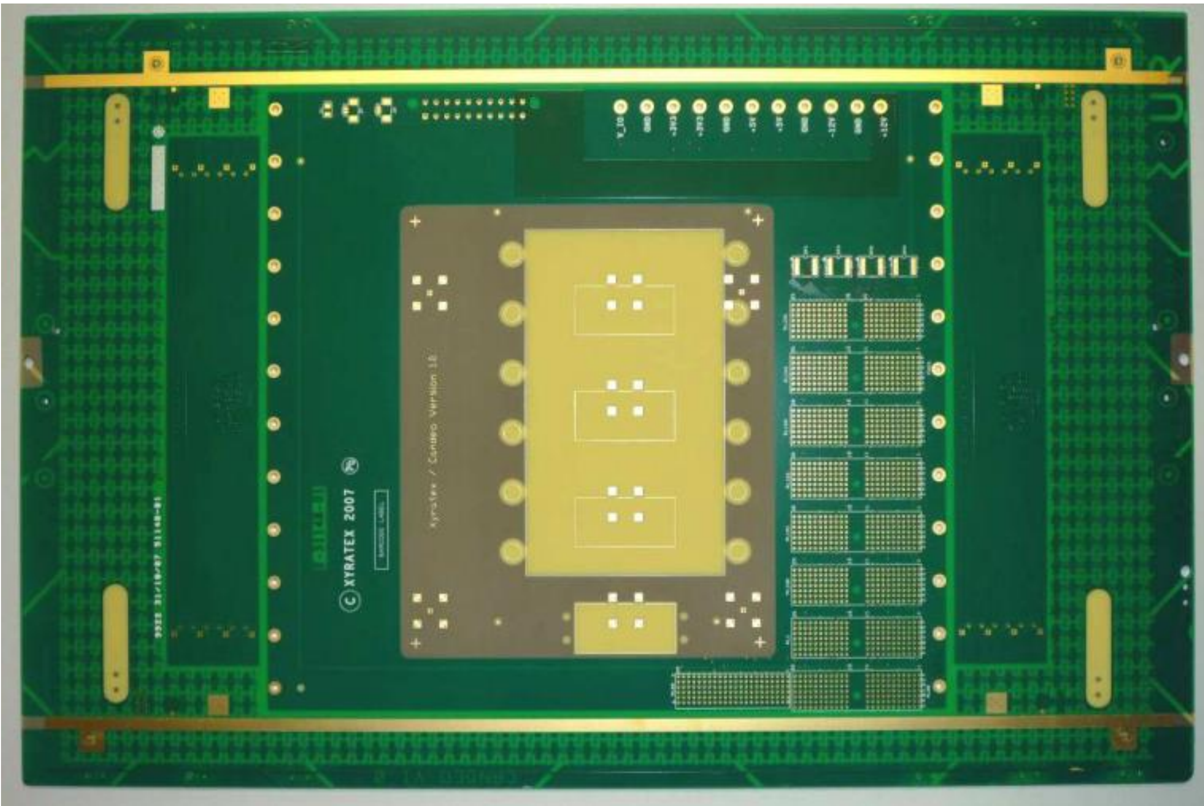

**Figure 4-25: StorConnOpt3 electro-optical backplane prior to slots being milled out**

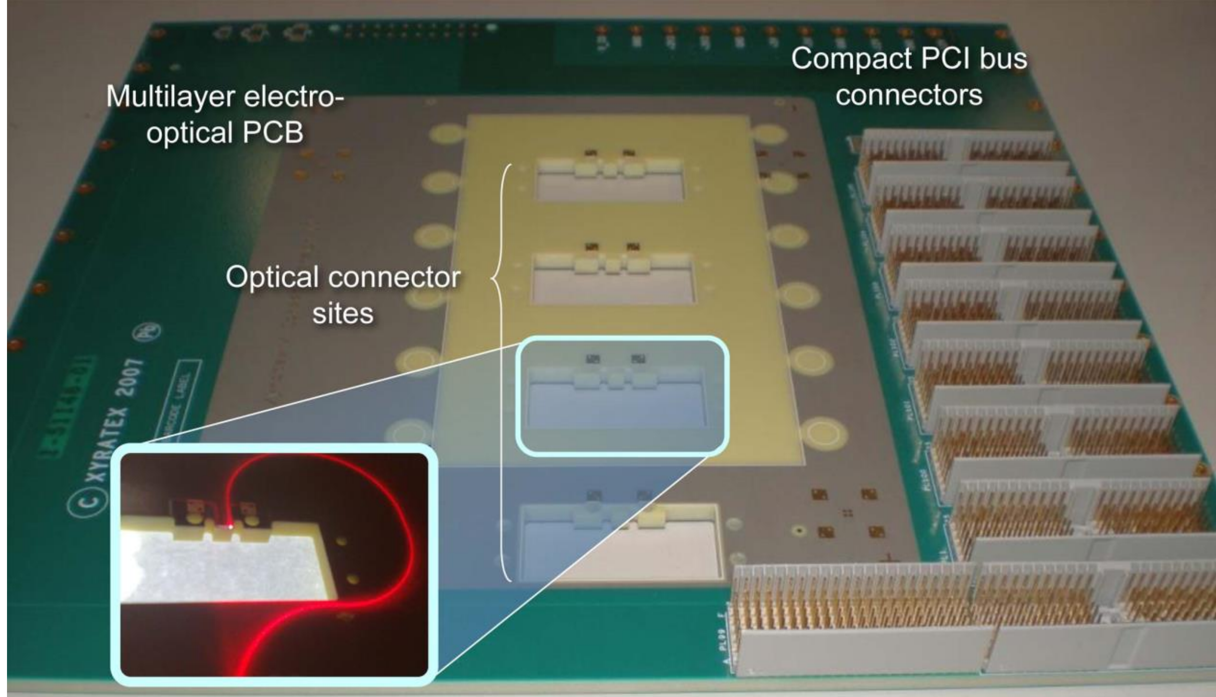


**Figure 4-26: StorConnOpt3 electro-optical backplane with optical connector slots milled out and electronic CompactPCI connectors populated. A close-up view of a connector aperture is shown with a single curved waveguide and its egress point illuminated with 635 nm visible light.**

### 4.3.4 Insertion loss measurements

The measurements for optical insertion loss of the optical interconnects were conducted jointly by Richard Pitwon and Kai Wang at UCL. An ST connector packaged 850 nm VCSEL was connected to a standard 50/125 µm step-index MM fibre with $NA_{fibre}$ of 0.22. The fibre core cross section and its NA were both smaller than the waveguide core cross section of 70 µm × 70 µm and the waveguide $NA_{waveguide}$ of 0.33 respectively. This reduced the fibre coupling loss during the butt coupling measurement. The launch fibre was 10 m long and was wound 20 times around a 38 mm diameter circular mandrel in order to maximise the distribution of optical power across the fibre modes and provide a worst case mode-filled near-field and far-field power distribution at the fibre launch facet (Figure 4-27).

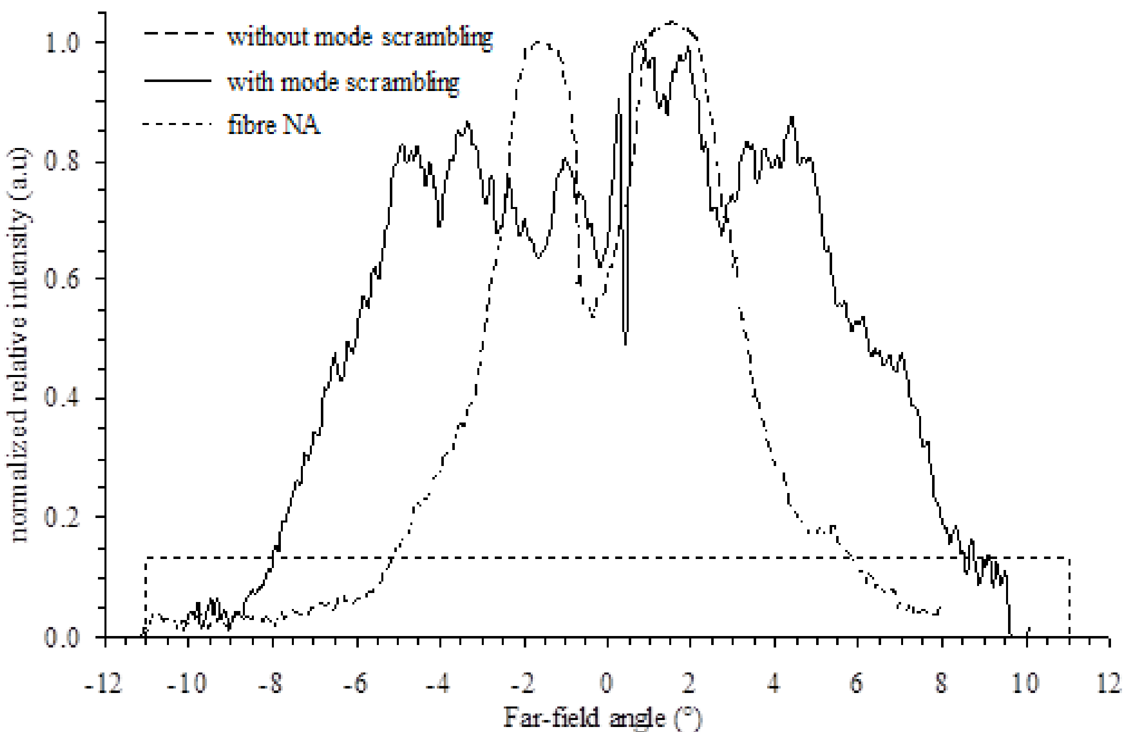


**Figure 4-27: Far field pattern of launch fibre with and without mode filtering (Source: UCL)**

The fibre was then butt coupled to the waveguide input facet.

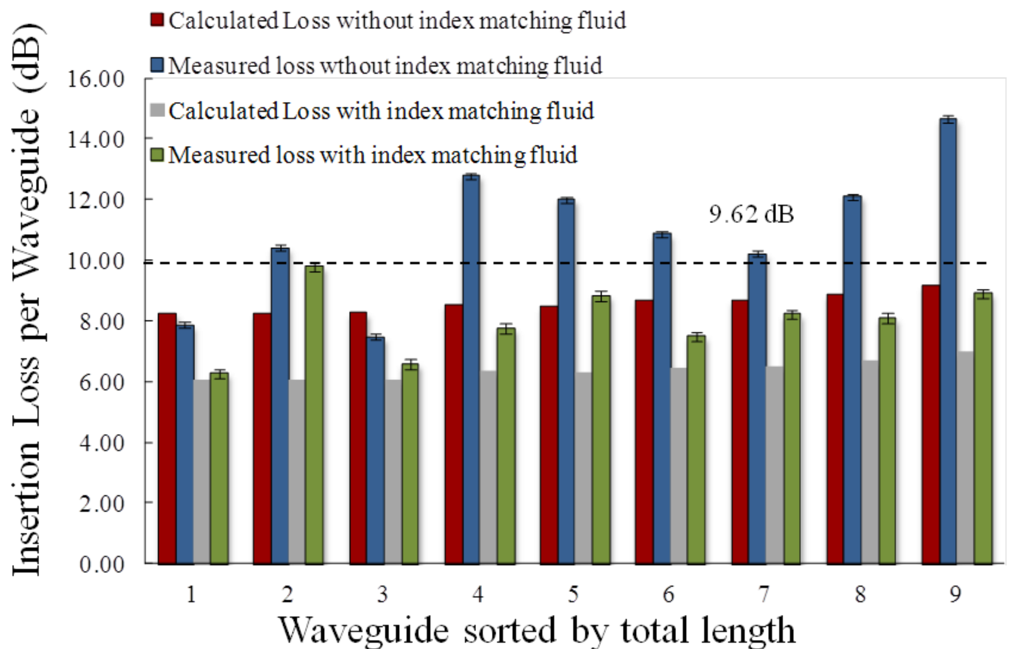


**Figure 4-28: Predicted and measured values for insertion loss on a group of nine waveguides on the OPCB. The dashed line shows the 9.62 dB threshold**

The optical power at the launch facet of the fibre was measured to be -1.48 ± 0.02 dBm. The launch fibre was mounted on a group of motorised translation stages with sub-micron step resolution in three axes, x, y and z, to accurately align the fibre to the waveguide and to optimise coupling into the waveguide. A thin silicon photodetector with an 8 mm aperture was required to fit through the waveguide interface engagement aperture of the StorConnOpt3 backplane and aligned to the output facet of the waveguide in order to capture the light received through the waveguide and measure the waveguide insertion loss (Figure 4-29).

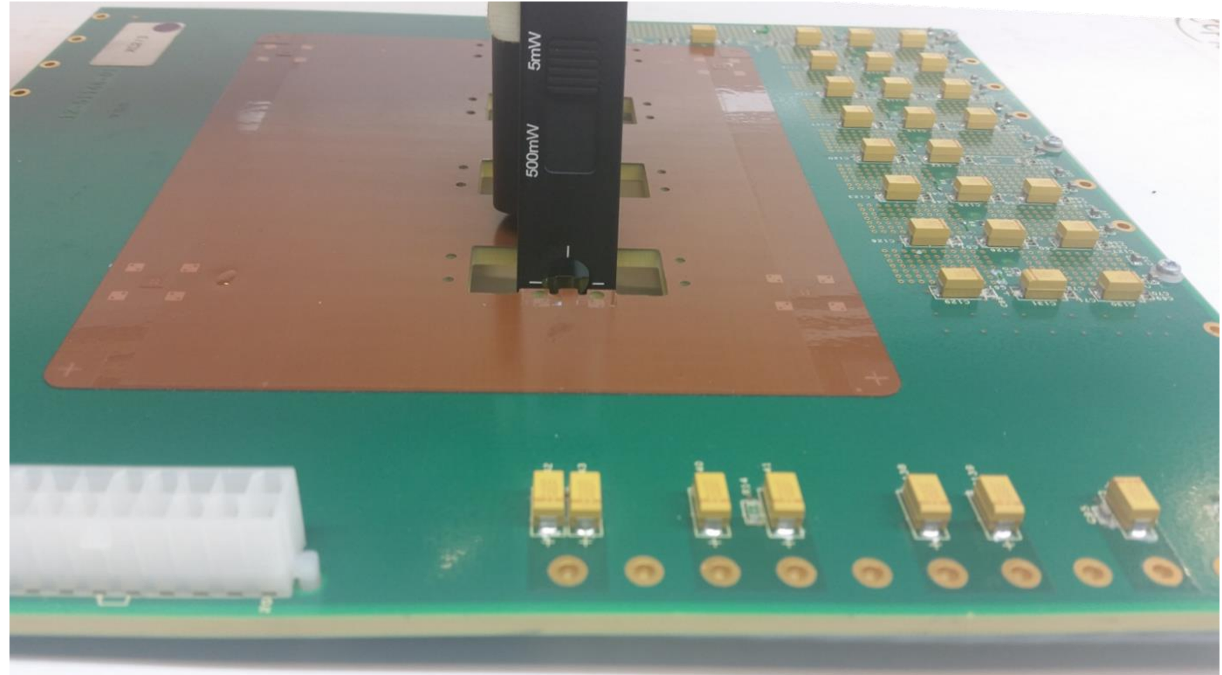


**Figure 4-29: Use of large area thin silicon photodetector to fit into the mid board waveguide connector interface slots in order to directly measure optical output from waveguides**

The detector was calibrated against the light condition in the dark room where the measurements were conducted. The loss measured for each waveguide included the coupling losses at the fibre-waveguide interface and waveguide-PD interface, absorption by the propagation medium, bend transition losses due to modal mismatch between waveguide segments of differing radii of curvature and losses incurred at crossing junctions. Though the FirstLight board contained 12 waveguides in total, the longest four waveguides were excluded from experimental characterisation due to damage during assembly. Insertion loss measurements were made on the remaining eight waveguides both with and without index matching fluid applied. The measurement results are shown in Figure 4-28 along with the original calculated values and highlight the strong dependence of loss on end facet roughness. The waveguides in the optical interconnect layout were designed to never exceed an insertion loss of 9.62 dB, the threshold required to achieve communication at a bit error rate of less than $10^{-12}$. Figure 4-28 includes the comparison of the calculated optical losses without index matching fluid (red column) and the measured insertion losses of waveguides without index matching fluid applied (blue column), which are on average 2.36 dB higher than the calculated predictions. The insertion losses of certain waveguides, i.e. 2, 4, 5, 6, 7, 8 and 9 are in excess of the 9.62 dB error free threshold. The higher insertion losses are partially due to the higher scattering losses at the waveguide end facets, which depend on the surface roughness of the end facet in question. The surface roughness in turn depends on how the end facets were cut and polished. In this case the

waveguide end facets are located at the edges of the milled midboard optical engagement apertures. Though the end facet roughness of these waveguides could not be measured without damaging the OPCB, waveguides cut using similar milling techniques to those deployed on this board exhibit RMS surface roughness values ranging from 183 nm to 350 nm [155]. The grey columns in Figure 4-28 represent the calculated insertion losses with index matching fluid, the values for which are 2.40 dB lower than the waveguide losses calculated without index matching fluid applied. The green columns in Figure 4-28 represent the measured insertion losses with index matching fluid applied. The measured insertion losses on the eight waveguides were on average 2.93 dB lower than those measured without applying the index matching fluid (blue columns). The index matching fluid, however, is not a practical means of reducing the roughness of waveguide end facets in applications involving repeatable connection to the waveguides as the fluid tends to dry out gradually and may accumulate dust, which in turn will cause the end facet scattering losses to increase in an unpredictable manner. Therefore a more durable method using waveguide core polymer to smooth the end facets was developed and demonstrated to permanently reduce the end facet roughness. This process is shown in Figure 4-30.

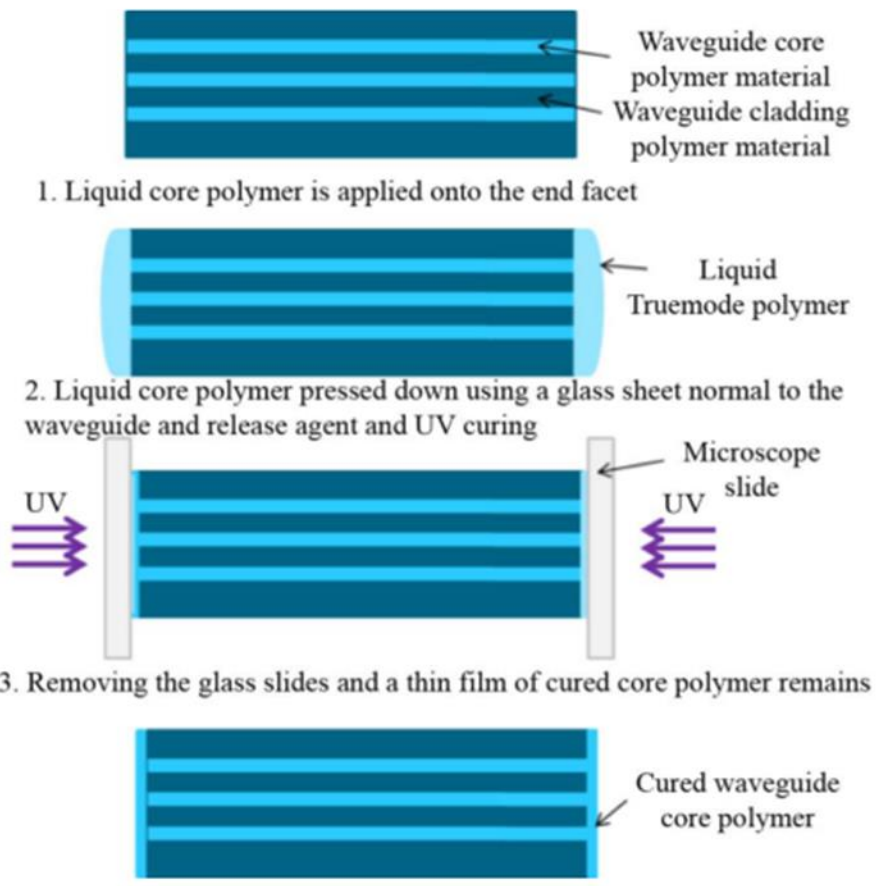


**Figure 4-30: Process to deposit smooth thin film over waveguide end facets in order to reduce coupling loss** [155]

The deployment of this technique, reported by Bagshiahi et al.[155], improved waveguide transmissivity by 0.49 dB on average compared to waveguides with index matching fluid applied.

Figure 4-31 shows a view of the OPCB backplane in the laboratory with a waveguide illuminated with 650 nm visible light. In this set-up, shown on the left hand side, the MT interface of a fibre-optic test jumper is aligned actively such as to butt-coupled directly to the unpopulated waveguide interface.

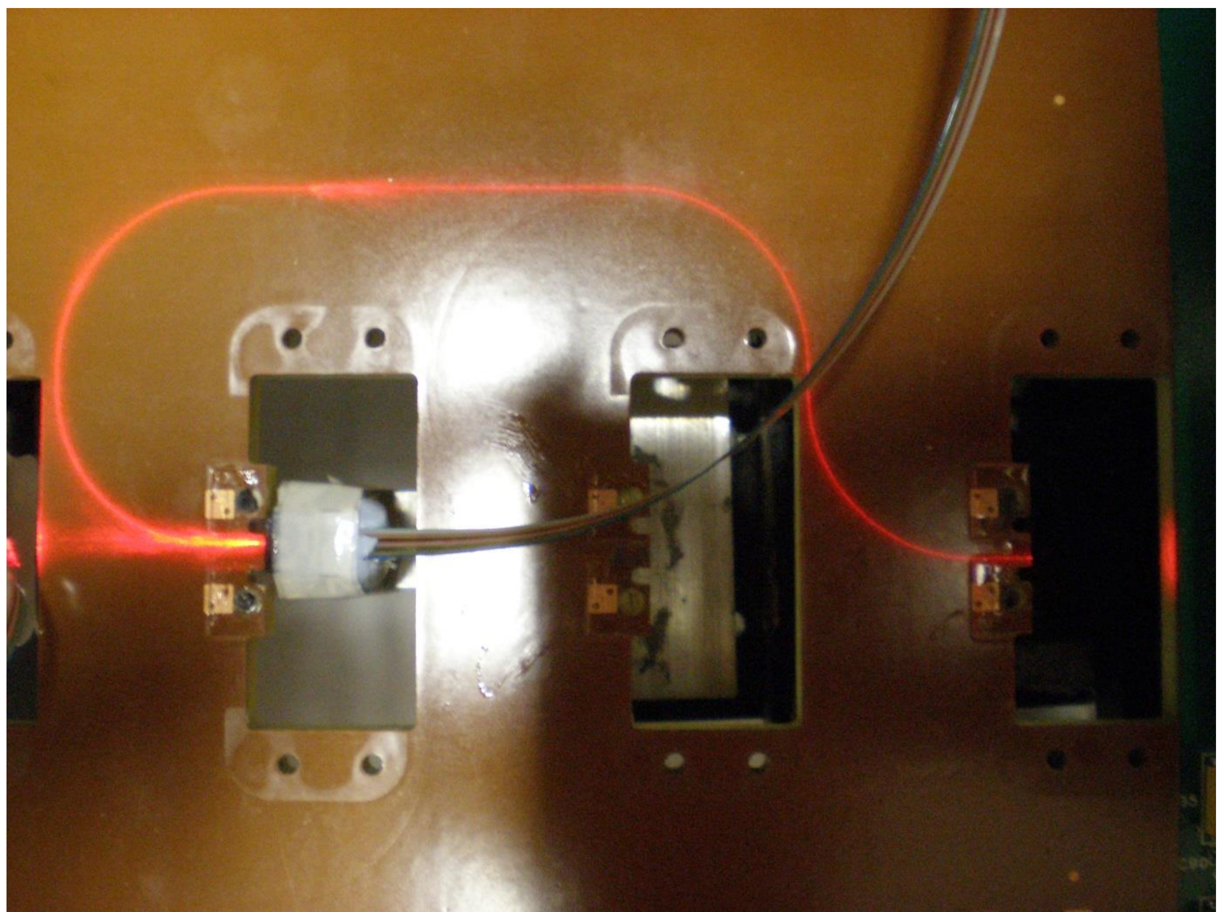

**Figure 4-31: Top view of OPCB backplane with one waveguide illuminated with 650 nm light**

### 4.3.5 Optical waveguide signal integrity characterisation

Bit Error Rate (BER) tests were carried out across the eight waveguides directly using an Anritsu signal analyzer MT1810 and a SPF+ driver and receiver unit and applied index matching fluid at both the launch and exit facets of the waveguides. The SFP+ unit [153] had the same sensitivity as the XFP model employed in the system, but excluded the clock recovery unit, so that the quality of the raw signal could be measured. A PRBS (Pseudorandom binary sequence) $2^{31}$-1 pattern length Ethernet LAN traffic 10.3125 Gbit/s bit rate was generated by the MT1810 and was used to modulate the SFP+ transceiver. The optical signal was guided by a MM fibre,

which was butt coupled to a waveguide channel on the prototype backplane. The output from the waveguide was captured by another MM fibre, which was connected to the receiver port of the SFP+ unit on the BERT. We were able to test individual channels with this arrangement. An error rate of less than $10^{-12}$ was achieved through each of the eight waveguides under test.

## 4.4 Passive alignment and assembly method

One crucial requirement for the commercial deployment of optical PCB technology is a low-cost technique for the high-yield assembly of optical interface components onto the optical layers. In order to enable high volume assembly it is preferable that such techniques be passive and repeatable. A proprietary fabrication technique and method of passively aligning and assembling parallel optical microlenses to embedded polymer waveguide arrays was successfully developed [24,30,31]. These form a critical part of the pluggable in-plane connection interface between arbitrary external optical devices, either passive or active, and a PCB embedded optical circuit.

### 4.4.1 Improvements over Storlite design

Unlike the Storlite connector interface, on which the waveguide array was offset from the MT pin datum, the waveguide array on the FirstLight version was in line with the MT pin datum thus increasing optical interface stability against pivoting effects.

As before the MT pin slots were required to be incorporated into a precision moulded mount (Figure 4-32a), which was self-aligned to the waveguides with very high accuracy using registration features in the waveguide core layer situated at either side of the waveguide array. In order to accommodate the need for the waveguide array to be inline with the MT pin datum, slots must be machined in the PCB either side of the waveguide array to accommodate the “feet” of the MT mount (Figure 4-32b).

The challenge to the PCB fabricator is that due to manufacturing tolerances the registration stubs must be reduced in size and brought closer to the end waveguides to allow the registration wings of the waveguide receptacle to engage both laterally and have enough room between the outer edges of the registration features and the edge of the PCB recess to engage vertically. A method of further improving robustness of the precision alignment method in future is outlined in section 7.2.4.

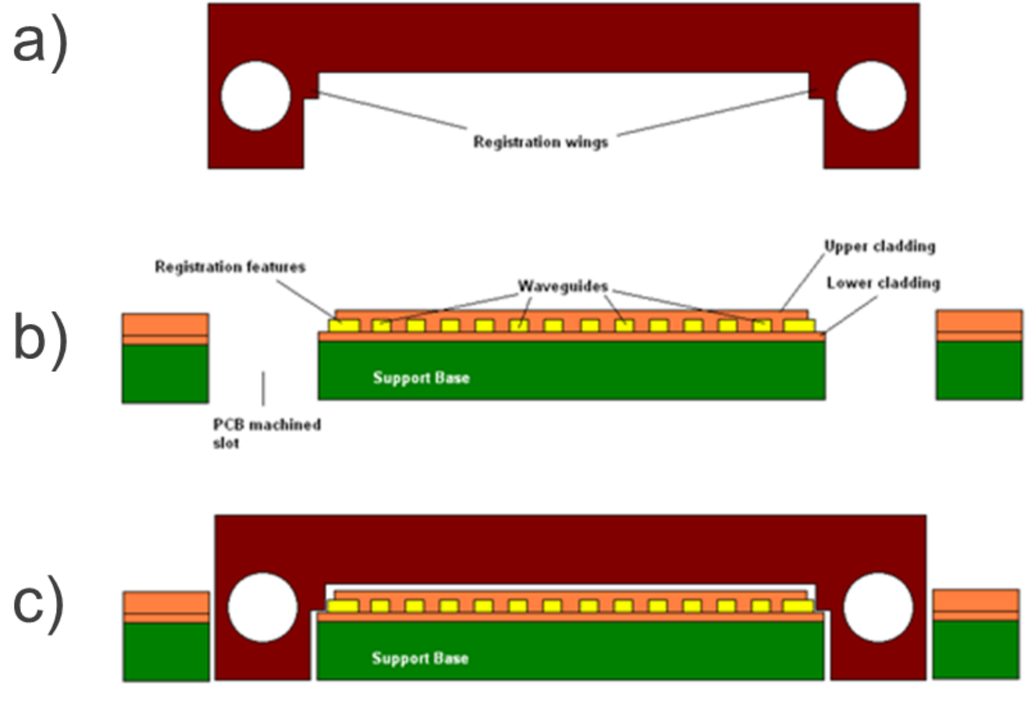


**Figure 4-32: Requirements specification for FirstLight waveguide receptacle: a) waveguide receptacle with 1x12 MT compliant interface, b) waveguide array with self-alignment features and PCB machined recesses, c) waveguide receptacle passively aligned over waveguides using self-alignment features**

### 4.4.2 Fabrication of passive alignment features

The complete fabrication process for the passive alignment features on the optical layer is outlined in Figure 4-33. The procedure involves the fabrication of passive mechanical registration features in the core layer during the same process step in which the waveguide cores themselves are patterned. Effectively these serve as additional "dummy" waveguides, which are positioned on either side of the signal waveguides and as a result their positional accuracy with respect to the signal waveguides is as high as those of the signal waveguides to each other. Instead of uniformly curing the upper cladding however, it must be selectively cured to ensure that the central signal waveguides are completely clad while the registration waveguides are not. The notches either side of the waveguide interface were milled after the waveguides were fabricated using a visual alignment system.

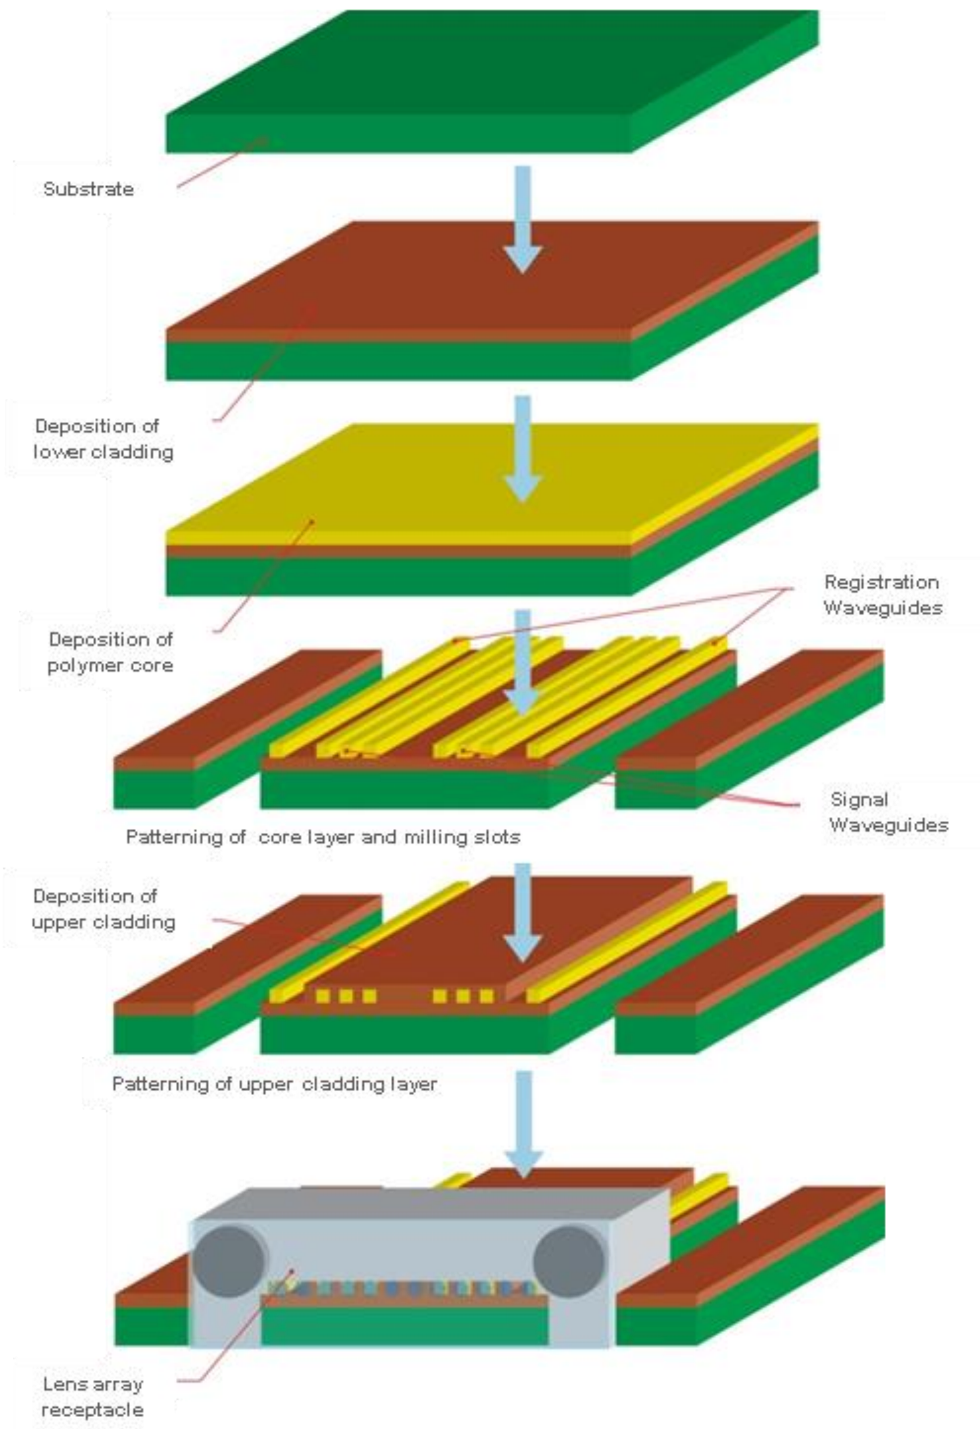


**Figure 4-33: Fabrication process for the passive alignment features on the optical layer**

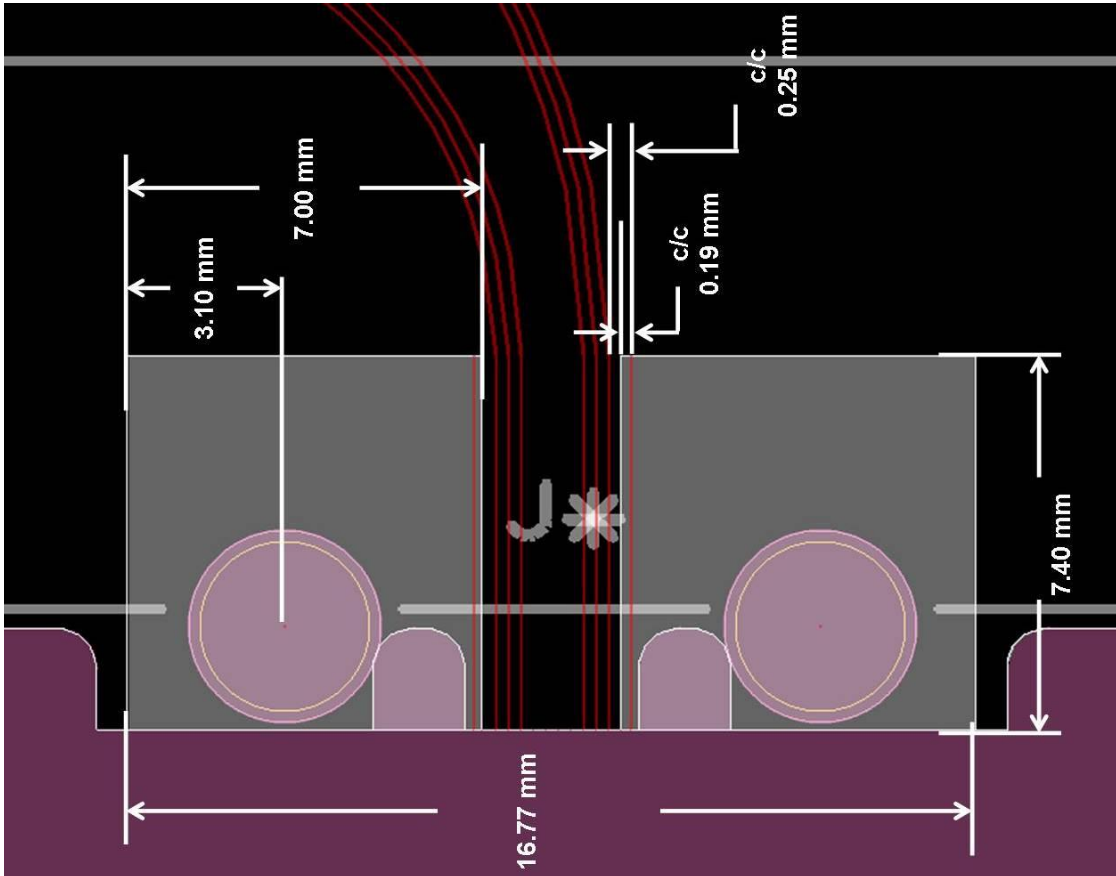


**Figure 4-34: View of the PCB layout of the FirstLight waveguide interface using the Cadence Allegro PCB Layout software tool.**

Figure 4-34 shows a view of the PCB layout for the FirstLight electro-optical backplane rendered using the Cadence Allegro PCB Layout software tool by Kai Wang of UCL. The sections are shown (light grey) are those areas in which the upper cladding must be removed i.e. not cured to provide mechanical access to the mechanical registration waveguides. Six communication waveguides (two sets of three waveguides) are shown in red between the clearance areas. These will be covered by upper cladding. Two peripheral straight “dummy waveguide” stubs just within the light grey clearance area serve as mechanical registration features.

This clearance allows for direct passive mechanical registration of optical components to the waveguides with very high precision. In addition, the fabrication tolerances required to pattern the upper cladding for this purpose are far lower than those required to pattern the waveguides themselves. It is only important that the outer edges of the registration waveguides, which form the mechanical datum, be mechanically exposed. Preferably, the upper cladding should partially cover the registration waveguides in order to provide structural reinforcement and reduce the risk of the registration waveguides delaminating under the strain. However, this is not strictly necessary and as shown in Figure 4-35 and Figure 4-36 the registration waveguides in the FirstLight OPCB were

left completely uncovered without any adverse effects. This technique can be implemented using most waveguide fabrication processes. The positional tolerance of the mechanical registration features with respect to each other has been measured to be ±3 µm for lateral misalignment in-plane and ± 4 µm normal to the PCB plane [2], [156].

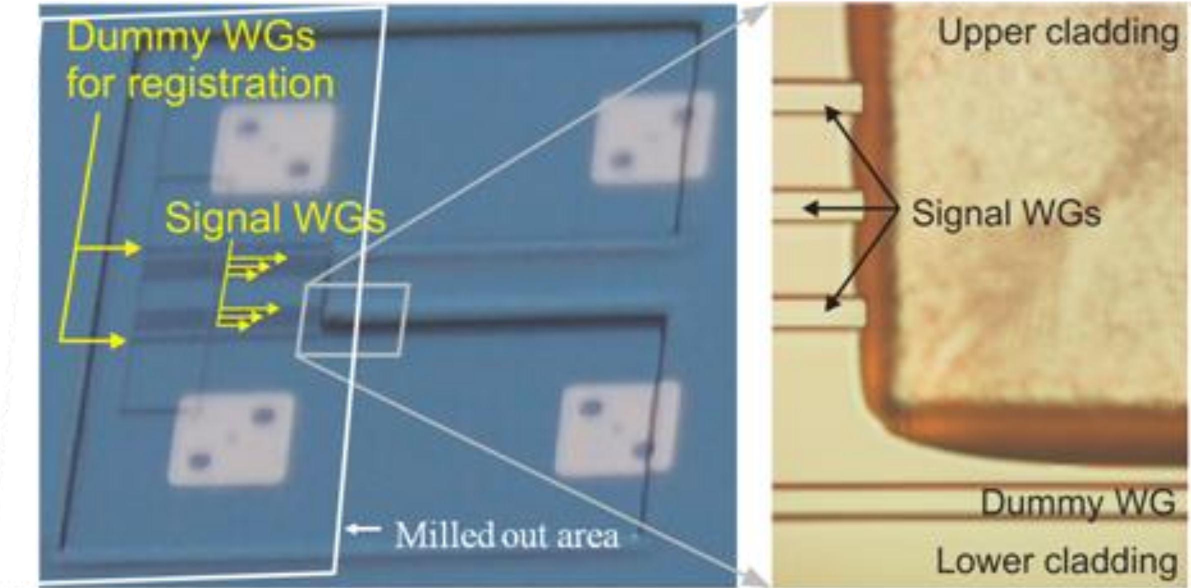


**Figure 4-35: Photographs of upper cladding opening showing waveguides for signal transmission as well as connector alignment features (dummy waveguides). (Source: IBM Research – Zürich)**

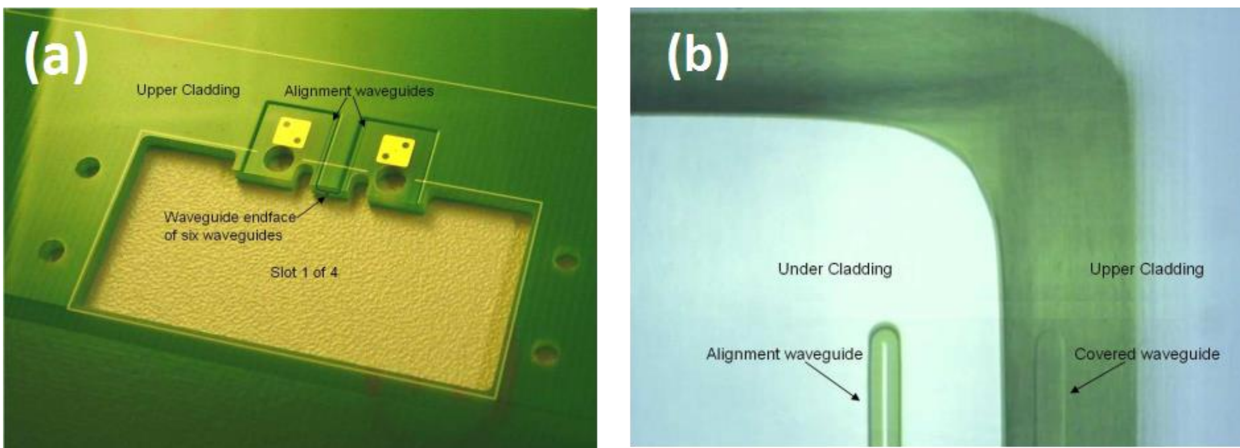


**Figure 4-36: (a) Clearance areas in upper cladding layer to provide mechanical access for optical connector receptacle and waveguide end facet**

**(b) An exposed waveguide structure serves as a passive alignment feature for the assembly of the optical connector receptacle**

**(Source: Varioprint)**

### 4.4.3 Lensed waveguide receptacle design

The author developed the design concept and requirements specification for the lensed connector receptacle, which comprised a custom moulded receptacle and discrete commercial geometric convex microlens array from Omron. The author worked with connector company Samtec to render the final design for the waveguide receptacle, which was fabricated by Samtec, who were partnering with Xyratex at the time.

The custom receptacle included compliant structures to allow it to mechanically engage with the registration waveguides on the board and a recess to accommodate a standard high performance "mechanical transfer" connector (MT) compliant lens array. The MT compliant interface on the lens array included two 0.7 mm pin slots and 12 micro-lenses arranged between the slots on a 250 µm pitch between lenses. In order to ensure that the lens array was accurately aligned within the receptacle, the receptacle also included two MT compliant pin slots, the dimensions of which matched those in the lens array. The lens array was fastened to the receptacle with a UV curable low shrinkage optical adhesive, Dymax OP-21.

The waveguide receptacle was fabricated through an injection moulding process in the commercial polycarbonate material Makrolon® AL2647.

Figure 4-37 shows the structure of the waveguide receptacle in layers while Figure 4-38 shows a dimensioned schematic of the receptacle.

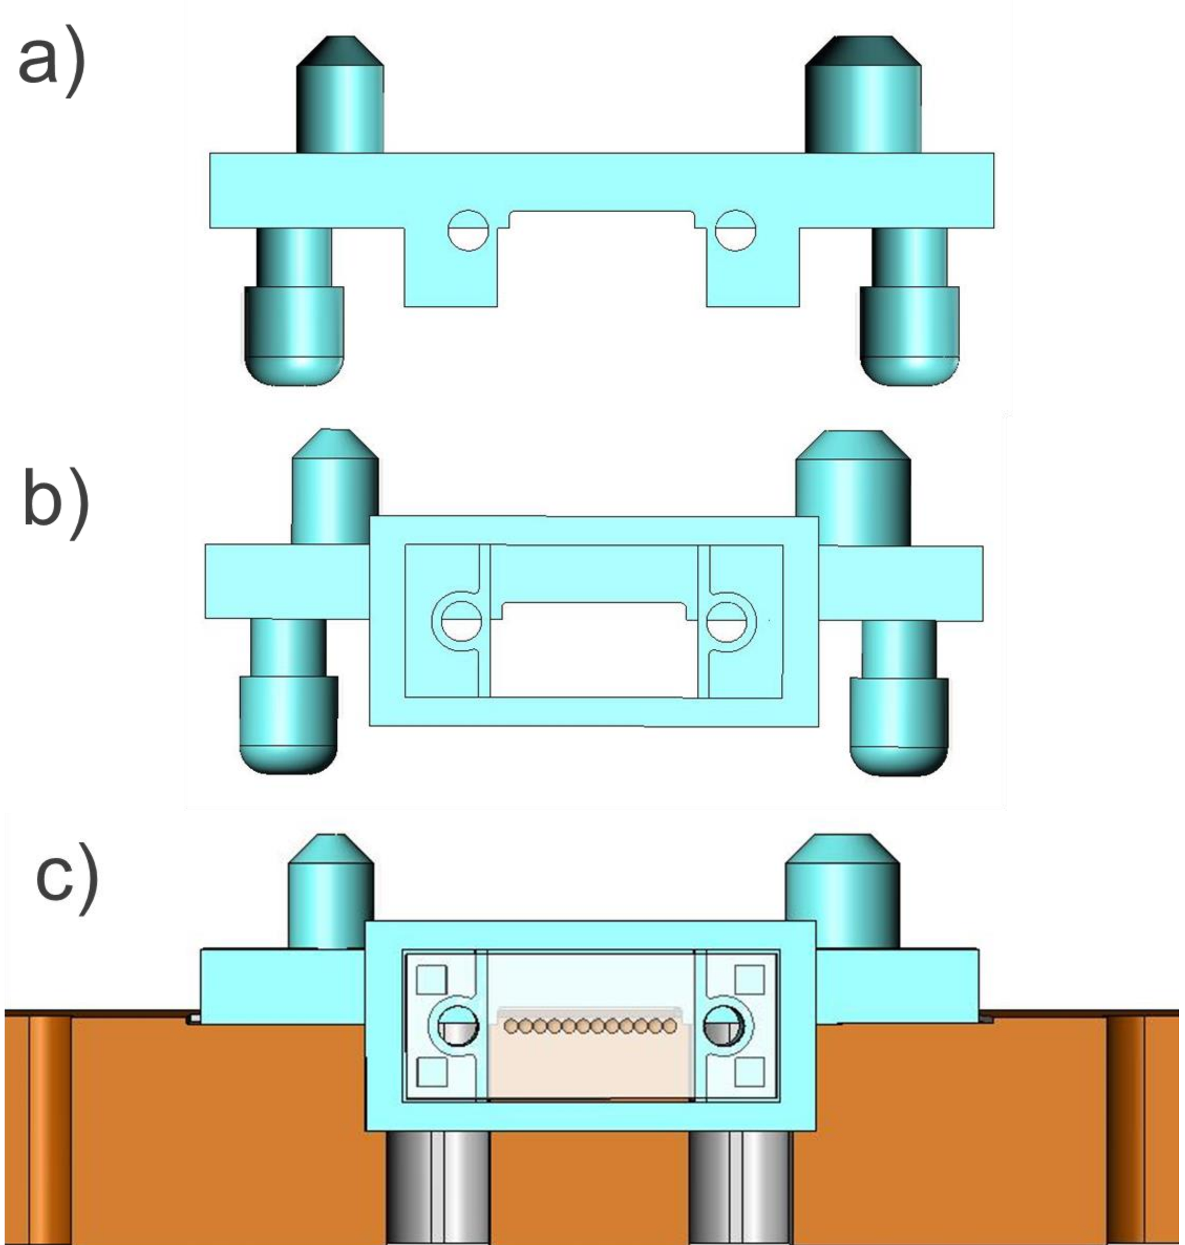


**Figure 4-37: Lensed waveguide receptacle front views: a) trimmed receptacle without lens array holder section, b) complete receptacle, c) complete receptacle with Omron microlens array integrated and receptacle assembled onto OPCB**

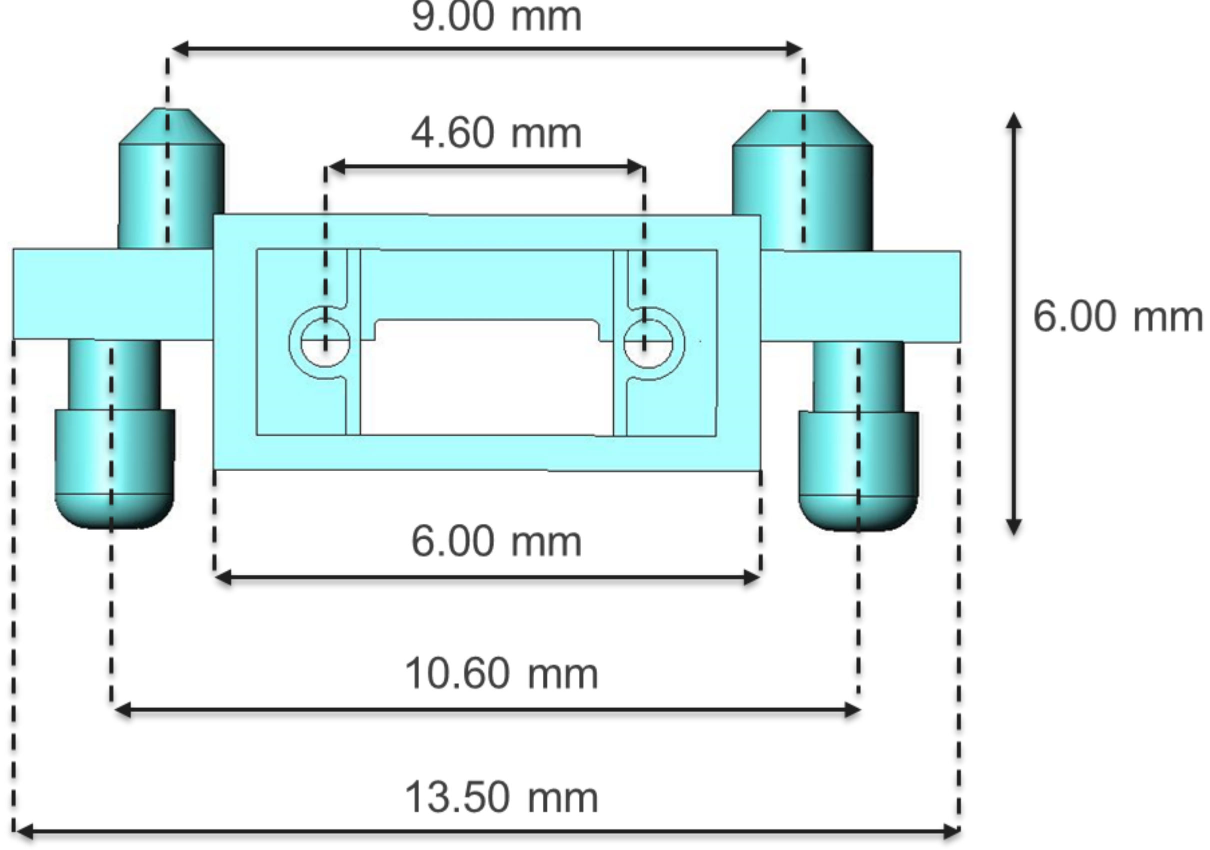


**Figure 4-38: Dimensioned drawing of waveguide receptacle**

Figure 4-39 shows a photo of a lens receptacle placed next to a British ten pence coin.

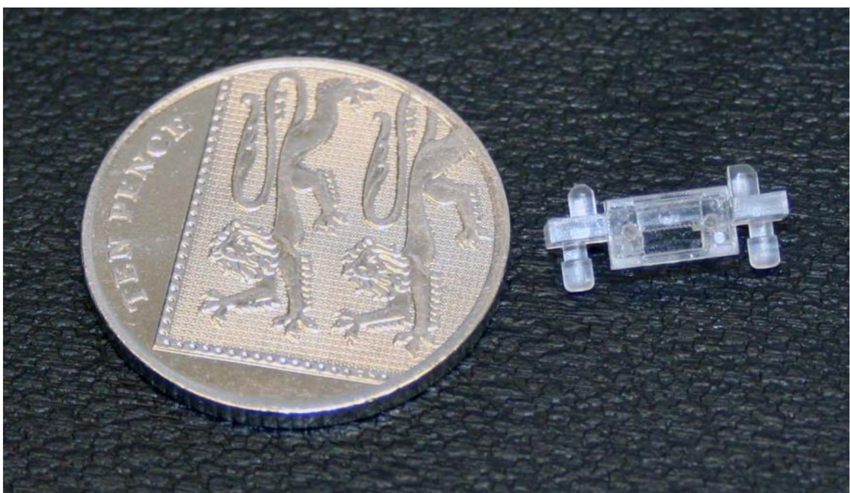


**Figure 4-39: Photo of complete micro-lens receptacle assembly next to British ten pence coin**

### 4.4.4 Assembly of microlens array into receptacle

Figure 4-40 shows a lens assembly jig designed by Samtec to align the lens array into the custom receptacle and hold it tightly in place during the UV curing process. It is critical that the lens be properly aligned to the custom receptacle otherwise it will not be properly aligned to the waveguides, to which the receptacle is fastened.

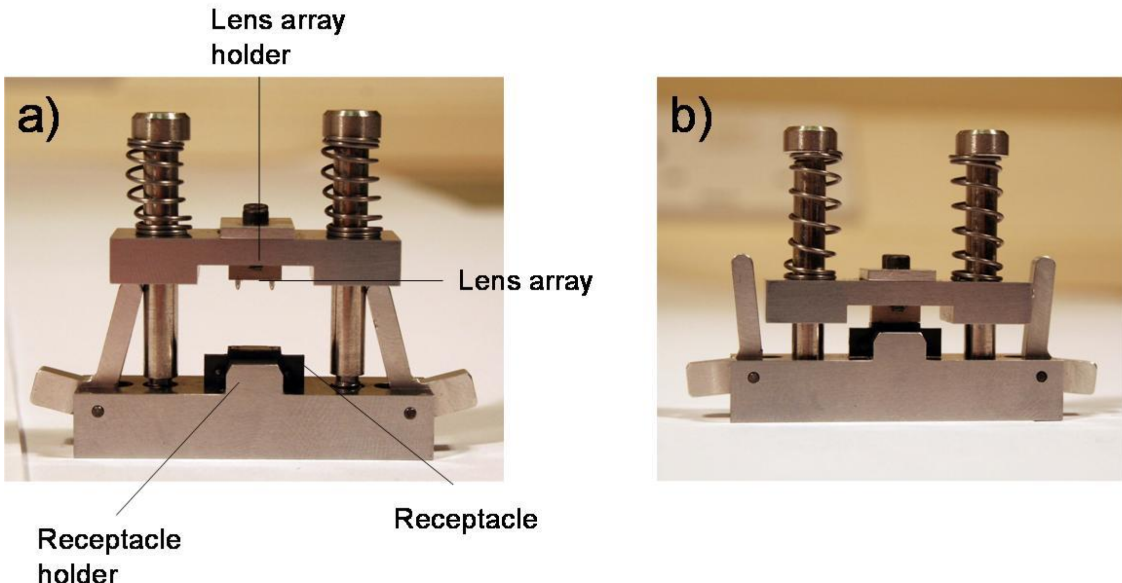


**Figure 4-40: Mechanical lens assembly jig A designed to align the Omron microlens array into the custom receptacle a) open position, b) engaged position**

While the jig is in the open position (Figure 4-40a), the lens array is mounted onto compliant pins in the jig and the receptacle is mounted into a compliant recess. When the jig is closed (Figure 4-40b), the lens array is pressed into the lens receptacle and held under a strong spring tension. A UV source was subsequently applied to cure the adhesive between the lens and receptacle. The lens receptacle included recesses to contain and channel the adhesive away from sensitive areas such as the MT pin holes and the microlenses, while maximizing contact between those areas of the lens plate and receptacle which are neither in the signal path nor the mechanical registration path. It was decided to fasten the lensed receptacle to the optical PCB using a UV curable optical adhesive.

### 4.4.5 Assembly of waveguide receptacle onto OPCB

As alignment of the receptacle to the waveguides is critical, it was important that the receptacle lay flat on the smooth exposed lower cladding and that none of the adhesive seeped underneath it.

Figure 4-41a shows a photo of two connector recesses in StorConnOpt3 electro-optical backplane with waveguide interfaces unpopulated. Figure 4-41b shows the same recesses with a lensed waveguide receptacle passively aligned onto one of the waveguide interfaces. This could be done by hand as the thickness of the core layer was sufficient to feel the registration structures and gently push the receptacle into position.

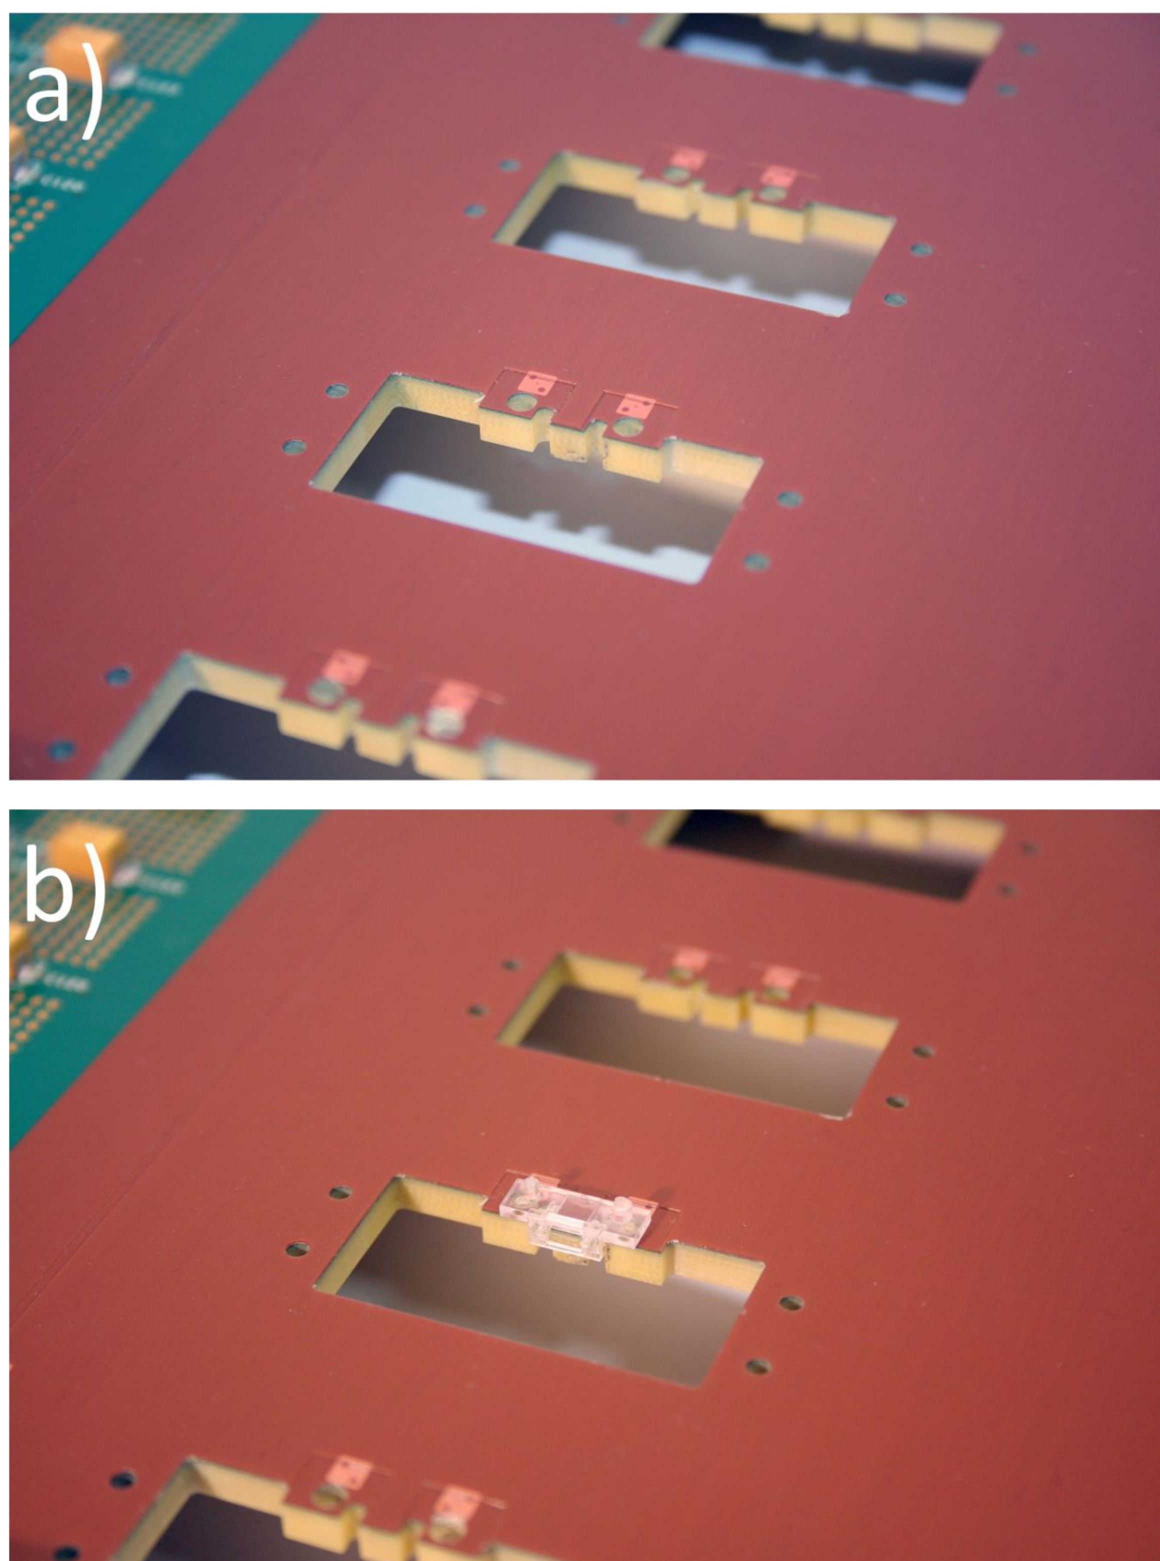


**Figure 4-41: a) OPCB waveguide interface without receptacle, b) OPCB waveguide interface with receptacle passively aligned, but not glued in**

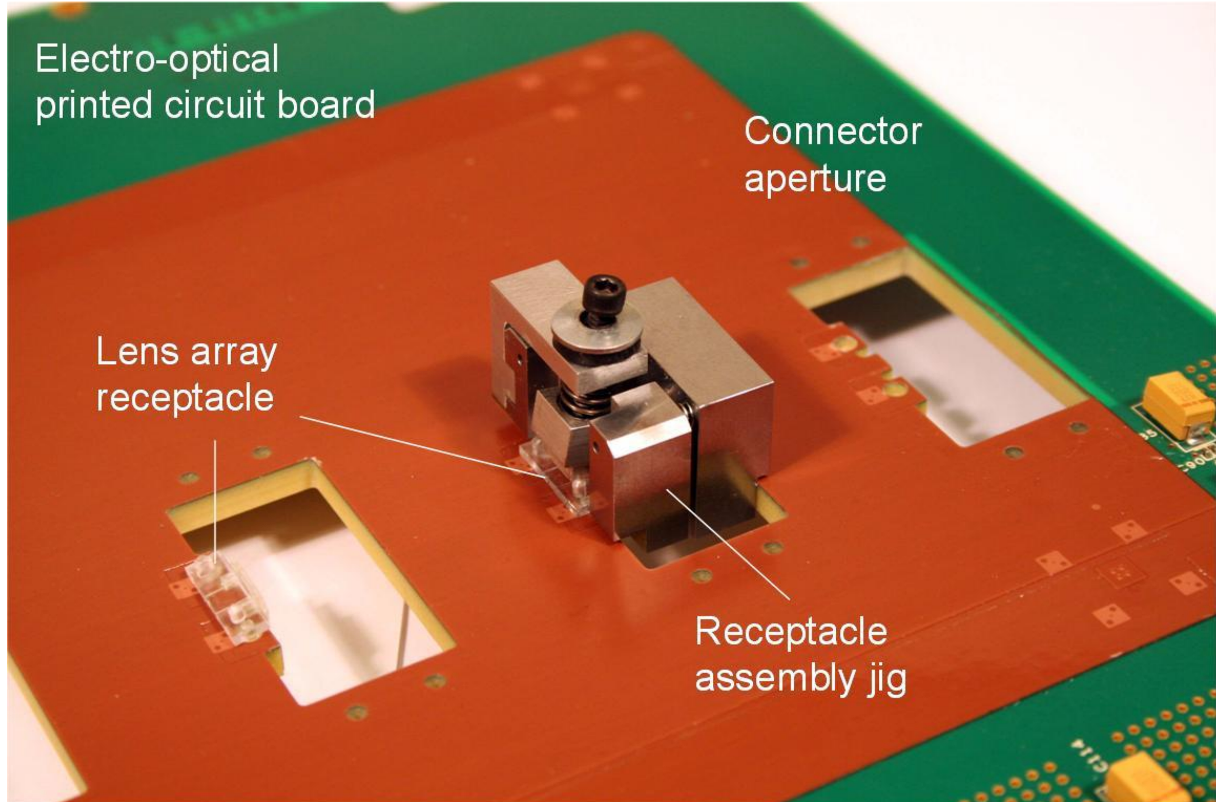


**Figure 4-42: Receptacle assembly jig B required to register lensed receptacle with waveguide alignment features and hold in place during adhesive curing process.**

Figure 4-42 shows a mechanical jig B designed by Samtec to hold the passively aligned lensed receptacle tightly in position onto the optical PCB while the adhesive is applied as fillet bonds around the edges of the receptacle and subsequently cured. As with the lens assembly jig A described above, the receptacle assembly jig B contained recesses to hold the passively aligned lensed receptacle in place over the waveguides under strong spring tension. The shape of jig B was customised to be snap-fit into the connector aperture on the board.

Once the lensed or “secondary” receptacle is assembled onto the OPCB, a larger primary receptacle is aligned accurately to the secondary receptacle by means of alignment stubs protruding from the secondary receptacle and subsequently bolted over the optical engagement aperture. Figure 4-43a shows an optical engagement aperture with just a secondary receptacle assembled and, behind it, one with both a secondary and primary receptacle assembled. The purpose of the primary receptacle is to provide coarse mechanical alignment of the OPCB connector during the mating process, while the secondary receptacle provides the precise optical alignment of the connector interface with the waveguide interface. Figure 4-43b shows an OPCB connector plug during the mating process about to engage with the primary connector receptacle mounted on the opposite side of the board.

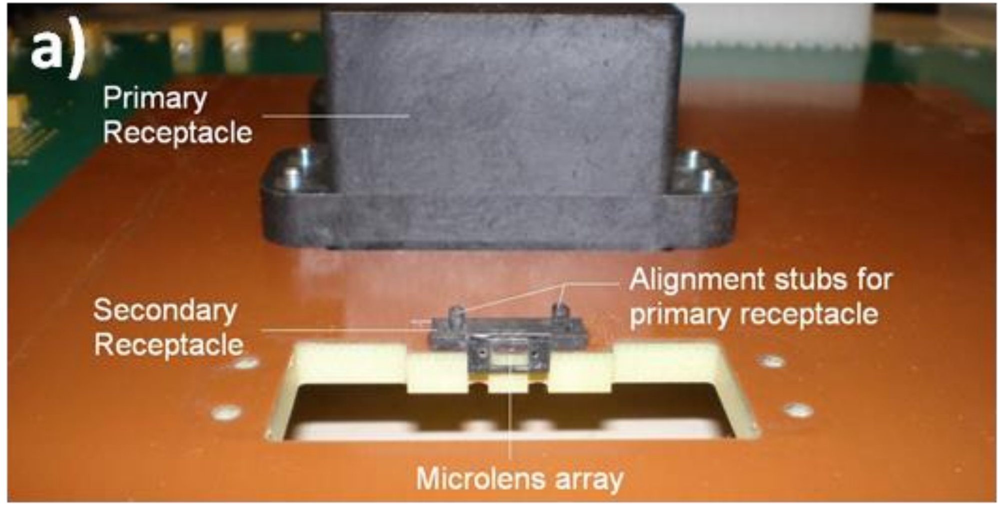


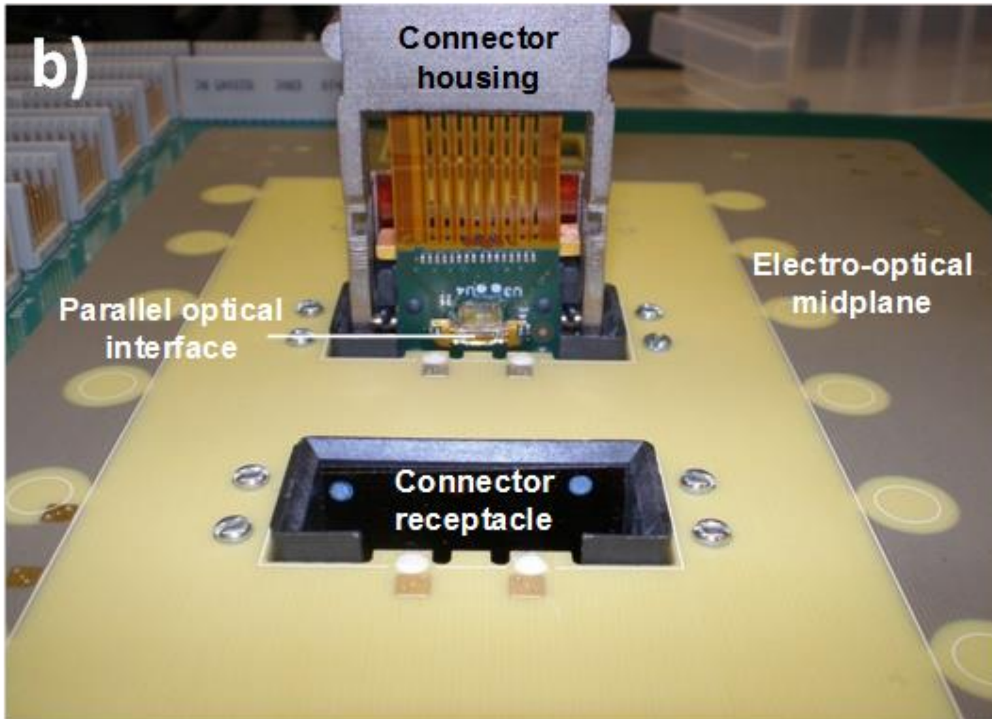


**Figure 4-43: a) Assembly of primary receptacle and secondary receptacle in optical engagement slots, b) Pluggable optical connector prior to engagement with OPCB. Parallel MT compliant optical interface is visible as connector starts to engage with connector receptacle**

## *4.5 StorConnTest3 – FirstLight 10 Gb/s test daughtercard*

The author designed the StorConTest3 peripheral test daughtercards (Figure 4-44) to relay external 10.3 Gb/s 10 GbE LAN test data to each other optically across the StorConnOpt3 electro-optical backplane through the pluggable StorConn3 optical connectors. Each test card included a reconfigurable crosspoint switch (Figure 4-44e) to map test data from four commercial 10 Gigabit Small Form Factor Pluggable (XFP) ports on the front end (Figure 4-44f) to the StorConn3 transceiver housed in the pluggable connector module (Figure 4-44b) mated to the StorConnOpt3 electro-optical backplane (Figure 4-44a). The switch also supported multicasting, whereby test data on any of its inputs could be copied to multiple outputs. This way one external test stream could be broadcast to all four VCSEL transmitters in the connector simultaneously allowing it be characterised while fully stressed. An FPGA (Figure 4-44d) was present on the board to allow user communication with the XFPs, crosspoint switch and StorConn3 transceiver. A PCI bridge chip (Figure 4-44c) allowed a user communications interface to be established between a single board computer and all the line cards via the electrical Compact PCI bus and connectors on the StorConnOpt3 backplane.

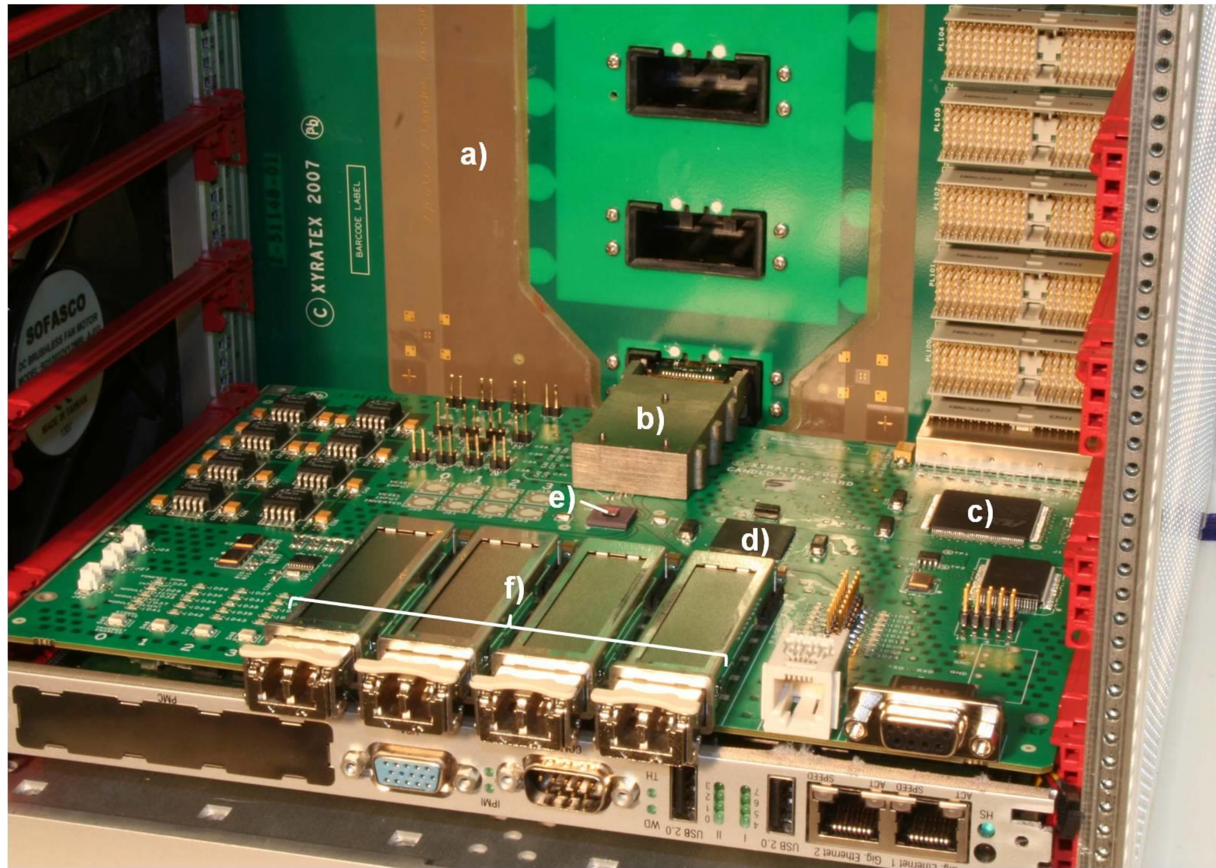


**Figure 4-44: Demonstration assembly with one test card inserted. a) Electro-optical backplane, b) FirstLight connector plugged into backplane receptacle, c) PCI bridge chip next to CompactPCI connector connected to electro-optical backplane providing electrical power and low speed electronic signal interface, d) FPGA to regulate test card, e) crosspoint switch, f) commercial XFPs**

### 4.5.1 Functional overview

Figure 4-45 provides a functional diagram of the StorConnTest3 daughtercard, which includes the following key features:

- The card houses four host ports to accommodate commercial XFP transceivers required for the transfer of 10.3 Gb/s test data between external protocol analyser and the adjoining StorConn3 transceiver circuit.
- Mezzanine connector to deliver all electrical and electronic signals including power, high and low speed and static control signals to the StorConn3 circuit.
- FPGA to regulate the functions of the StorConnTest3 XFPs and the StorConn3 transceiver.

- Compact PCI bridge chip and Compact PCI connector, which allows a computer motherboard connected to the Compact PCI bus of the same backplane to communicate with the regulating FPGA on each StorConnTest3 daughtercard connected to the backplane Compact PCI bus

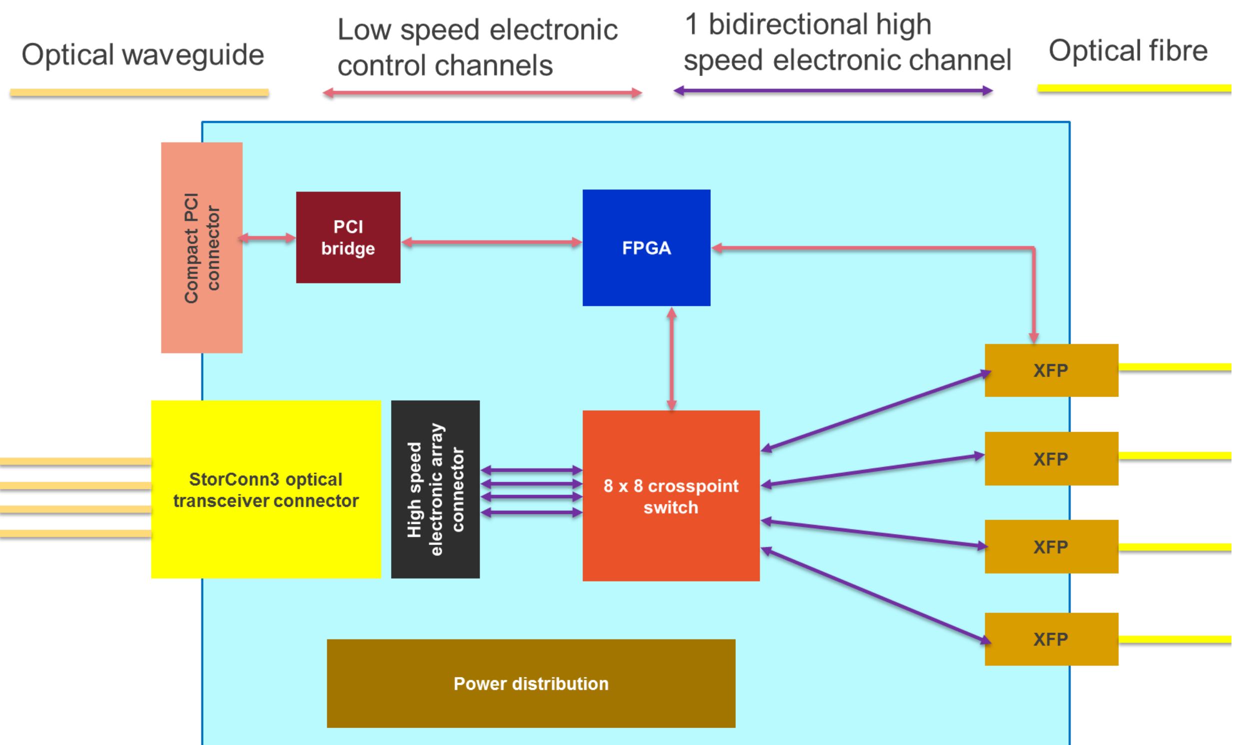


**Figure 4-45: StorConnTest3 functional diagram**

### 4.5.2 StorConnTest3 component listing

Table 4-3 lists the main components of the StorConnTest2 card, while Figure 4-46 shows the component layout.

**Table 4-3: StorConnTest3 key component listing**

| Component | Supplier | Part No. | Description |
|---|---|---|---|
| 8 x 8 crosspoint switch | Vitesse | VSC3008 | 10 Gb/s capable crosspoint switch with eight input ports and eight output ports allowing any of one or more output ports to be connected to a given input port |
| XFP cages / connectors | Tyco | 788862-1 / 1489951-1 | Cages and connector bezels to accommodate different XFPs on host board including 850 nm multimode XFPs and 1310 nm singemode XFPs |
| FPGA | Altera | EP2C8 | Field programmable gate array allowing bespoke user programmed functionality to be loaded onto the chip |
| FPGA configuration EEPROM | Altera | EPC4 | EEPROM to programme the FPGA after each power cycle (FPGA is volatile) |
| Comact PCI bridge | PLX Technology | PLX9054 | Bus master I/O accelerator chip to provide a signal bridge between the Compact PCI backplane and the regulating FPGA |
| PECL oscillator 161.132 MHz | Epson Electronics | Q3803CA | Supplies clock signal for XFP CDR signal conditioning at 10.3 Gb/s |
| PECL differential clock driver | On Semiconductor | MC100EP14 | 1 : 5 PECL clock distribution buffer to relay single 161.132 MHz clock source to all four XFPs for CDR conditioning at 10.3 Gb/s |
| Four way right-angled connector | Molex | 53109-0410 | Auxiliary four way power connector to enable board to be powered outside the test enclosure |

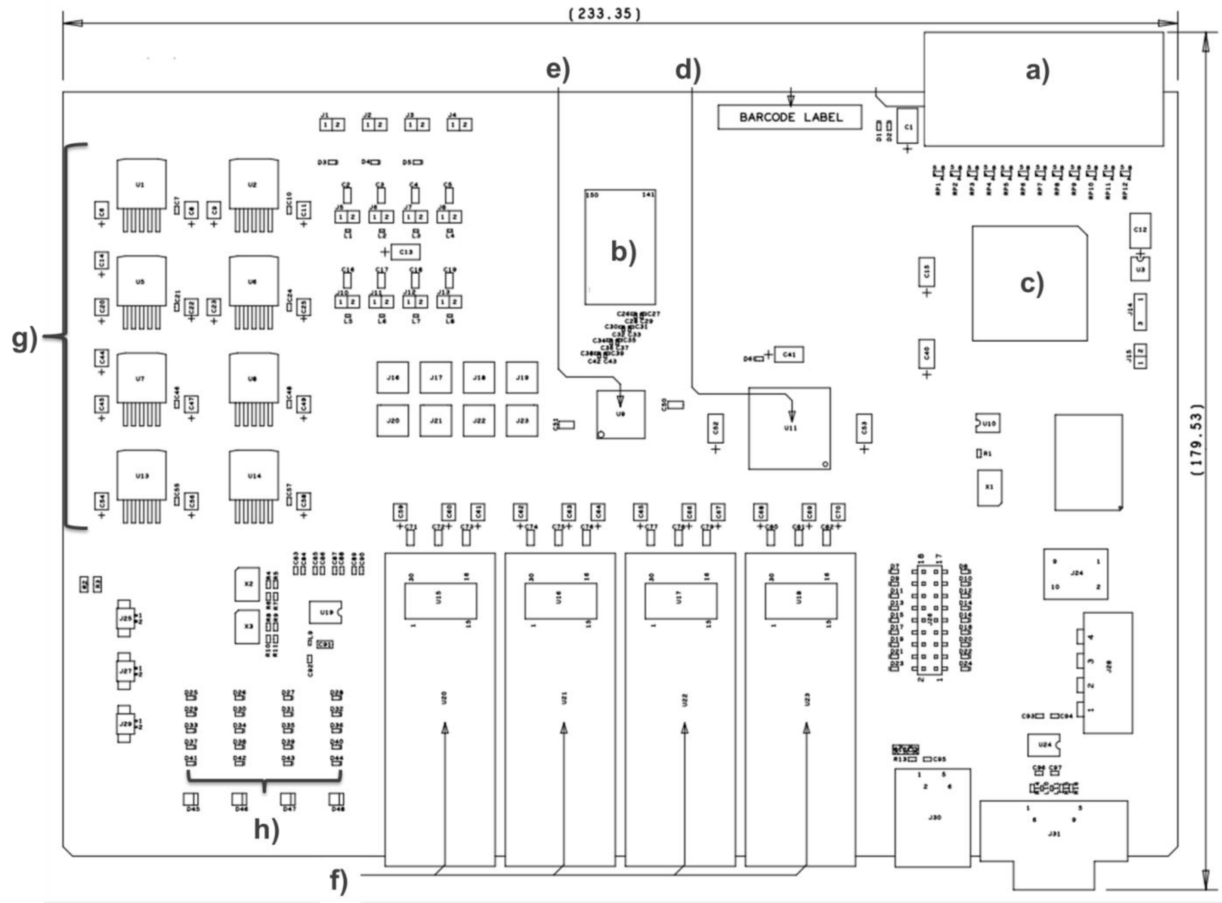


**Figure 4-46: StorConnTest3 board dimensions and component layout: a) CompactPCI connector plug, b) high speed electronic array connector, c) PCI bridge chip, d) FPGA to regulate test card, e) crosspoint switch, f) four commercial XFP cages, g) linear voltage regulators to provide different voltage supplies to card components, h) four banks of indicator LEDs for each XFP**

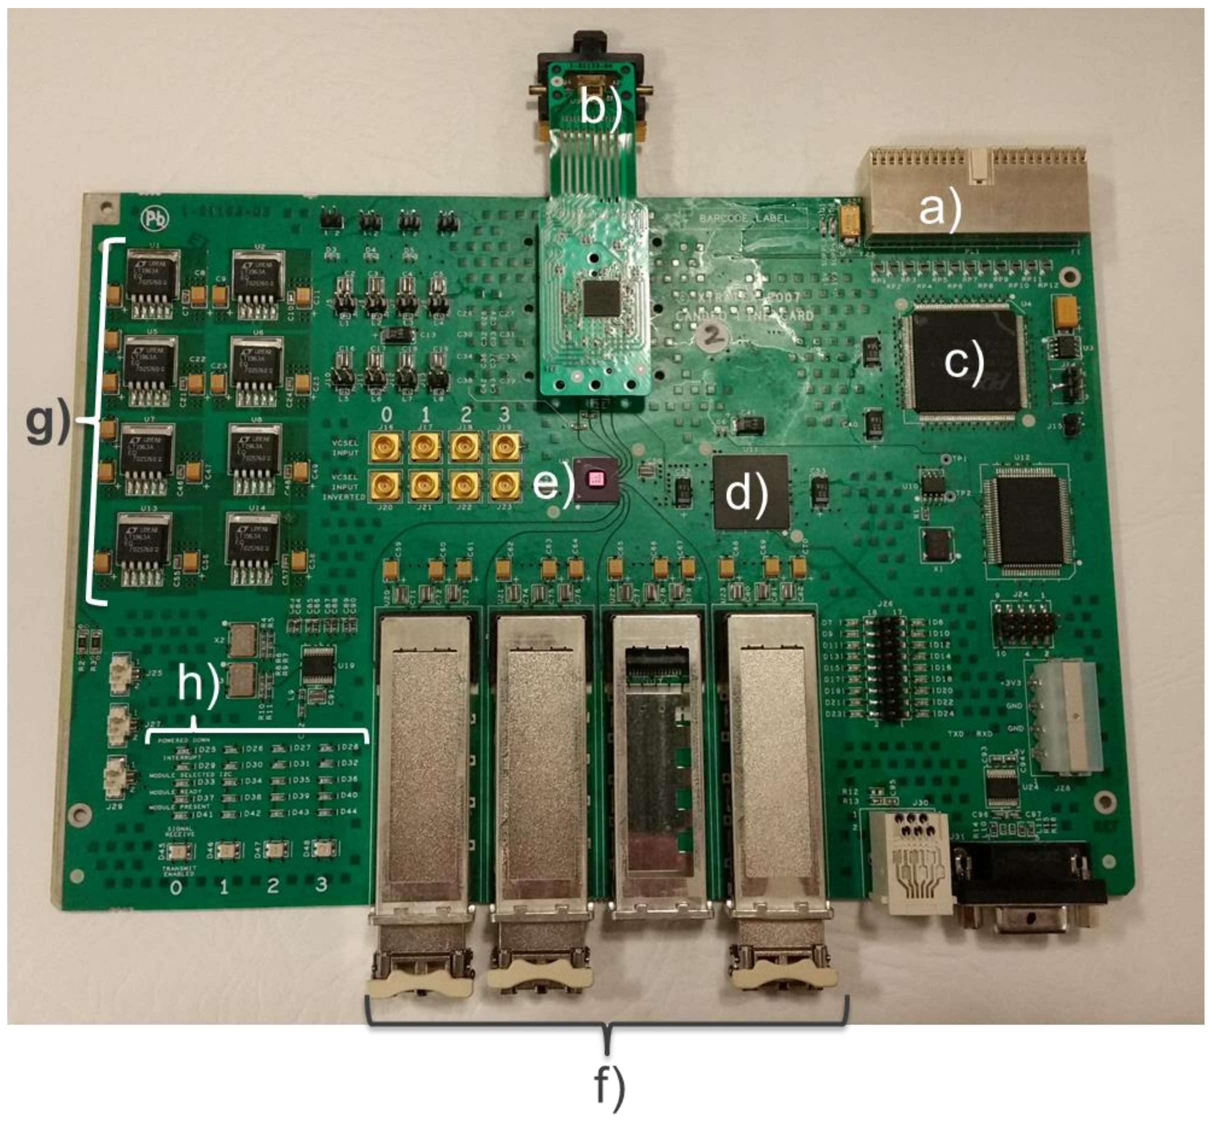


**Figure 4-47: StorConntest3 card with StorConn3 transceiver circuit (without connector module): a) CompactPCI connector plug, b) StorConn3 transceiver circuit, c) PCI bridge chip, d) FPGA to regulate test card, e) crosspoint switch, f) four commercial XFP cages, g) linear voltage regulators to provide different voltage supplies to card components, h) four banks of indicator LEDs for each XFP**

#### *4.5.2.1. Raw card stack-up*

The circuit is housed on an eight layer PCB, of which two layers, accommodating the high speed signal traces are composed of a Rogers dielectric material – RO4350B and the remaining layers accommodating low speed signals and power planes of a standard FR4 composite.

The high speed controlled impedance tracks are routed only on layer one and layer eight and designed to maintain a differential impedance of 100 Ω ± 8% along the trace.

Figure 4-48 shows the complete lay-up of this circuit board.

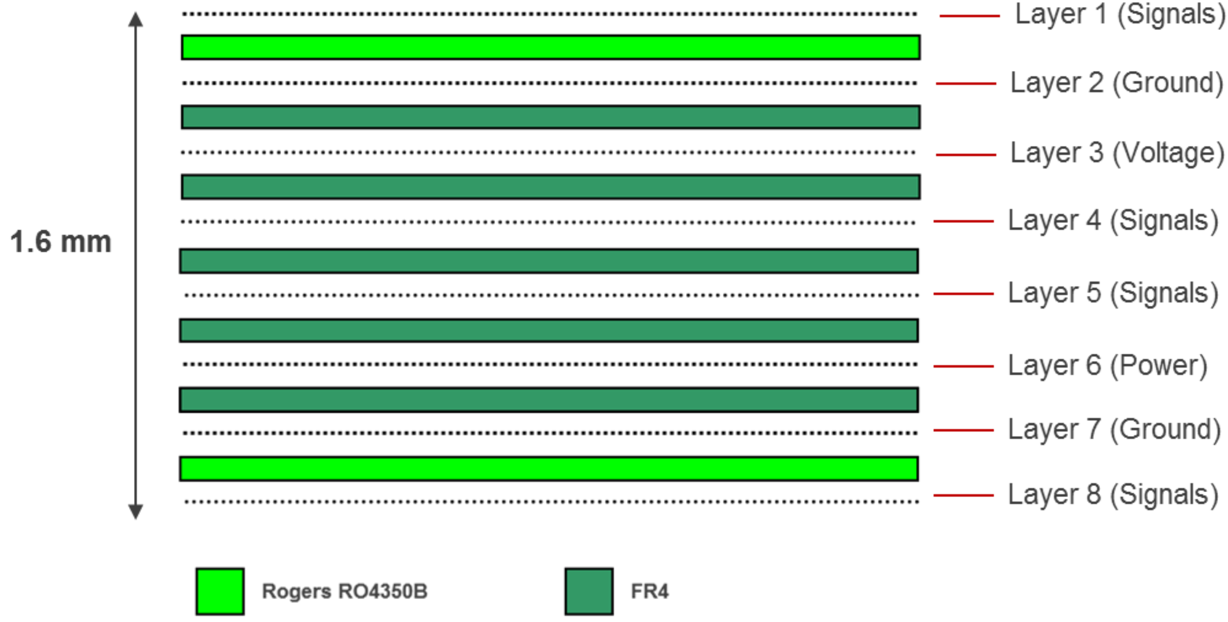


**Figure 4-48: StorConnTest3 raw card stack-up**

#### *4.5.2.2. High speed PCB layout guidelines*

The high speed differential copper signal traces are laid out across the top and bottom layers. These layers are composed of a Rogers dielectric material offering greater impedance control and a lower loss tangent at 10 Gb/s in comparison to FR4. The use of vias is unavoidable if crossovers on the high speed transmit and receive lines are to be prevented. In order to minimise unwanted reflections, any vias on high speed traces pass from top to bottom layers and not to intermediary layers, in order to avoid stubs. A ground via is located in the vicinity of each high speed via. The purpose of these “stitched” vias is to accommodate the return path of the signal at *all* points along the signal path. All aspects of the circuit layout have been adapted to minimise the lengths of the high speed traces.

### 4.5.3 FirstLight demonstration platform

To evaluate the viability of these technologies, the author designed a demonstration platform (Figure 4-49), which comprised a 10 U (445 mm) high Compact PCI chassis with a single board computer, the StorConnOpt3 electro-optical backplane and four peripheral StorConnTest3 test daughtercards, each housing a pluggable StorConn3 optical connector.

**Figure 4-49: FirstLight demonstration platform fully populated with all four test line cards and powered up**

Although the backplane had been designed to accommodate eight Compact PCI slots, only five were populated. As shown in Figure 4-50, the bottom slot A was reserved for the single board computer motherboard, while slots

B, D, F and H were reserved for StorConnTest3 cards. Due to the height of the StorConn3 connector module, there was insufficient space to allow StorConnTest3 cards to be plugged into adjacent Compact PCI slots.

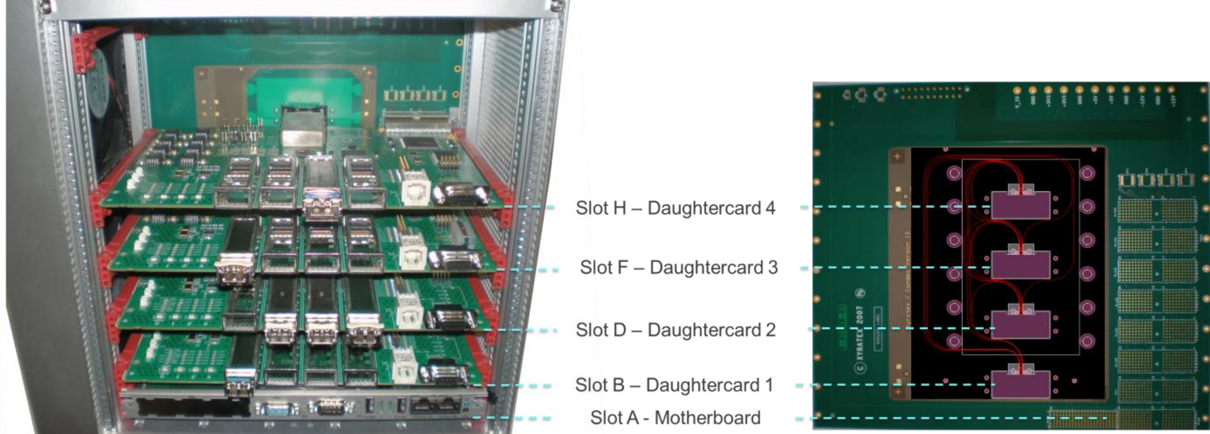


**Figure 4-50: Compact PCI slot population by singleboard computer motherboard and four StorConnTest3 daughtercards. Slots C, E and G are unpopulated as the daughtercards require a separation of at least 2 slots**

### 4.5.4 FirstLight platform FPGA firmware design

The purpose of the singleboard computer motherboard located in slot A of the FirstLight platform was to run an operating system, through which the user can communicate with and configure any of the four StorConnTest3 daughtercards connected to the Compact PCI bus of the same StorConnOpt3 backplane.

The author designed and developed an extensive firmware programme for each StorConnTest3 FPGA to regulate and interpret the PCI commands received from the singleboard computer via the shared Compact PCI bus. The author wrote the programme using VHDL (VHSIC Hardware Description Language), which was loaded onto the Altera EP2C8 FPGA by the Altera EPC4 FPGA configuration EEPROM with each power cycle.

The firmware code incorporated a comprehensive command map that would allow the user to read and write to registers in the FPGA, which in turn would read-back or configure an extensive range of parameters including:

- General purpose static control and status LED values
- StorConn3 transceiver parameters
- Crosspoint switch parameters
- Host side XFP parameters

A comprehensive description of the complete PCI command map and required PCI command sequences is provided in Appendix - FirstLight platform firmware coding.

### 4.5.5 FirstLight Graphical User Interface (GUI)

The author developed the requirements specification for a graphical user interface (GUI) to run on the operating system of the single board computer, allowing selective user access to any StorConnTest3 daughtercard connected to the same StorConnOpt3 backplane and permitting the user to configure the parameters of the selected StorConnTest3 card. The FirstLight GUI was developed by Xyratex software engineer Mike Horgan (Figure 4-51).

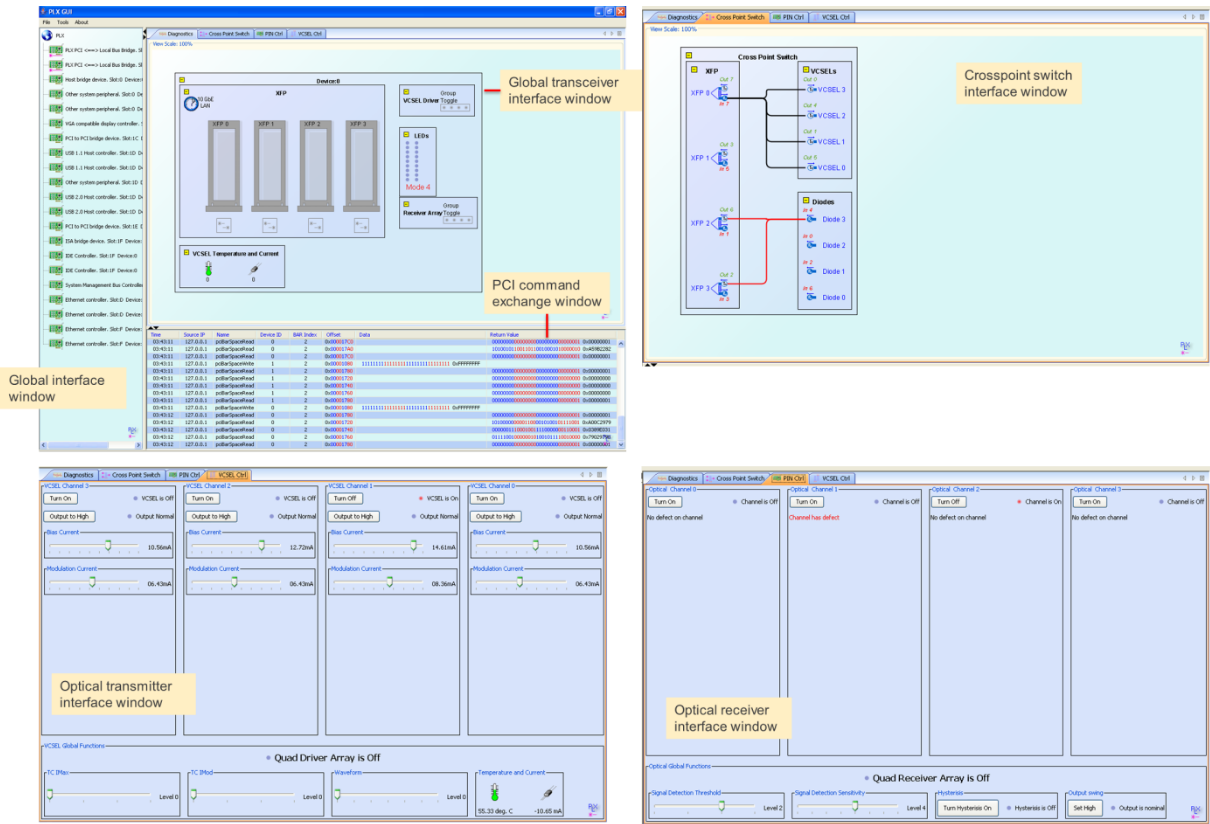


**Figure 4-51: FirstLight GUI overview**

#### *4.5.5.1. Diagnostic interfaces*

***Global interface window***

- List of all StorConnTest3 daughtercards on the C-PCI bus i.e. docked in the C-PCI chassis based on card specific ID as stored in PLX EEPROM
- Choice list to select between any StorConnTest3 daughtercards on the C-PCI bus

***PCI command exchange window***

- PCI Reset button
- Displays the 32-bit word transactions between host and line card
- Displays information on the BAR (space 0) offset for the selected PLX device (as stored on its EEPROM)

#### *4.5.5.2. Control interfaces*

***Global transceiver interface window***

- Buttons for both Power Down and Transmit Disable on each of the four XFPs of the selected line card
- XFP I²C interface control:
    - Read back on all critical XFP information e.g. Vendor ID
    - Allows read back of any user selected address on the I²C XFP EEPROM
    - Allows data write to any user selected address on the I²C XFP EEPROM

#### *4.5.5.3. StorConn3 optical transmitter interface window*

- Button for global VCSEL Driver Enable
- Buttons for individual VCSEL channel enables
- Scroll bars for modulation current and bias current on each VCSEL
- Scroll bar for temperature compensation
- Temperature read-back display (polled every five seconds when no other activity)

#### *4.5.5.4. StorConn3 optical receiver interface window*

- Button for global optical receiver enable and squelch enable
- Buttons for individual photodiode channel enables
- Scrollbar for waveform control

#### *4.5.5.5. Crosspoint switch (CPS) interface window*

- Global and channel selective signal input equalisation settings:
    - no equalisation
    - minimum equalisation
    - medium equalisation
    - maximum equalisation
- Global or individual output settings:
    - Force all outputs to 1, 0 or normal operation
    - Output power nominal or high
    - Pre-emphasis enabled or disabled
    - Pre-emphasis adjustment for varying line lengths between range 0 to 15 (450 ps to 700 ps respectively)
- Switch Control:
    - Choice lists allowing mapping of switch inputs to switch outputs
- Switch Map:
    - Displays current switch configuration i.e. what inputs are mapped to what outputs

### 4.5.6 FirstLight test and measurement results

#### *4.5.6.1. Direct VCSEL transmitter characterisation*

Both the bias current and the modulation current range of the VCSELs could be programmed through the VCSEL driver chip. 10.3 Gb/s test data was conveyed through the optical interface, the output captured through the fibre interface and passed to the Tektronix CSA. The variation in jitter with modulation current range was measured at two bias currents 8.94 mA and 10.56mA and is shown in Figure 4-52. In both cases the jitter reaches a steady minimal state once the modulation current range is set at ≥15 mA.

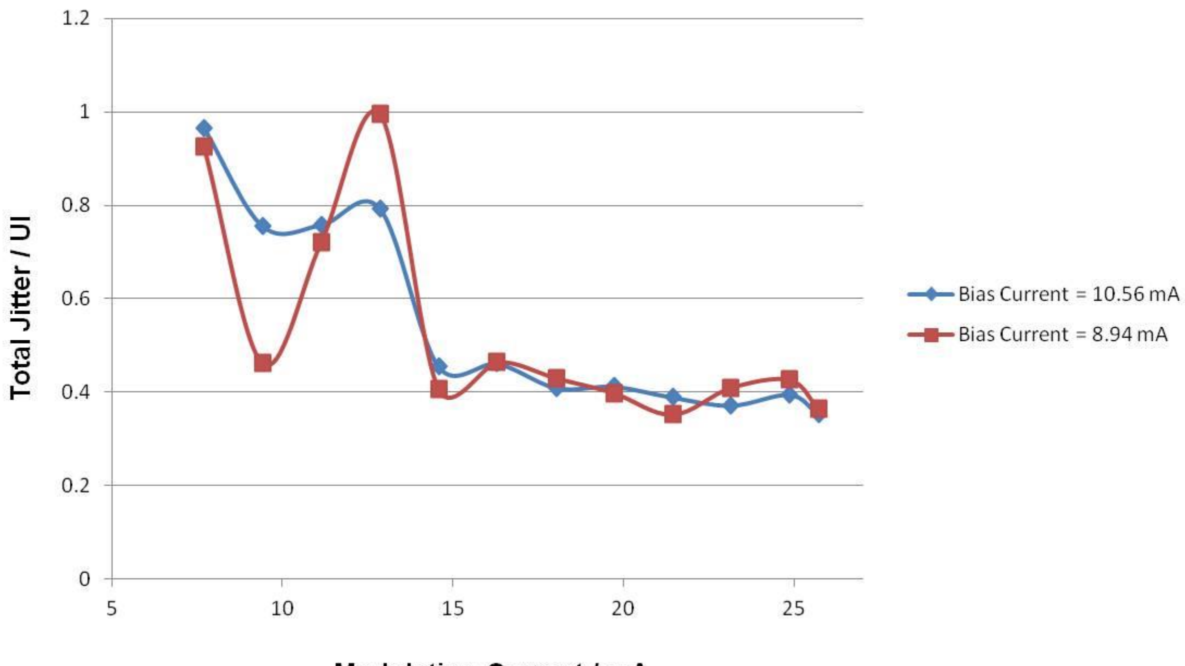


**Figure 4-52: Characterisation of optical jitter from VCSEL channel one operating 10.3 Gb/s with modulation current**

#### *4.5.6.2. High speed data transmission loop-back test across demonstration platform*

An external Xyratex proprietary 10 Gb Ethernet LAN traffic source was arranged to convey a 10.3 Gb/s test data stream along a fibre-optic cable to one of the commercial XFP devices on the front end of a peripheral test card in the demonstration platform. The XFP device converted the optical data stream to a serial electronic data stream on the test card, which was then mapped by the crosspoint switch to one of the VCSEL transmitters in the connector attached to that card and reconverted into an optical data stream. As the connector was optically engaged to the backplane, the optical data stream was launched into a waveguide and conveyed to the receive element of another connector on a different test card in the chassis. The data was then converted to a serial

electronic data stream, mapped to an XFP port on that test card and reconverted to an optical data stream on the output of the XFP device.

Finally a fibre-optic cable was connected between the XFP output port to a Tektronix CSA8000B communications signal analyzer where the test data was characterised.

Figure 4-53 shows the demonstration assembly.

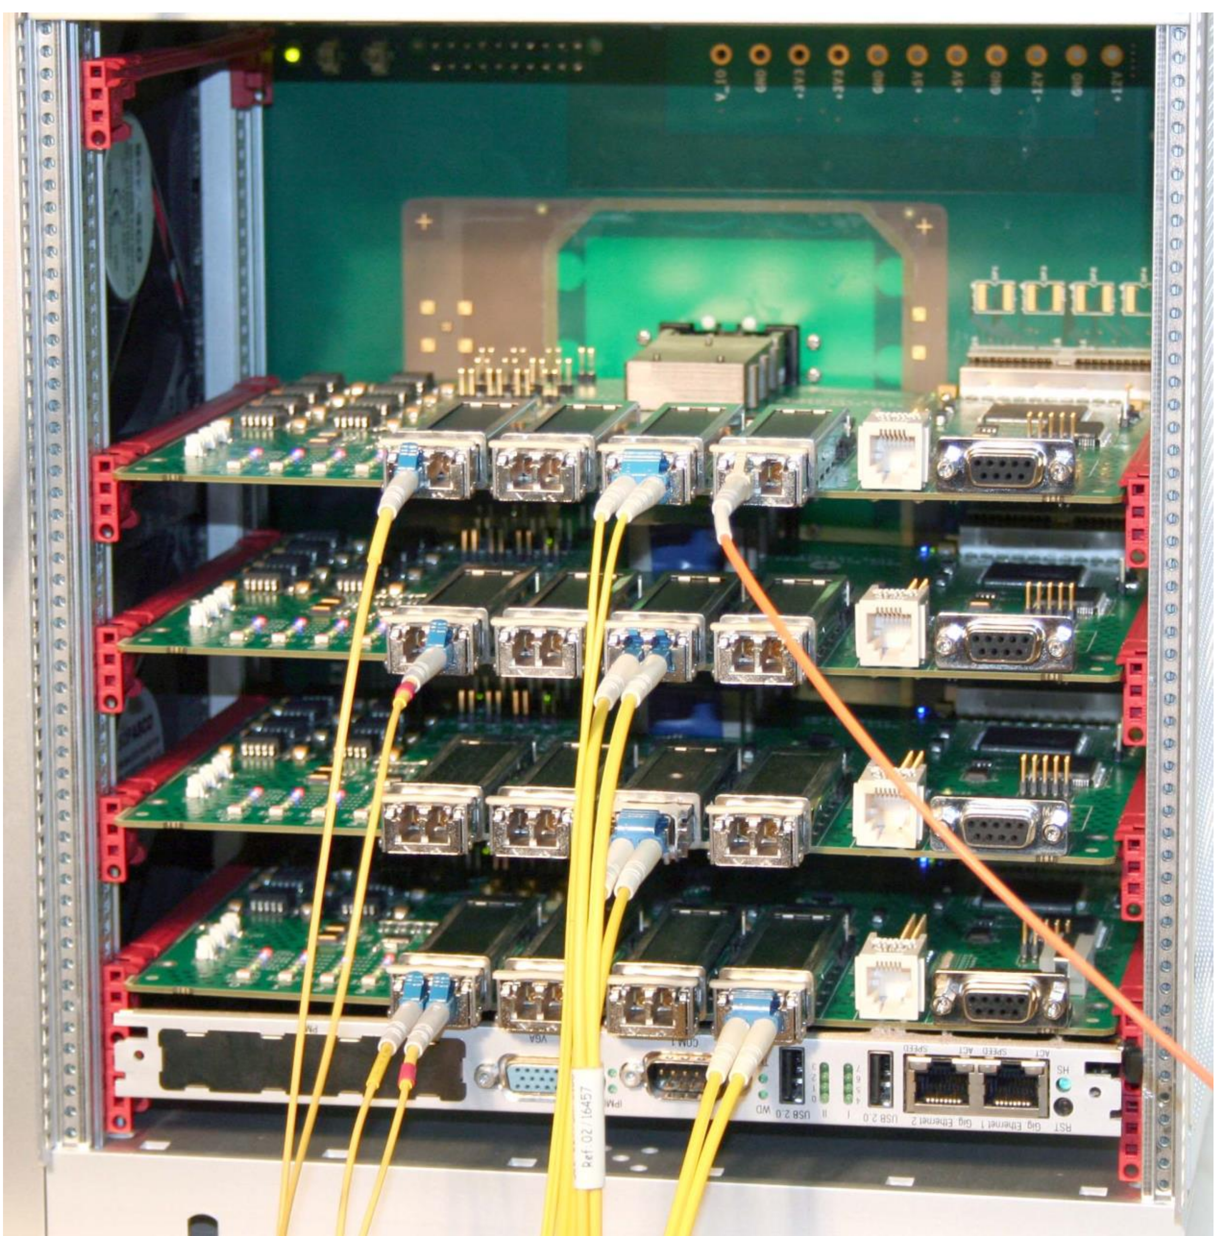

**Figure 4-53: FirstLight demonstration platform with optical test cables attached to XFPs**

Figure 4-54 to Figure 4-61 show the eye diagram captured across the highlighted waveguide link.

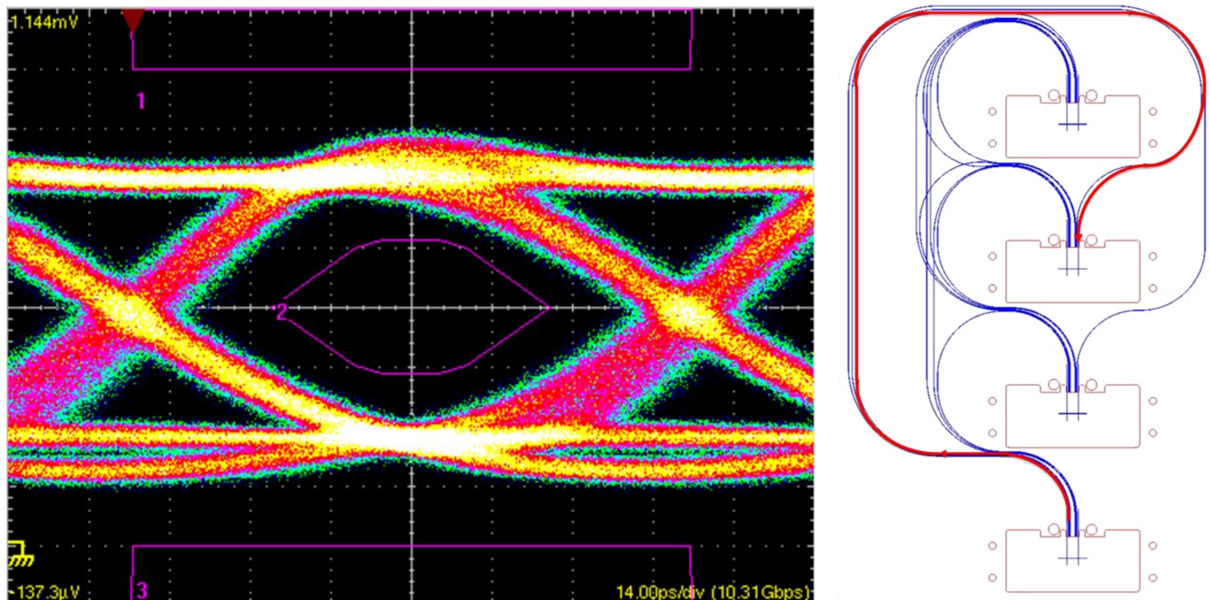


**Figure 4-54: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel one**

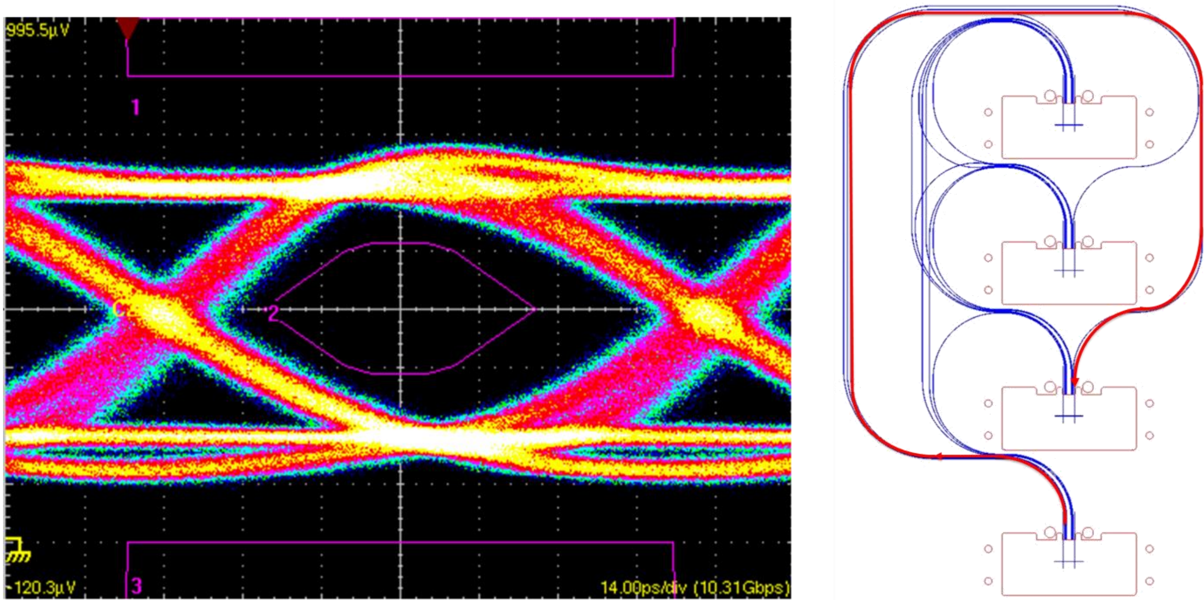


**Figure 4-55: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel two**

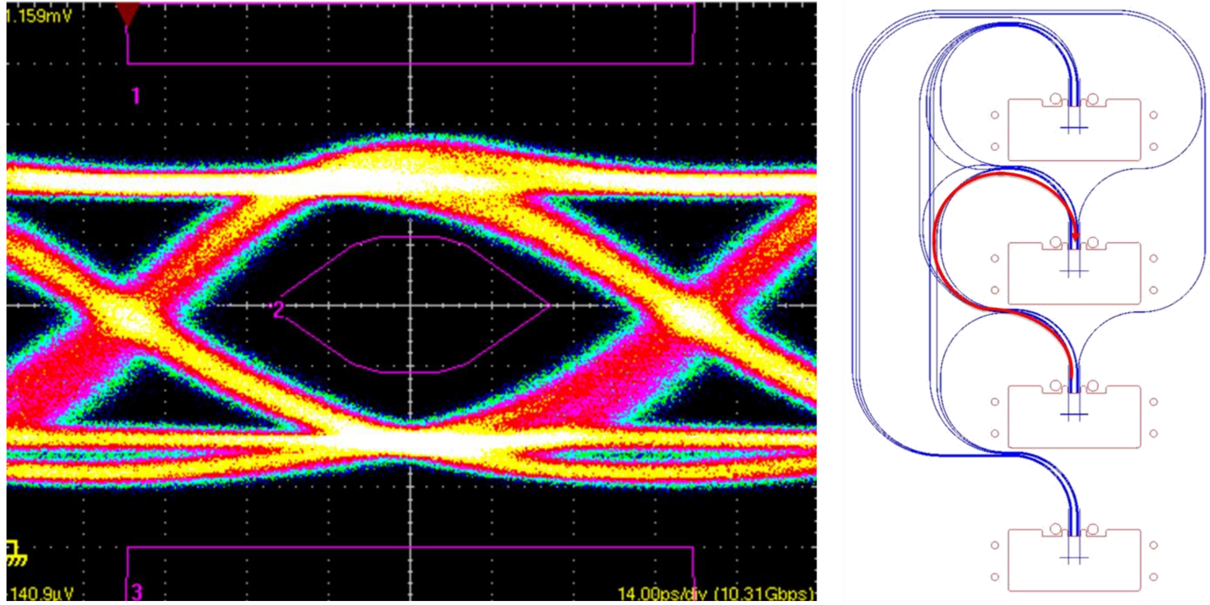


**Figure 4-56: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel three**

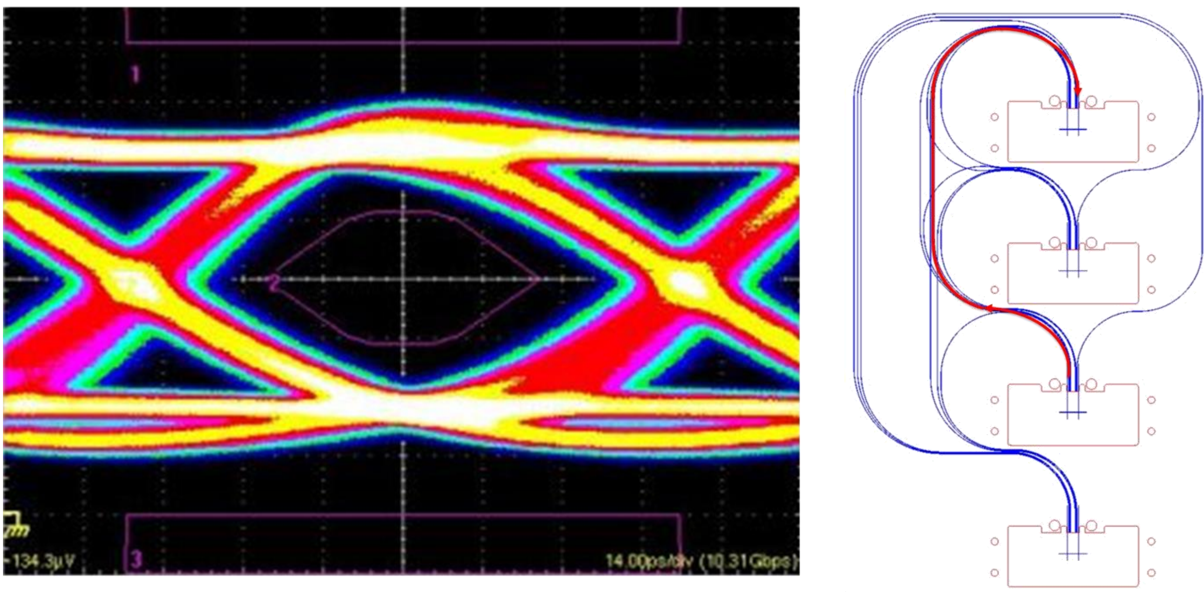


**Figure 4-57: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel four**

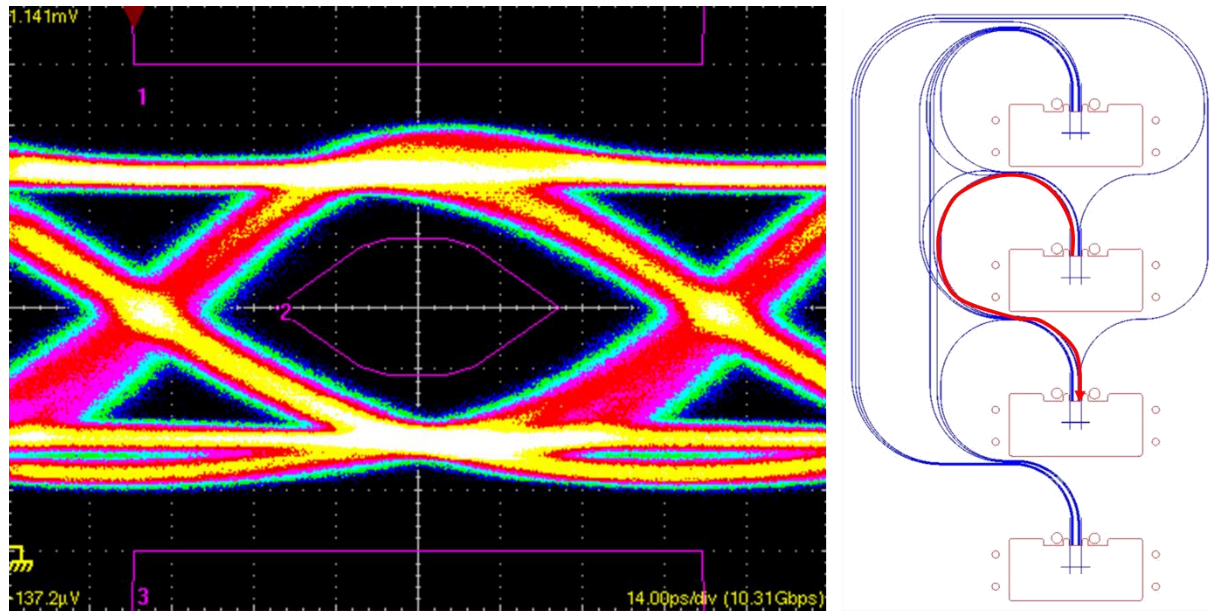


**Figure 4-58: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel five**

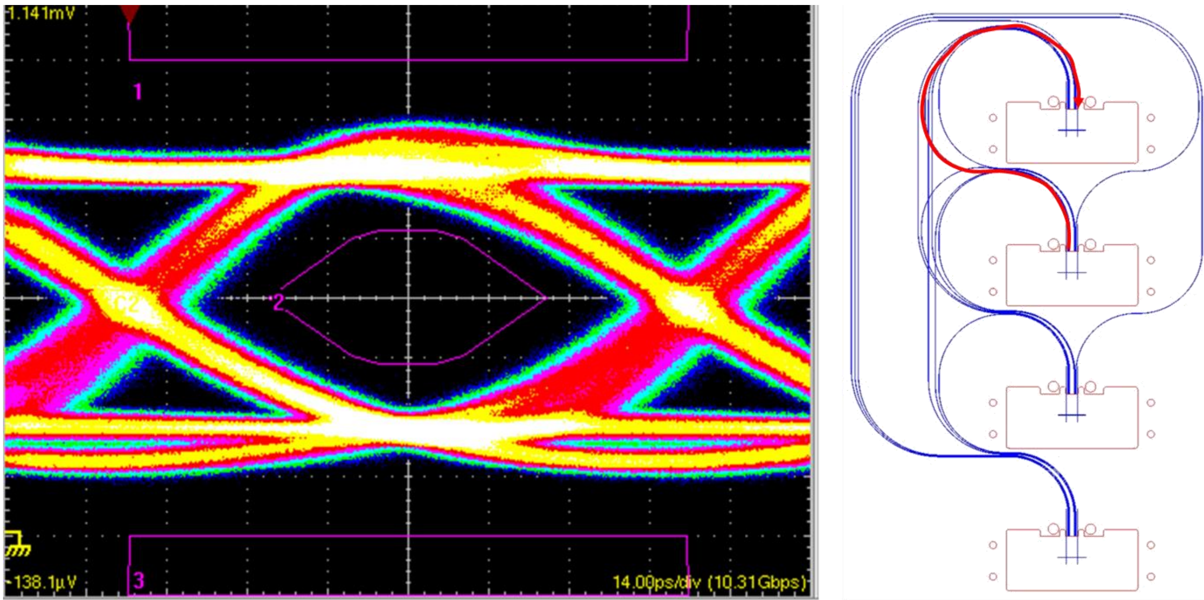


**Figure 4-59: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel six**

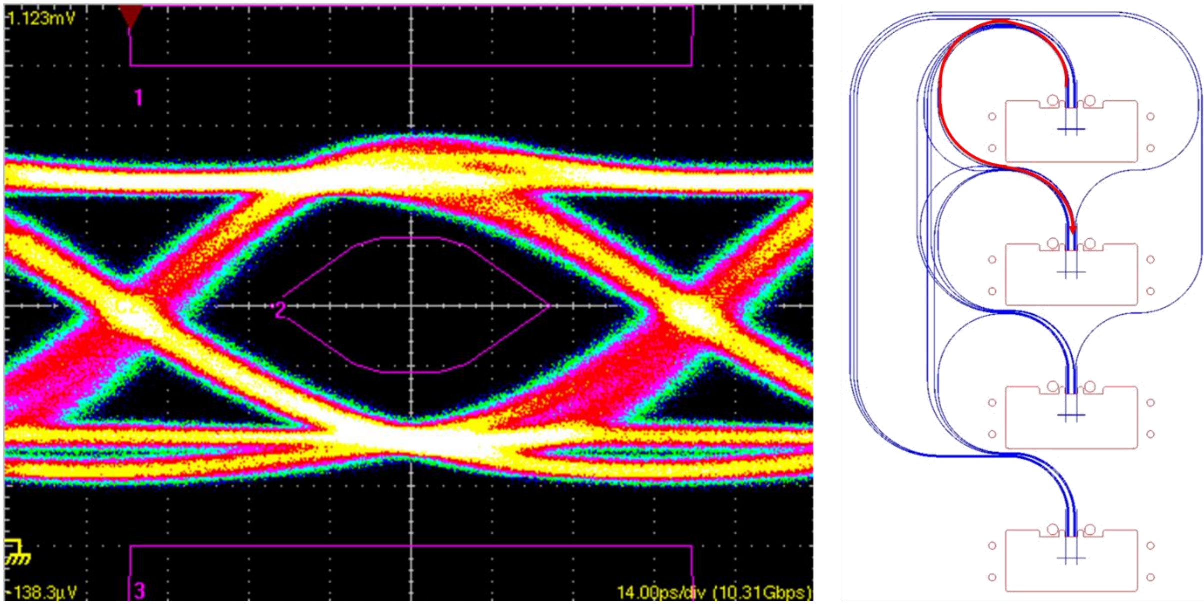


**Figure 4-60: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel seven**

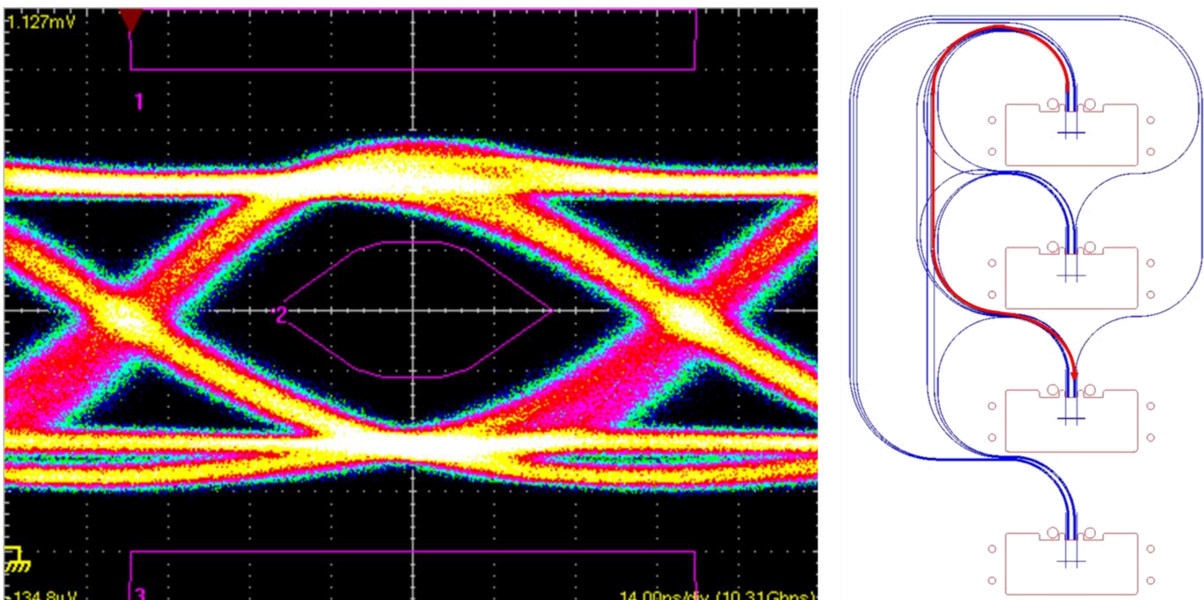


**Figure 4-61: Eye diagram of 10.3 Gb/s optical signals received from demonstration platform on selected waveguide under test - Channel eight**

In total eight waveguides were tested as described and 10.3 Gb/s test data was successfully conveyed between all test cards and their prototype connectors with an acceptable level of signal recovery. Figure 4-54 - Figure 4-61 show the eye diagrams corresponding to the eight waveguides under test with a schematic overlay in each case

identifying the waveguide and direction. The average total peak to peak jitter on the eight communication links including front end clock and data recovery through the exit XFP was measured to be 28.217 ps or 0.29 Unit Interval (UI), which is within the jitter thresholds for both transmitter input (0.61 UI) and receiver output (0.363 UI) as specified by the XFP MSA [150].

## *4.6 Summary*

In this chapter, the design and development of a second generation active pluggable optical connector was described, which allows the two stage engagement process, validated in the first generation connector, to be carried out without intermediary manual interference by the user, thus overcoming the critical flaw in the viability of the first generation connector.

An enhanced backplane waveguide receptacle was detailed, which exhibits no offset between the alignment slot datum and the channel interface thus improving accuracy compared to the first generation receptacle. Crucially, the receptacle housing provides a protective barrier over the waveguide interface reducing the chance of dust contamination. The housing on the receptacle includes a geometric micro-lens array suspended over the waveguide interface, while the housing on the connector includes a compliant micro-lens array suspended over the VCSEL and PD arrays. Together these form part of a dual expanded beam system, whereby, at the exposed mating interface between two lens arrays, the beam width is expanded to many times the size of the waveguide aperture, thus strongly reducing susceptibility to dust.

The final demonstration system was based on a Compact PCI bus backplane system, in which a computer motherboard and four daughtercards with active optical connectors were oriented parallel to each other and connected into an electro-optical PCB backplane. A complex optical waveguide pattern was integrated into the OPCB backplane to allow full optical connectivity between the four daughtercards.

The greatest disadvantage in the approach of using active optical connectors is that the optical receptacle on the receiving OPCB must accommodate not just the size of the optical interface, but also the body of the active connector mechanism, which includes the size taken up by the integrated transceiver. This means that only a small number of active optical connectors could be accommodated for each daughtercard due to size restrictions and therefore the number of optical channels supported on the OPCB would be far less than could be physically accommodated based on the size of the channels themselves. The relatively large size of the active optical connector therefore represents a bottleneck to the number of supportable channels on the OPCB.

If a passive optical connector system were instead used, whereby the transceiver would not be local to the

interface but positioned somewhere else on the board, preferably close to the signal source, such as in the case of a midboard transceiver, then the space consumed by the now passive connector on the OPCB would be substantially reduced.

The use of polymer waveguides aligns very well to low cost 850 nm VCSEL based transceivers, as most optical polymers exhibit the least absorption around this wavelength. However transceiver technologies, formerly confined to long distance telecommunications applications and operable at longer wavelengths such as 1310 nm and 1550 nm, are now also starting to migrate down the data centre levels in response to the need to accommodate increasing bandwidths over shorter distances. One can foresee a point at which it will become necessary for OPCBs to accommodate such interconnect, however optical polymers would be an unsuitable optical waveguide medium as they exhibit much higher losses at these longer wavelengths.

In the next chapter, a new generation of OPCB will be described based on planar glass waveguides, which exhibit lower losses at longer wavelengths, and on which board-to-board pluggability is achieved through passive rather than active optical connectors allowing much higher optical interface densities to be accommodated.

# 5 PLUGGABLE PASSIVE OPTICAL CIRCUIT BOARD CONNECTOR FOR PLANAR GLASS WAVEGUIDE BASED OPTICAL CIRCUIT BOARDS

## *5.1 Introduction*

While polymer waveguides are well suited to applications requiring 850 nm optical signals, such as low cost optical links based on commodity, directly modulated VCSELs, planar glass waveguides may be preferable in applications requiring longer wavelengths such as those around 1310 nm (O-band) or those around 1550 nm (C-band) due to their superior transmissivity at these wavelengths.

Deployment of electro-optical printed circuit boards (OPCBs) based on embedded planar glass waveguide technology in the system, would enable seamless optical connectivity from external fibre-optic networks directly onto the system motherboard or backplane. Furthermore the emergence of affordable longer wavelength

transceiver solutions based on photonic integrated circuit (PIC) technologies such as silicon photonics transceivers and switches [157]–[160] makes glass waveguides an attractive OPCB technology.

### 5.1.1 Glass waveguide based optical circuit boards

The Fraunhofer Institute of Reliability and Microintegration (Fraunhofer IZM) in Germany have spearheaded glass waveguide fabrication for embedded modules and OPCBs over the past decade [84], [103], [161]–[163]. During and after the SEPIANet project, the author has collaborated with Fraunhofer IZM extensively to help demonstrate the viability of glass waveguide based OPCBs in modern data centre systems [4], [83], [164], [165].

### 5.1.2 SEPIANet project

In March 2010, the European PIANO+ funding competition was launched, which was a trans-national call for project proposals on "photonics-based internet access networks of the future". It was co-funded by the participating national funding bodies from Austria, Germany, Israel, Poland and the United Kingdom and by the European Commission [166].

In 2010, the author wrote and submitted a proposal to the Piano+ competition for a project, which would bring together key European fabricators and allow us to develop a technology eco-system around planar glass waveguide based OPCBs. The project was called System Embedded Photonics in Access Networks (SEPIANet) [31] and comprised the following consortium of European organisations: Fraunhofer Institute of Reliability and Microintegration (Fraunhofer IZM), ILFA, V-I Systems, Conjunct and TerOpta with Xyratex as lead organisation.

As author and organiser of the SEPIANet proposal, the author then continued to serve as technical coordinator of the project and was responsible for defining the technology requirements for the SEPIANet optical backplane and connector platform.

### 5.1.3 SEPIANet technology overview

In this chapter, the key technology advances made on the SEPIANet project are described.

Section 5.3 describes a disruptive new method developed by Fraunhofer IZM of fabricating multimode waveguides with elliptical graded refractive index cross-sectional profiles within the surface region of thin glass foils, which can be scaled to a large area PCB form factor.

Section 5.4 describes how the commercial PCB fabrication processes were adapted by German PCB company ILFA to enable planar glass waveguide panels to be reliably embedded into an OPCB backplane comprising 4 conventional electrical PCB layers with an outer board area of 281 x 233 mm² and board thickness of 3.5 mm, encasing a smaller 199 x 160 mm² glass waveguide panel with a thickness of 500 µm.

Section 5.5 describes a full suite of passive optical connector components designed and developed by the author and Allen Miller at Xyratex to enable pluggable optical connectivity to the glass waveguide based OPCB backplane for fibre-to-board and board-to-board configurations. These included fibre ferrule receptacles to enable direct connection of conventional parallel optical fibre cables to the OPCB embedded waveguides, and pluggable edge connectors to enable daughtercards to be orthogonally plugged into an OPCB backplane. An active, automated assembly process was developed by Fraunhofer IZM and successfully deployed for aligning and attaching the optical fibre ferrule receptacles to the OPCB.

Section 5.6 describes the design, development and characterisation of the SEPIANet test and measurement platform comprising a sub-rack chassis, an OPCB backplane with pluggable optical board-to-board connectors and 5 pluggable test daughter cards. Finally, the results of a comprehensive test and measurement regime are reported, whereby PRBS $2^{31}$-1 optical test data at both 1310 nm and 850 nm was conveyed along the OPCB backplane embedded optical waveguides through the pluggable connector system from various optical test sources and validated for both fibre-to-board and board-to-board optical connectivity at data rates up to 32 Gb/s per channel exhibiting bit error rates of less than $10^{-12}$.

### 5.1.4 Objectives

- Co-develop a waveguide receptacle based on those already developed for polymer waveguides to be assembled onto planar glass waveguide based OPCBs allowing direct fibre-to-waveguide coupling
- Define the requirements and design specifications of passive board-to-board optical connector system, which incorporates the waveguide receptacle, but also makes use of the commercial parallel optical ferrule jumpers.
- Define the requirements and design specifications for the SEPIANet demonstration test platform, which incorporates a glass waveguide based OPCB backplane and peripheral test cards, which can be plugged directly into the OPCB backplane using the optical connectors developed.
- Define an appropriate testing regime for SEPIANet platform and carry out test and measurement

## 5.2 Optical waveguide layout design

In order to accommodate the interconnect topologies of target system enclosures in data centre environments, the optical layout for the OPCBs needed to contain waveguide groups with varied point-to-point geometries between edge or midboard optical connector interface points. A waveguide layout requirements specification was developed by the author for two interchangeable OPCB backplane variants, SEPPLANE1 and SEPPLANE2, which addressed this requirement.

SEPPLANE1 contained one embedded glass waveguide panel (Figure 5-1a) of size 199 mm x 160 mm, while SEPPLANE2 contained two smaller glass waveguide panels in the same plane (Figure 5-1b), each of size 79.25 mm x 160 mm. Both SEPPLANE1 and SEPPLANE2 contained edge connector interfaces, while SEPPLANE2 also contained midboard connector interfaces. SEPPLANE2 also allowed the challenge to the PCB fabricator of laminating multiple glass panels into a single PCB layer to be addressed and demonstrated. This capability would be instrumental in future to provide optical interconnect across larger, higher density backplanes, if individual glass waveguide panel sizes were constrained.

For demonstration purposes only, the designs included exposed windows in the PCB material to show sections of the glass waveguide panel embedded within, though in practise this would be a liability as the exposed sections could be damaged very easily during PCB assembly.

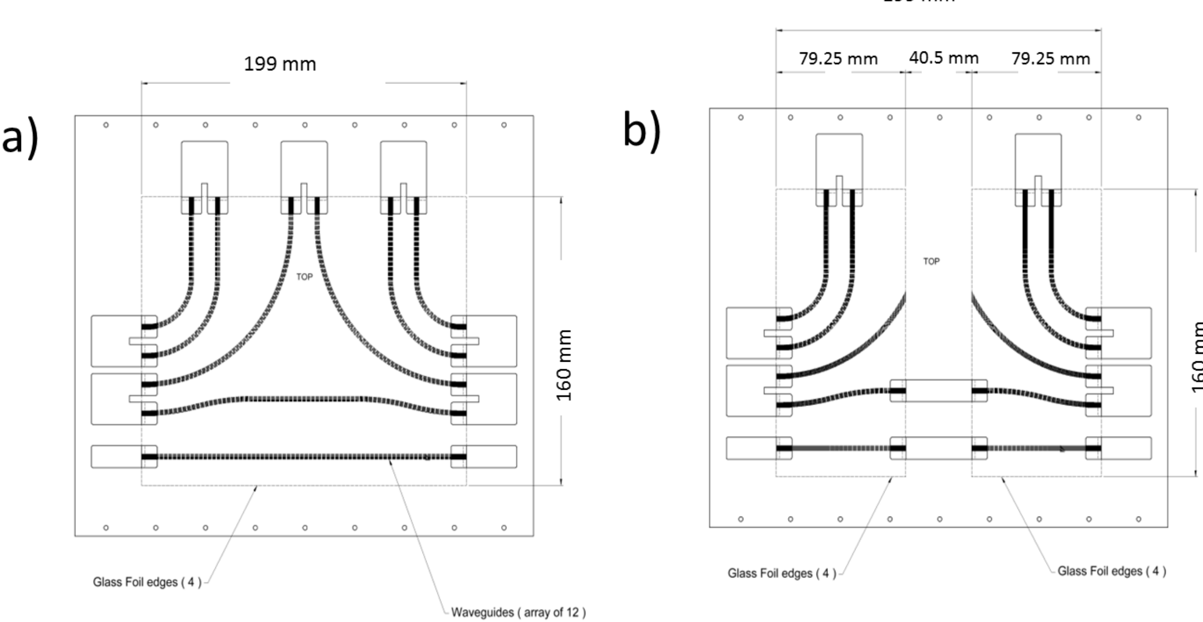


**Figure 5-1: a) SEPPLANE 1 glass panel dimensions, b) SEPPLANE2 glass panel and gap dimensions**

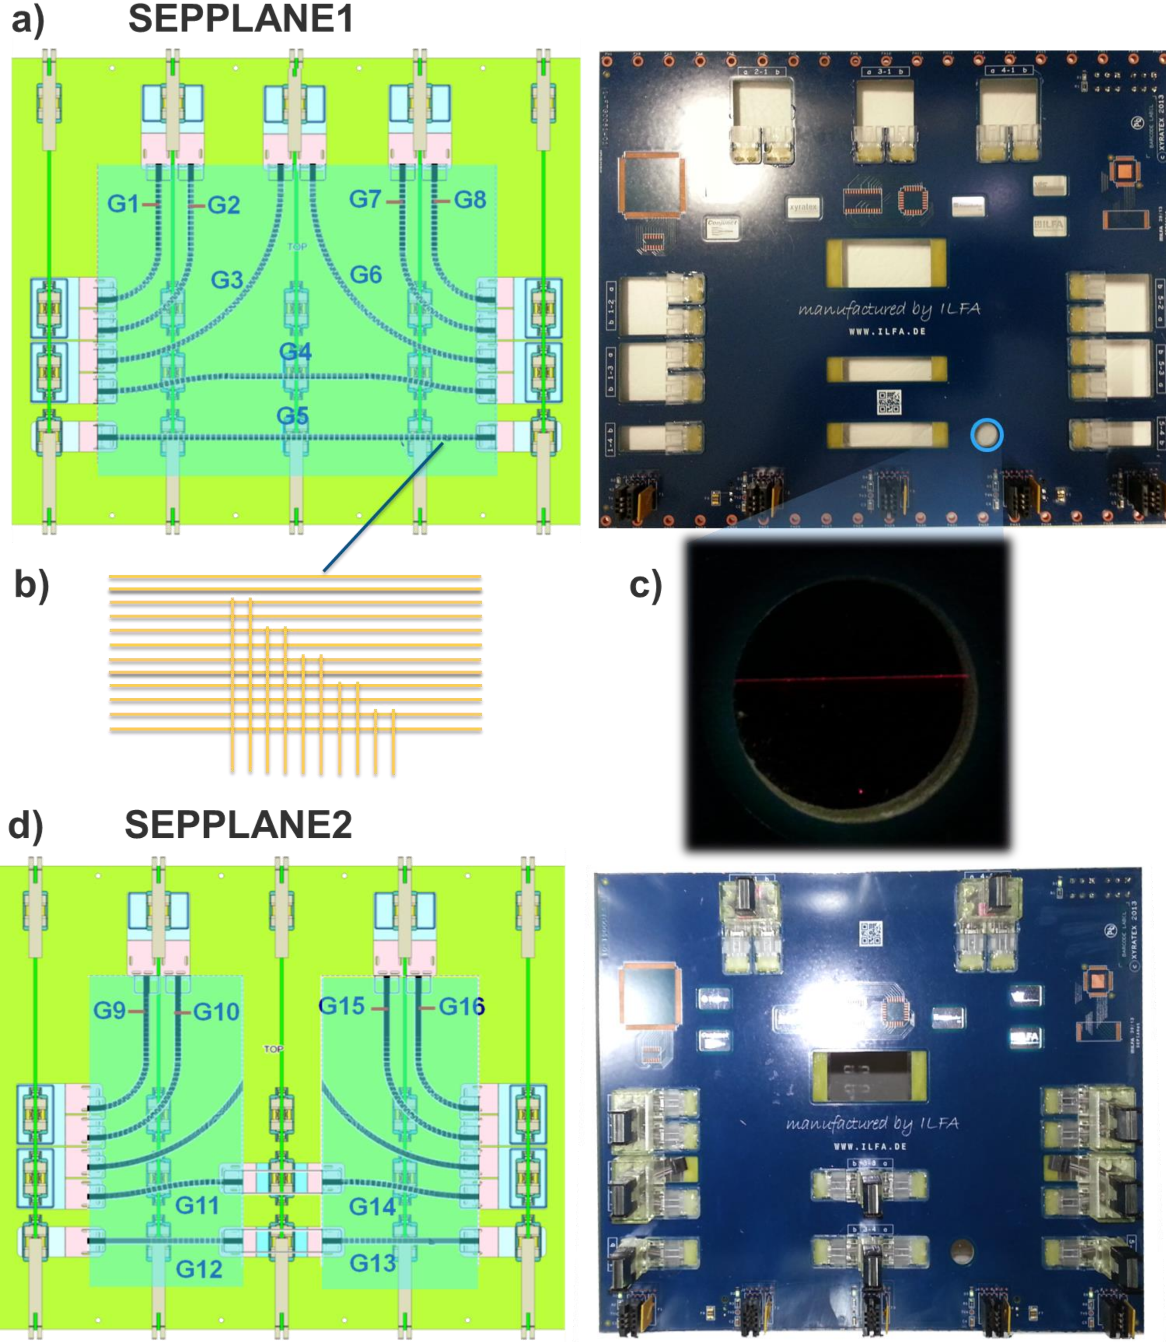


**Figure 5-2: Schematic layouts and photos of fully populated OPCB backplane variants: a) SEPPLANE1 with single glass waveguide panel and connector layout, b) iterated waveguide crossover section, c) top view of embedded glass panel showing single waveguide illuminated, d) SEPPLANE2 with dual glass waveguide panel and connector layout**

As shown in Figure 5-3 and Figure 5-4, the SEPPLANE1 and SEPPLANE2 layouts each consisted of 8 waveguide groups. Each group comprised a row of 12 waveguides with a centre-to-centre channel separation of 250 µm, which is compliant with conventional parallel optical fibre interfaces such as those based on MT

standards. The waveguide layout groups G1-G3, G6-G8 on SEPPLANE1 and G9–G10, G15-G16 on SEPPLANE2 were curved over a 90° arc with varied concentric bend radii. All curved groups included additional straight sections. Groups G4, G11 and G14 included an S-bend, while groups G5, G12 and G13 are straight. Groups G5 and G13 include a section with stepped waveguide stubs crossing the main waveguide group orthogonally.

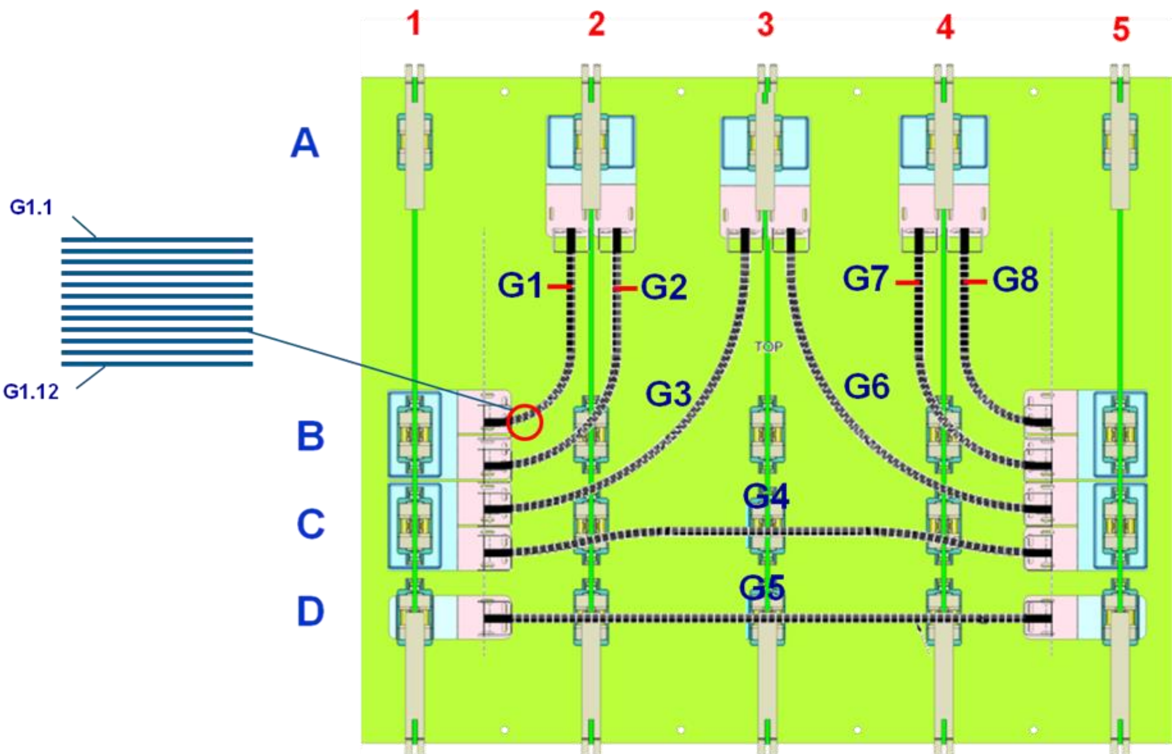


**Figure 5-3: SEPPLANE1 with waveguide groups G1 to G8, test card slots (vertical) 1 – 5 and connector positions A to D. Individual waveguides in each group are always numbered 1 – 12 going from top to bottom in the horizontal component section of the waveguide group**

Although the same waveguide layout was used on both SEPPLANE1 and SEPPLANE2, the size and shape of the glass panel cut-outs on SEPPLANE2 allowed the groups G4 and G5 on SEPPLANE1 to be split into two separate shorter groups G11, G12, G13 and G14 on SEPPLANE2.

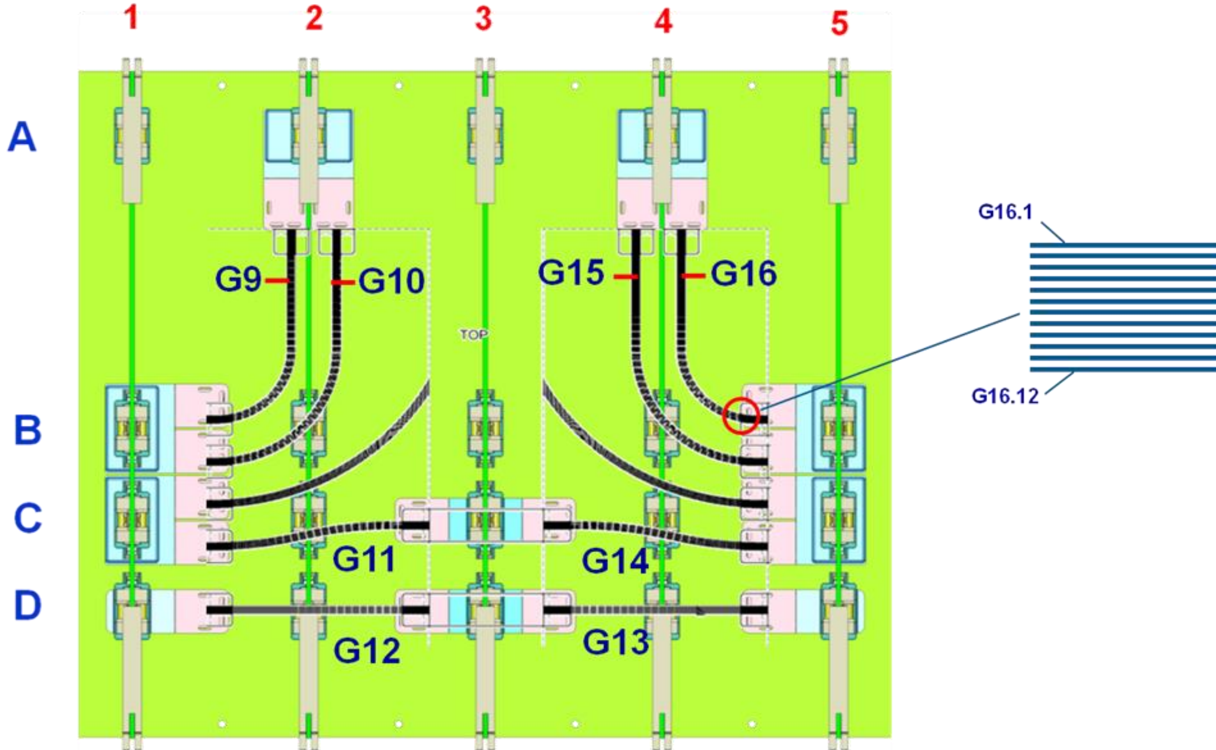


**Figure 5-4: SEPPLANE2 with waveguide groups G9 to G16, test card slots (vertical) 1 – 5 and connector positions A to D. Individual waveguides in each group are always numbered 1 – 12 going from top to bottom in the horizontal component section of the waveguide group**

The groups G3 and G6 on SEPPLANE1 were not used on SEPPLANE2 as the waveguides did not intersect orthogonally with the glass edge facet. Figure 5-5 shows schematically how non-orthogonal waveguide-to-interface intersections will result in a waveguide cross-section and channel-to-channel pitch both increased by a factor of 1/cosØ, where Ø is the angle formed between the interface and the line normal to the waveguide axes. Therefore when coupling to a standard MT ferrule compliant optical interface the channel centres will be mis-aligned.

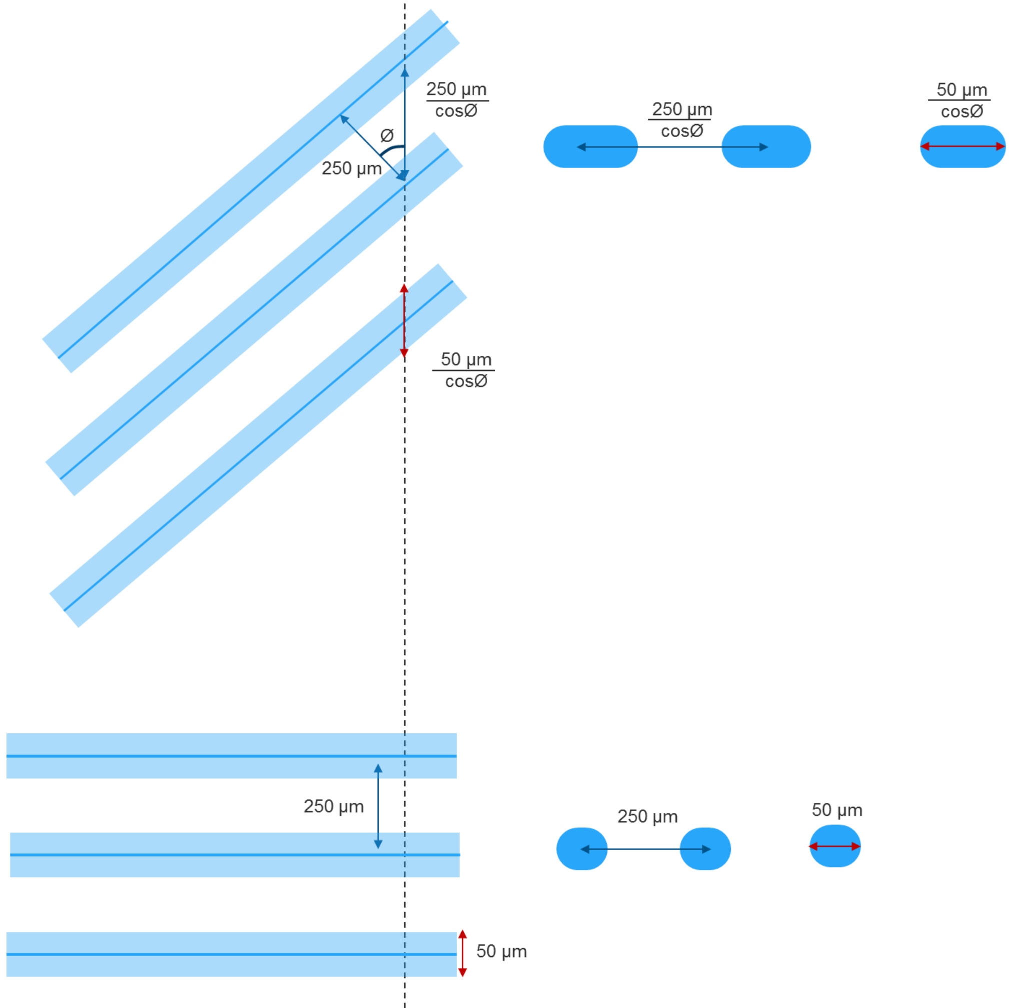


**Figure 5-5: Change of optical interface dimensions for non-orthogonal waveguide - interface intersections**

## 5.3 *Glass waveguide fabrication*

The optical waveguide layout designs were prepared by Xyratex and sent to Fraunhofer IZM, who needed to first fabricate the graded index waveguides on a glass panel, prior to sending it to the PCB fabricator ILFA to be embedded into an OPCB.

### 5.3.1 Glass waveguide fabrication process flow

A glass waveguide fabrication process was developed by Fraunhofer IZM for panel sizes of 210 × 297 mm². At the time of writing, commercial glass panel thicknesses could be as low as 30 µm [164], however the glass panels were inherently fragile and in order to reduce the risk of damage during handling, the first generations of waveguides were fabricated on more mechanically robust 500 µm thick panels.

According to the Fraunhofer IZM fabrication process, when the glass surface is exposed to a molten salt, a thermo-chemical ion exchange process takes place between the glass and the salt-melt, which gives rise to a localised increase in refractive index on the glass surface [167]. The glass matrix must contain a sufficient concentration of ions that are chemically exchangeable with a counterpart ion in the salt-melt, in order to achieve a refractive index change suitable for waveguiding. The type of glass chosen was Schott D263Teco, a borosilicate glass containing monovalent sodium ions.

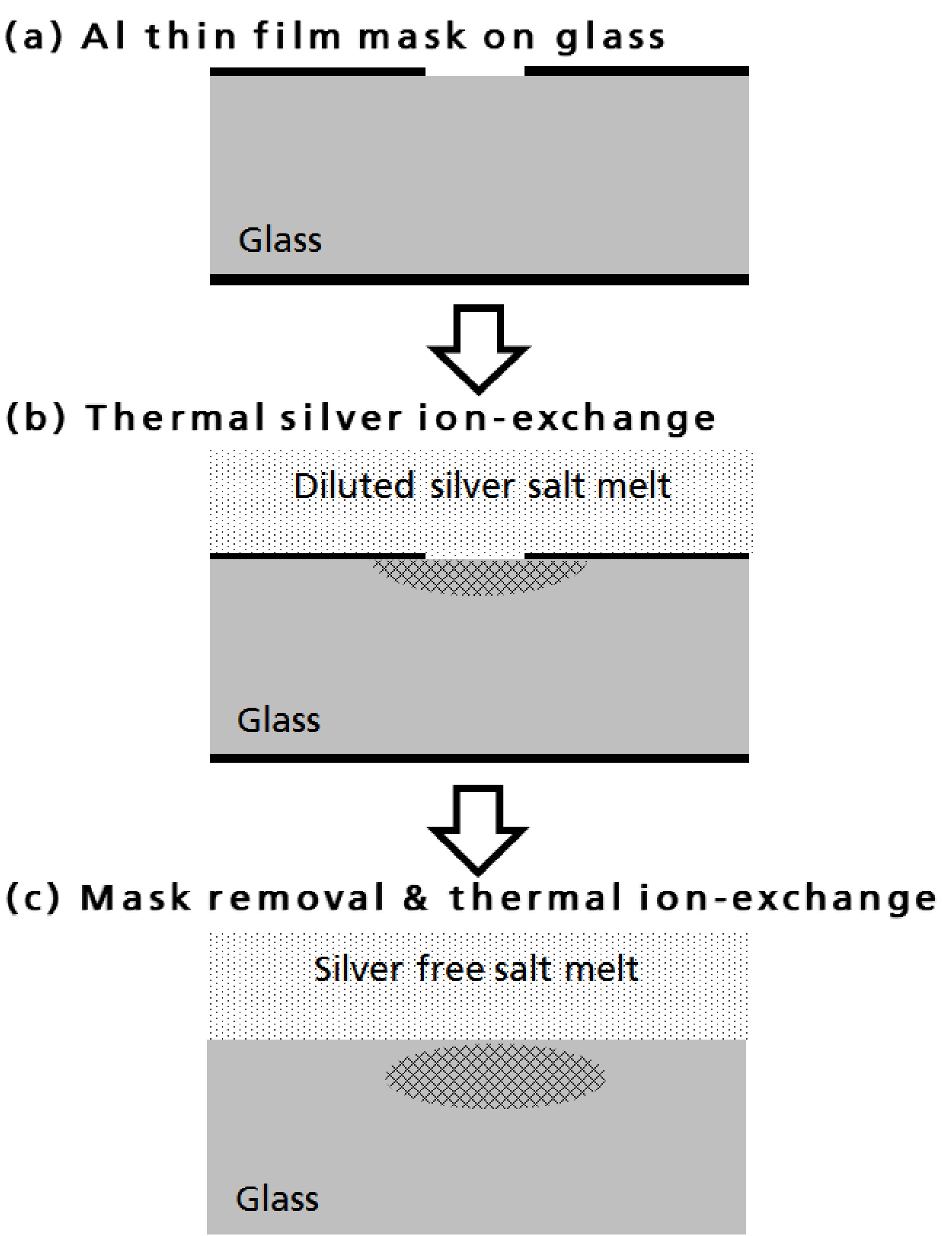


**Figure 5-6: Glass waveguide fabrication process flow (Source: Fraunhofer IZM)**

The waveguide fabrication process consisted of two thermal ion-exchange steps between the salt-melt and the glass. The positions of the waveguides were defined by a salt-melt resistant aluminium (Al) thin film mask, deposited on the glass surface as shown in the schematic process flow in Figure 5-6a. The masked glass was dipped into a hot silver salt melt in the temperature range of 300 °C to 400 °C, which served as the source of silver ions for the diffusion process. The difference in concentration between the sodium ions in the glass matrix and the silver ions in the salt melt precipitated a thermo-chemical ion-exchange process, whereby the silver ions diffused isotropically into the glass forming a flattened semi-cylindrical graded refractive index profile in the area of the mask opening (Figure 5-6b). The refractive index decreased with distance from the diffusion interface at the exposed glass surface from its peak value down to the value of the bulk glass refractive index. Key figures of merit, such as the refractive index increase and diffusion depth were determined by the silver salt concentration, process duration and temperature. The diffusion process ended when the masked glass was removed from the salt-melt. After removal of the mask layer by chemical wet etching, a second thermal ion-exchange process was performed, in which the panel was immersed into a silver free salt-melt as shown in Figure 5-6c. A reverse diffusion process took place as the silver ions were drawn back out of the glass, causing a decrease of refractive index at the diffusion interface on the exposed glass surface and shifting the position of the refractive index maximum from the glass surface deeper into the glass. By rounding off the refractive index in this manner, this final process step shapes the mode profile of the waveguide to both reduce propagation loss along the waveguide by confining the light within the bulk glass, and more closely match the mode profile of glass fibres, thus reducing the coupling loss between the waveguide and glass fibre when in physical contact.
A cross-section of fabricated glass waveguides with a centre-to-centre pitch of 250 µm is shown in Figure 5-7. Because of the isotropic thermal diffusion behavior, the final waveguide profile has an elliptical cross section with a graded refractive index distribution. The process parameters applied in this case caused the refractive index maximum to form at a depth of 18 µm below the glass surface giving rise to a maximum difference between peak refractive index and that of the bulk glass in the range of 0.016. The cross section in the vertical dimension matches that of a 50/125 µm graded index multimode fibre with N.A= 0.2. In Figure 5-7b), x = 100 µm represents the top surface of the glass.

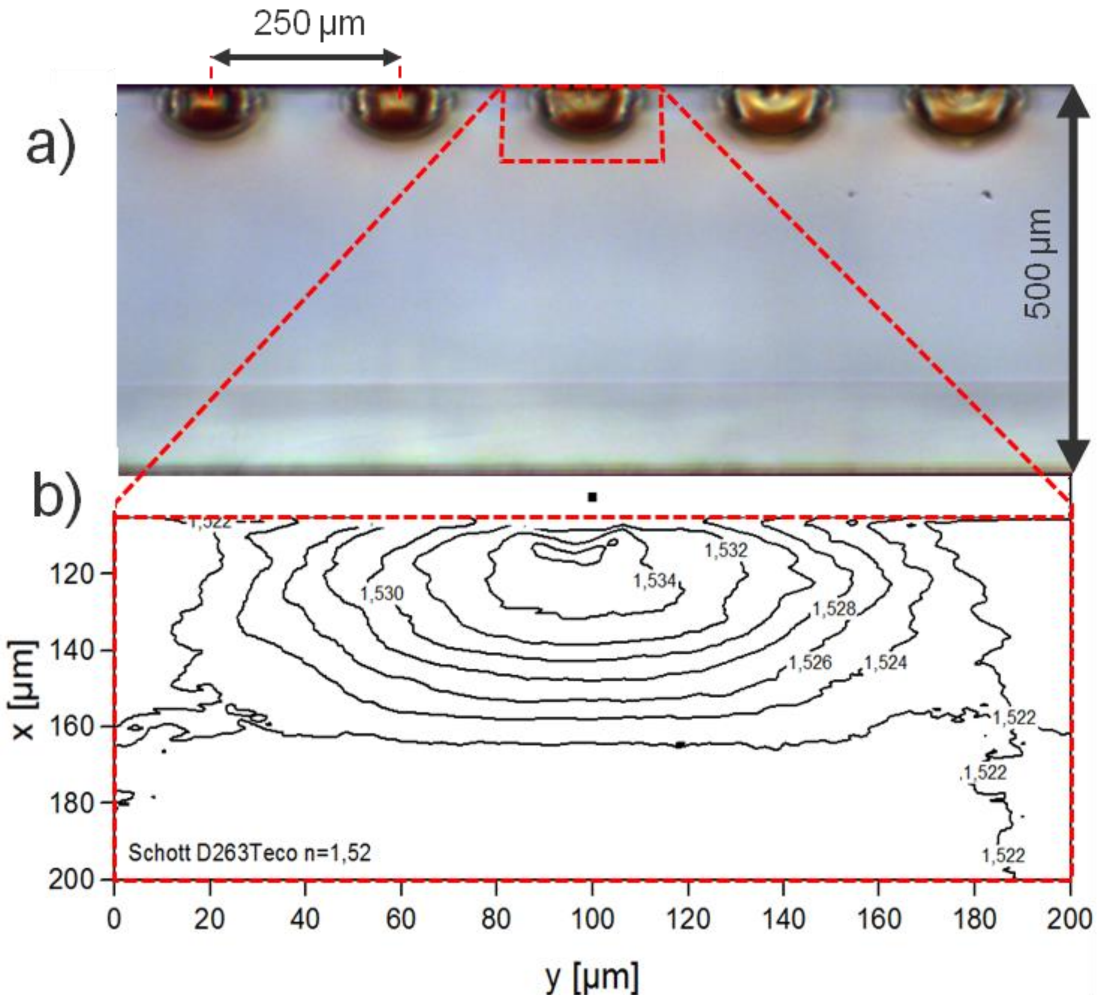


**Figure 5-7: a) Cross-section of glass waveguide sample with thickness of 500 µm and centre-to-centre waveguide pitch of 250 µm. b) Refractive index profile at measurement wavelength of 678 nm as measured by a refractive near field (RNF) scan (Source: Fraunhofer IZM)**

### *5.3.1.1. Glass waveguide panel fabrication*

For the SEPPLANE backplane waveguide layout, a functional waveguide area was reserved for a 199 mm×160 mm central portion of a 297×210 mm² glass panel. The layout included fiducials for process alignment as well as partner logos. In Figure 5-8 the glass thin film mask layout is shown in black with mask openings in white.

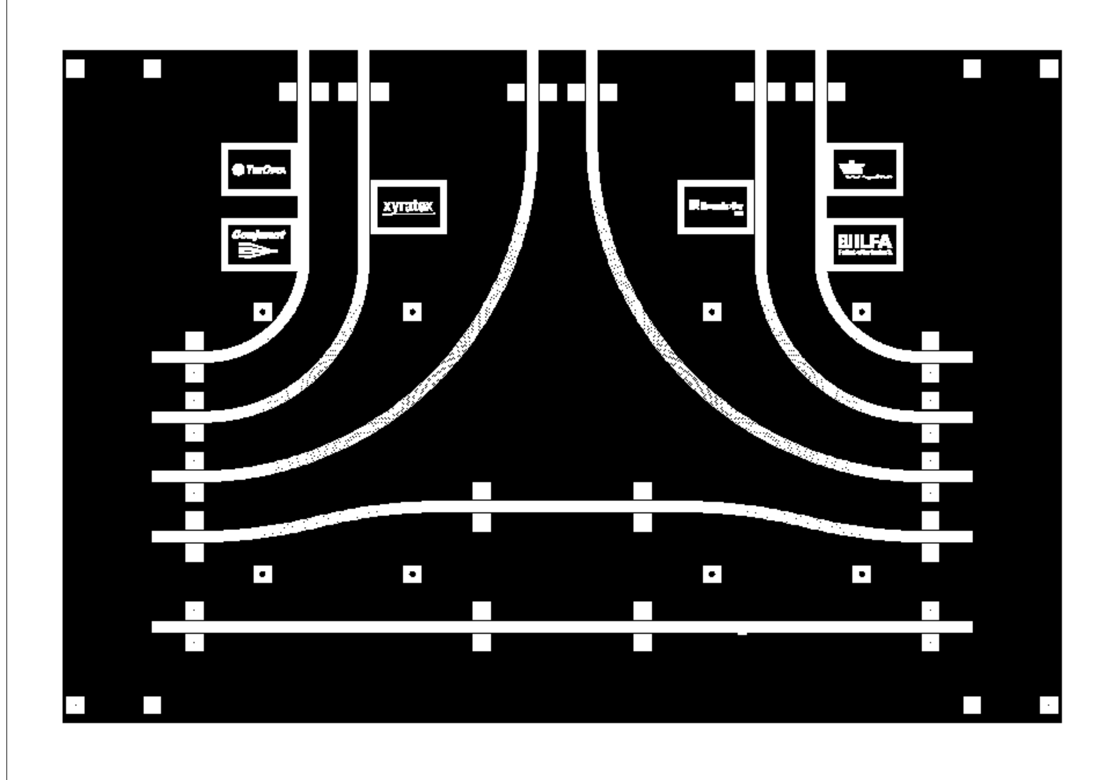


**Figure 5-8: Layout diffusion mask on glass containing waveguide layout, partner logos and alignment marks (Source: Fraunhofer IZM)**

After a cleaning treatment step, an aluminium thin film of 400 nm thickness was sputtered on both sides of the glass panel. A photoresist layer was then deposited uniformly over the aluminium layer through a dip-coating process. In order to convey the desired waveguide layout pattern onto the photolithographic layer, a 14 inch chrome-on-glass mask was placed in contact with the glass panel on the vacuum chuck of an Orbotec Paragon 9000 laser direct imaging (LDI) system and contact mask exposure carried out. A positioning mount was used to align the smaller glass panel to the centre of the 14 inch mask as shown in Figure 5-9.

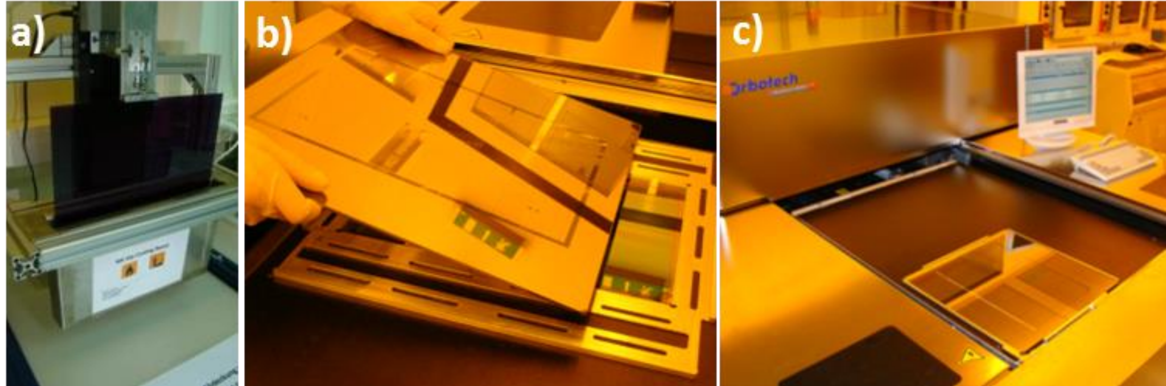


**Figure 5-9: a) Dip-coating, b) positioning mount and c) mask exposure with LDI (Source: Fraunhofer IZM)**

After exposure, the photoresist was developed and the underlying aluminium layer was structured with acid before removal of the photoresist. In order to account for the isotropic diffusion characteristics under the mask opening, a mask gap of 30 μm width was required to create a 50 μm glass waveguide. The aluminium diffusion mask layer on the glass panel is shown in Figure 5-10.

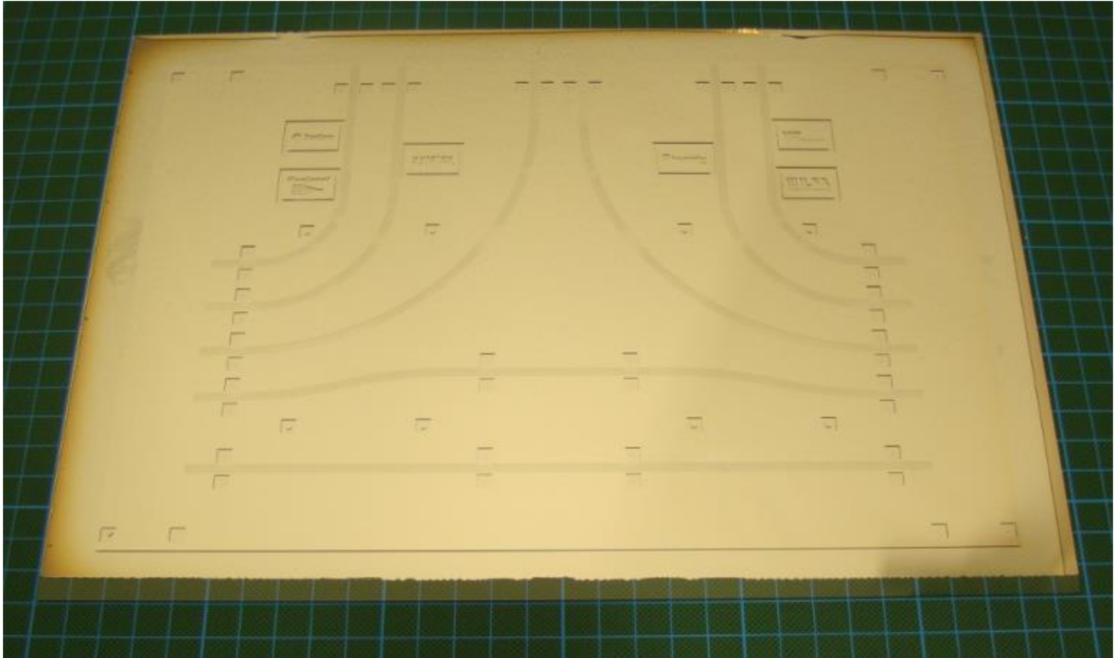

**Figure 5-10: Aluminium diffusion mask layer on glass panel (297 x 210 mm²) (Source: Fraunhofer IZM)**

Multiple aluminium masked glass panels were inserted vertically into the hot salt-melt and the ion-exchange process described above carried out to create the sub-surface graded index waveguide layout defined by the

mask pattern. The aluminium mask was then removed from the glass surface except around the partner logos and fiducial marks required for alignment and assembly (Figure 5-11b). In the second ion diffusion step, the glass panel without the mask layer was inserted vertically into a silver free salt-melt, whereby, as previously described, the reverse ion-exchange process caused the refractive index maximum (waveguide core centre) to shift to a point 18 µm below the glass surface.

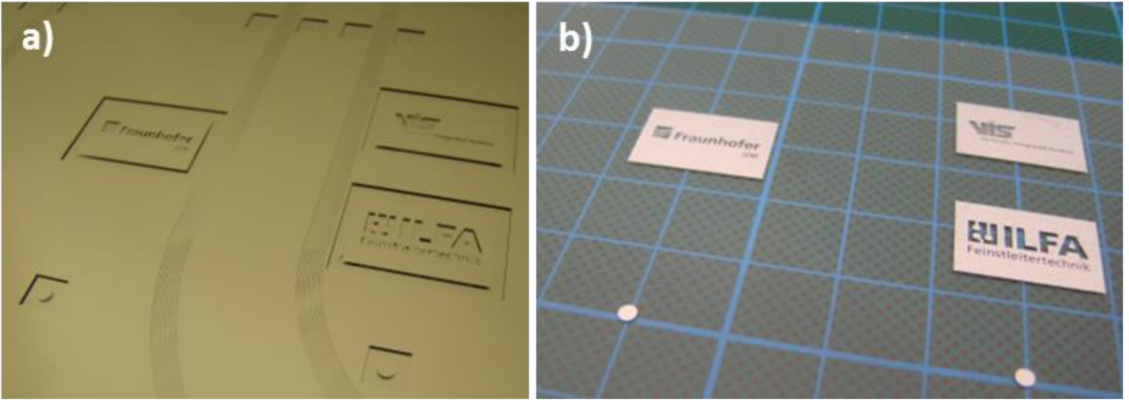


**Figure 5-11: a) Detail of mask layer on glass, b) protected partner logos and fiducial marks after removing aluminium mask layer (Source: Fraunhofer IZM)**

Finally, the 297 x 210 mm² glass waveguide panel was trimmed with a CO2-laser cutting process [168] down to the panel size required for the OPCB backplane type: a) for SEPPLANE1, a single panel of 199 x 160 mm² or b) for SEPPLANE2, two smaller panels of 79.25 x 160 mm², which would be embedded with a horizontal separation of 40.5 mm. When compared to diamond reel cutting, the CO2-laser cutting process produced around a four-fold increase in edge strength and trace stability in the processed glass panels. The glass edges processed with this cutting method exhibited an RMS surface roughness of 40 nm, which was sufficiently low to ensure low-loss optical coupling [155] to the waveguides.

### 5.3.2 Waveguide characterisation by Fraunhofer IZM

The waveguide characterisation of the glass panel by Fraunhofer IZM prior to embedding into the OPCB backplane stack-up is hereby included as a reference for later measurements.

Each waveguide core had an elliptic graded index profile, with the longer axis parallel to the glass surface due to the isotropic ion diffusion along the mask gap. The propagation and coupling losses on the fabricated glass

waveguide panel were characterised by Fraunhofer IZM at a wavelength of 1310 nm using the "cut-back method", which will now be explained.

Propagation loss is defined as the inherent attenuation of the light as it propagates along a waveguide.

Coupling loss is defined as the amount of light lost upon entering or leaving the waveguide due, for the most part, to Fresnel and scattering effects at the waveguide interface.

Insertion loss is the total loss measured across the waveguide, which includes both coupling loss and propagation loss. The cut-back method allows the propagation loss and the coupling loss on a given waveguide to be separated by comparing the insertion loss results on different lengths of the same waveguide sample. If the same cleaving process is used, then it can be reasonably assumed that the loss characteristics at each waveguide interface are the same and therefore can be assumed to represent a constant loss value on each insertion loss measurement with only the length dependent propagation loss contributing to the linear increase in insertion loss (expressed in dB) with increasing waveguide length. The coupling loss from both input and output interfaces can be extrapolated as the value of loss at the intercept between the linear insertion loss profile with the y-axis.

The results of the measurements using the cut-back method are shown in Figure 5-12. The propagation loss at 1310 nm was measured to be 0.05 dB/cm and coupling loss was 2.15 dB for a 50 µm core graded index multimode (GI MM) fibre launch. The results are comparable with those of previous work [83]. In addition, the waveguides were characterised at a wavelength of 850 nm with propagation loss measured to be 0.41 dB/cm and coupling loss 2.7 dB for a GI MM fibre launch. The higher propagation loss at 850 nm is most likely due to the formation of silver ion clusters in the glass matrix, the sizes of which give rise to stronger Rayleigh scattering of 850 nm light than at 1310 nm light. Fraunhofer IZM expects that the propagation losses can be further reduced in future by a factor of 2 – 4 by process improvement and glass selection. Following the work reported in this thesis, Fraunhofer IZM started evaluating Gorilla glass from Corning, as a candidate for planar waveguide fabrication.

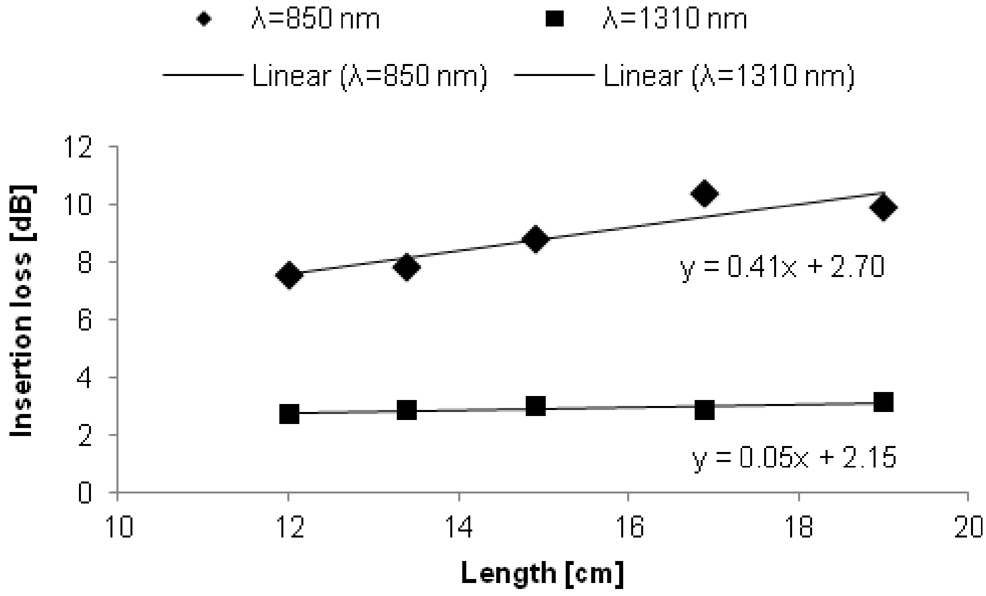


**Figure 5-12: Propagation loss and coupling loss measurements based on best fitting line of cut-back measurements (Source: Fraunhofer IZM)**

The insertion loss measured using both SM (9/125µm step-index fibre, NA=0.11) and MM (50/125µm, graded-index fibre, NA=0.2) launching conditions for all waveguide groups G1-G8 and corresponding bend radii are summarized in Table 5-1. The propagation losses for the MM launch were extrapolated based on the coupling loss of 2.15 dB predicted in the cut-back measurements shown in Figure 5-12 and are in agreement with the propagation loss predicted by those measurements. The results indicated an increase in propagation loss with decreasing bend radius. The waveguide crossings on group G5 with a maximum number of 10 intersections did not show significant increase of insertion loss. These results confirm estimations based on previous research [83]. Furthermore, comparison of waveguide groups with the same geometry indicate that this fabrication process gives rise to good uniformity across the panel, though improvements are still possible to reduce the value of deviation.

**Table 5-1: Insertion, propagation and bend losses at 1310 nm of waveguide groups in SEPPLANE1 panel dependent on bend radii and launching condition (Source: Fraunhofer IZM)**

| Group | Bend Radii | Insertion loss with SM launch | Insertion and Propagation loss with MM launch | Calculated loss in bend section for MM launch |
|---|---|---|---|---|
| G1 | 25–27.75 mm | 1.21 ± 0.24 dB | 2.62 ± 0.010 dB<br>0.052 dB/cm | 0.23 ± 0.10 dB |
| G2 | 41–43.75 mm | 1.10 ± 0.14 dB | 2.75 ± 0.10 dB<br>0.052 dB/cm | 0.36 ± 0.10 dB |
| G3 | 88.125 – 90.875 mm | 1.15 ± 0.06 dB | 2.92 ± 0.27 dB<br>0.049 dB/cm | 0.75 ± 0.27 dB |
| G4 | 2 S-bends: RoC = 134.0 mm, arc angle =14.04° | 1.76 ± 0.21 dB | 2.98 ± 0.09 dB<br>0.042 dB/cm | 0.49 ± 0.09 dB |
| G5 | no bend, crossings | 1.72 ± 0.17 dB | 3.08 ± 0.21 dB<br>0.047 dB/cm | - |
| G6 | 88.125 – 90.875 mm | 1.22 ± 0.26 dB | 2.83 ± 0.20 dB<br>0.043 dB/cm | 0.66 ± 0.20 dB |
| G7 | 41 – 43.75 mm | 0.96 ± 0.14 dB | 2.74 ± 0.13 dB<br>0.05 dB/cm | 0.35 ± 0.13 dB |
| G8 | 25 – 27.75 mm | 1.04 ± 0.11 dB | 2.90 ± 0.20 dB<br>0.082 dB/cm | 0.51 ± 0.20 dB |

## *5.4 Electro-optical circuit board fabrication using low temperature lamination processes*

On the SEPIANet project, German PCB fabricator ILFA GmbH was responsible for embedding the glass waveguide panels produced by Fraunhofer IZM into a PCB stack-up to ensure that the integrity of the glass panel and the properties of the waveguides embedded therein could withstand the thermal or mechanical stresses inherent to the lamination process. The incorporation of different materials into a PCB layer stack-up can lead to problems due to dimensional expansion of the individual layers during the thermal bonding process. The disparity between the coefficients of thermal expansion (CTE) in different materials and the high bonding temperature of up to 250 °C during lamination cycles can lead to substantial built-in stress and/or bowing and twisting within the PCB.

Combining different materials with similar horizontal CTEs reduces inter-material stresses and is thus inherent to PCB fabrication. This is evident when comparing the CTEs of copper ($\alpha_{Cu}$ = 16.5 $10^{-6}$ $K^{-1}$) and FR4 ($\alpha_{FR4}$ = 12…14 $10^{-6}$ $K^{-1}$), which are of a similar order. However the incorporation of materials with very different CTEs, such as glass ($\alpha_{D263Teco}$ = 7.2 $10^{-6}$ $K^{-1}$) will lead to substantial built in stresses under normal high temperature lamination processes, which, given the inherent fragility of glass and the influence of mechanical strain on refractive index, would be highly detrimental to optical transmission and signal integrity. Figure 5-13 shows a photo of an early OPCB prototype developed using a standard hot lamination process with a section of the embedded glass foil exposed in a square recess in the surrounding materials. The glass section is visibly bulging due to the inherent mechanical stresses caused by the different material expansions exhibited between the glass and the FR 4 materials bounding it as the whole assembly was heated up.

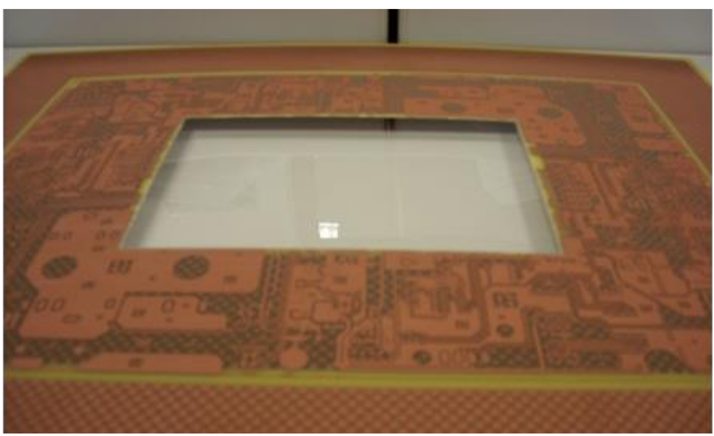

**Figure 5-13: Glass foil laminated into PCB stack-up with conventional hot lamination technique shows bulging (Source: ILFA GmbH)**

ILFA therefore developed a substrate bonding technology based on 50 µm thick, low temperature activated adhesive foils to avoid thermally induced material expansion of the different materials, which is detailed in [4], [165].

By combining lamination process temperatures of around 40°C and high pressure loads during the lamination process, the relative thermal expansion of the individual layers in the stack-up was not sufficient to give rise to deleterious levels of built-in stress within the manufactured OPCB. Therefore this low temperature bonding process was deployed in the manufacture of the OPCBs on the SEPIANet project, none of which subsequently exhibited the CTE mismatch induced built-in stresses observed in the early prototypes.

The final OPCB stack-ups of SEPPLANE1 and SEPPLANE2 with a total thickness of 3.5 mm are shown in Figure 5-14. The full PCB boards of both variants measured 28 mm x 23.5 mm. Both variants consisted of two electrical packages with solder mask, each with two electrical layers, and one optical package with either one glass substrate for SEPPLANE 1 or two glass substrates in the same plane for SEPPLANE 2. Figure 5-14d shows an unpopulated SEPPLANE 1. The OPCBs were both optically and electrically functional with the encased fragile glass panels shielded from direct external mechanical forces. The prototype exhibited no built-in stress, with no apparent bulging or twisting of the embedded glass, thus showing the advantages of the newly developed cold lamination technology developed by ILFA.

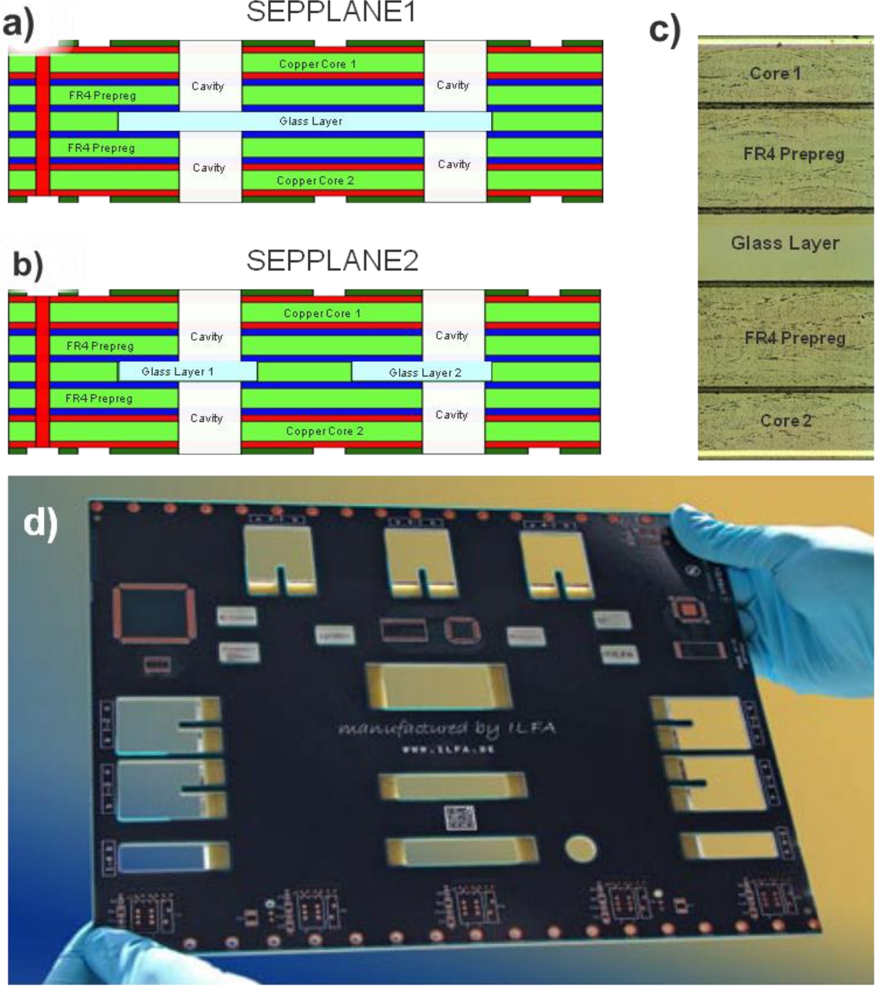


**Figure 5-14: a) SEPPLANE1 stack-up with one single panel optical and two electrical packages with solder mask, b) SEPPLANE2 stack-up with one dual panel optical and two electrical packages, c) Photo of OPCB cross-section, d) photo of unpopulated OPCB backplane circuit board**

Although the size of waveguide panels deployed on these backplanes was restricted to 199 x 160 mm² by available mask sizes used for the LDI and ion-exchange process, ILFA have successfully integrated glass panels with a form factor of 570 mm x 420 mm and have also demonstrated that up to four smaller glass panels can be integrated into the same EOCB layer. Furthermore ILFA and IZM have demonstrated that they can vertically embed up to eight glass layers into a PCB stack-up [169].

## *5.5 Pluggable optical connector design and assembly*

A complete passive optical connector system was designed and developed to enable test daughtercards to be plugged to a glass waveguide based OPCB and to enable parallel optical fibre cables to be directly connected to the embedded waveguide interfaces themselves (Figure 5-15). The design was tailored to satisfy the opto-mechanical requirements of the fragile glass waveguide panel embedded in the electro-optical circuit board stack-up.

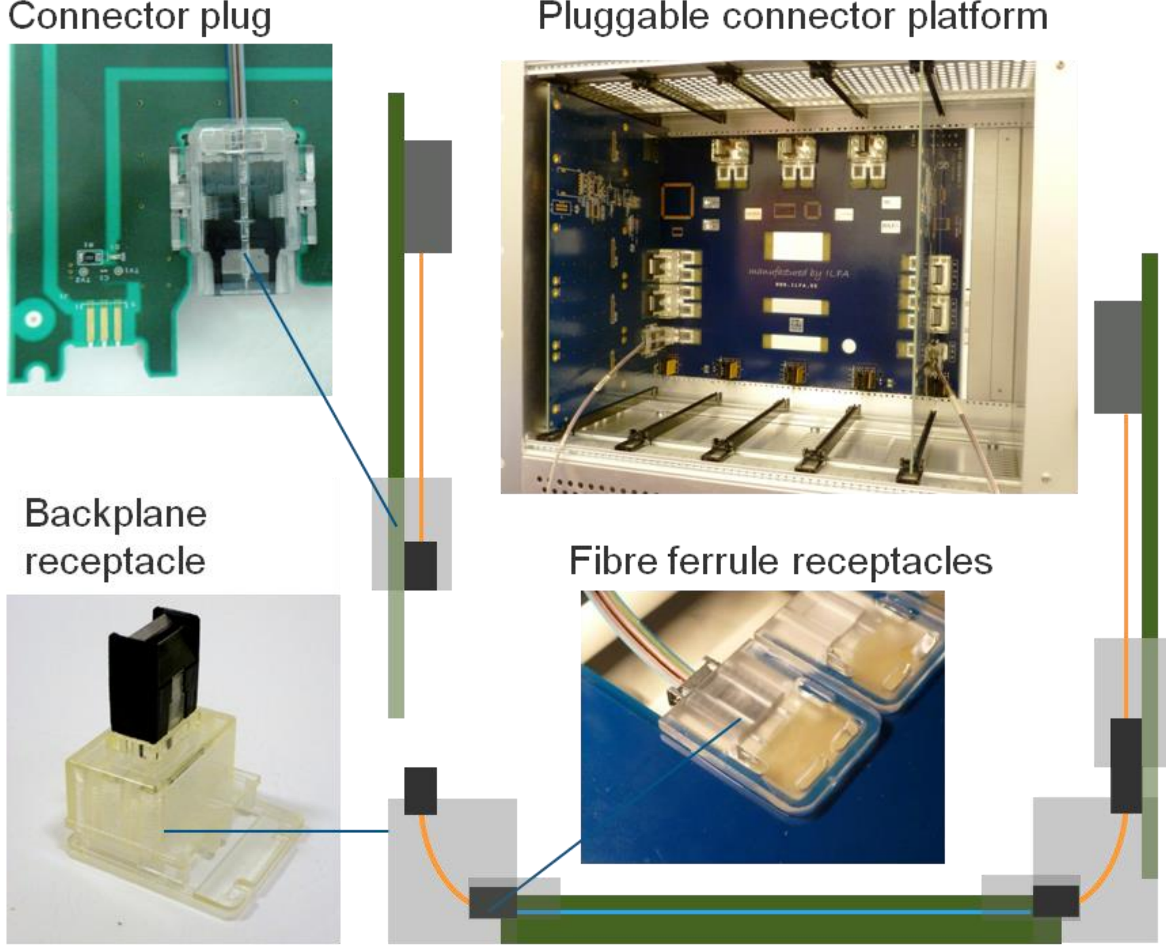


**Figure 5-15: Complete pluggable connection system comprising connector plug on connecting edge of daughtercard, backplane receptacle and fibre ferrule receptacles**

### 5.5.1 MT ferrule compliant waveguide receptacle

The author developed the design requirements specification for the waveguide receptacle to be directly assembled onto the glass foil edge over the embedded waveguide array. Various views of the receptacle are shown in Figure 5-16.

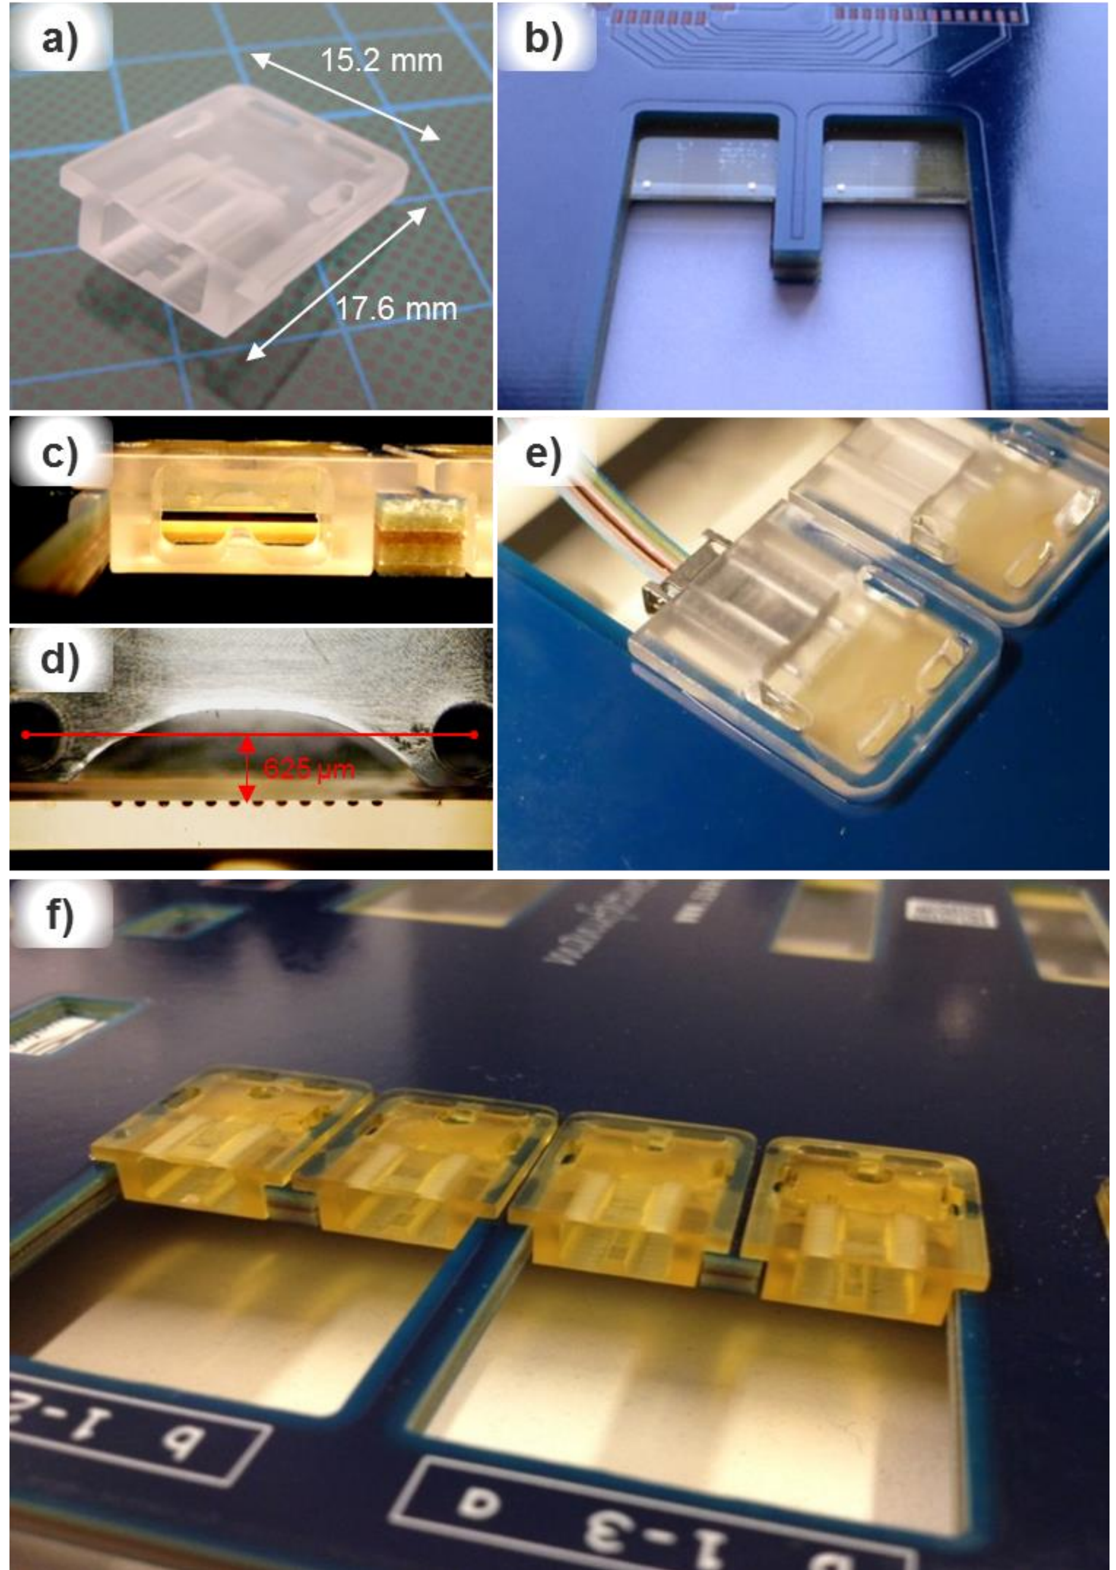


**Figure 5-16: a) MT ferrule receptacle made of polycarbonate material, b) exposed dual glass waveguide array interfaces, c) front view of waveguide array interface with MT compliant receptacle port assembled, d) vertical offset between MT pins datum and waveguide array, e) dual ferrule mounts with MT fibre jumper plugged into one, f) four adjacent waveguide ferrule receptacles made of Ultem™ material**

#### *5.5.1.1. Receptacle design*

The receptacle was designed to accommodate a commercial MT style ferrule (Figure 5-16c), however, in order to avoid cutting notches into the glass panel either side of the waveguide group, the receptacle needed to accommodate an offset between the alignment pins and the waveguide array, which could be matched by a commercial 6 x 12 MT ferrule. When the ferrule receptacle is placed on the top surface of the glass panel, the datum between the alignment pins of the connecting ferrule must be vertically offset from the embedded waveguide array as shown in Figure 5-16d. This offset must be 625 µm, in order to match the offset in a 6 x 12 MT (72 way) fibre interface between the datum between the alignment pins and the lowest row of the ferrule plugged into the receptacle.

This is the same design principle on which the first generation waveguide receptacle for polymer waveguides was based as described in Chapter 3.

The waveguide receptacle design was realised by opto-mechanical engineer Allen Miller at Xyratex and the prototypes were fabricated in polycarbonate through a precision machining process by high precision component fabricator Optima Pacific Optima Asia Pacific Sdn Bhd.

#### *5.5.1.2. Material choice*

Receptacles were fabricated from two different materials: polycarbonate (Figure 5-16a), which appeared as an opaque white colour, and Ultem™ (Figure 5-16f), which appeared as a yellow colour. Ultem™ proved to be an unsuitable material for the Fraunhofer IZM assembly process described below in Section 5.5.3 as it is highly absorptive of UV light and as such prevented curing of photosensitive adhesive through the receptacle. Polycarbonate however was much less absorbing of UV light and as such allowed better curing of the adhesive through the receptacle.

As such polycarbonate was the preferred choice for high precision waveguide receptacles on this project.

### 5.5.2 Board-to-board connector plug and backplane receptacle design

A board-to-board connection system was invented and designed by Allen Miller at Xyratex.

The connector plug section was mounted on the edge of a connecting daughtercard (Figure 5-17a) and comprised an internal brace and an outer housing section. The brace supported an MT style parallel optical fibre ferrule and attached fibre cable (Figure 5-17b) and contained an interlocking shutter system to reduce dust contamination of the ferrule interface.

The backplane receptacle comprised a shuttered head section (black), which houses its own ferrule terminated fibre ribbon and a clear housing section (Figure 5-17d).

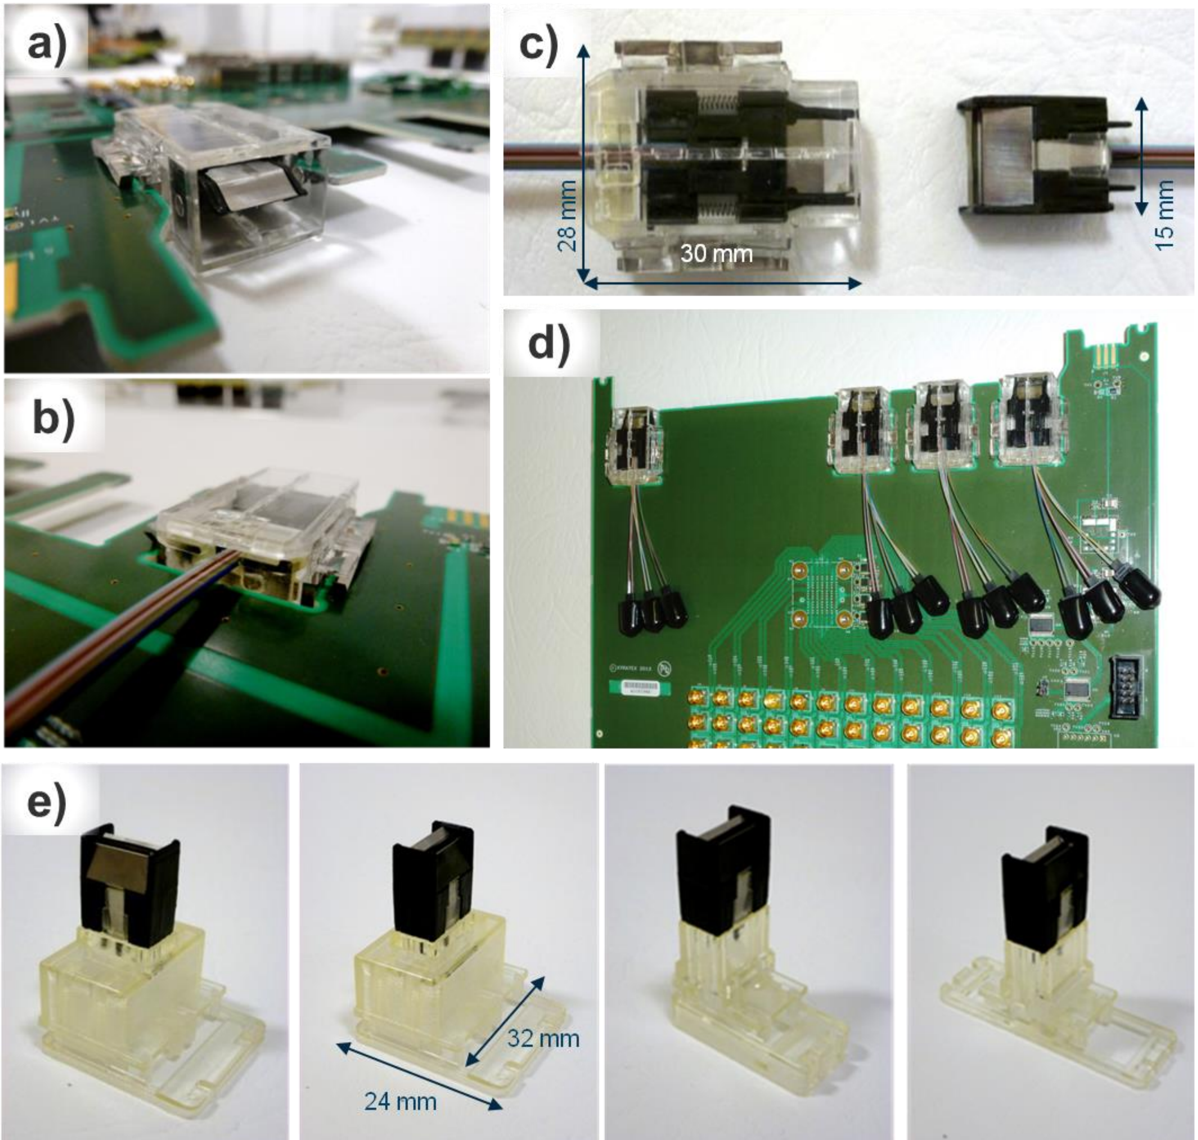


**Figure 5-17: a) Pluggable connector assembled on card edge (front view), b) rear view of pluggable connector showing fibre ribbon from internal MT ferrule, c) pluggable connector and receptacle head, d) connector plugs assembled on edge of test daughtercard, e) 4 variants of full backplane receptacle with receptacle head (black) in different orientations relative to receptacle body (clear)**

During engagement, the coarse alignment structures in the plug housing section allow it to capture the compliant receptacle head, (Figure 5-17c), mutually opening the shutters of both plug and receptacle and bringing the high precision alignment structures of the internal commercial MT ferrules into engagement. An internal spring

system on the plug brace provides a degree of over-travel to ensure the ferrule interfaces are held together under force when the daughtercard is engaged.

The backplane receptacle housing (Figure 5-17e) is designed to incorporate a short dual fibre patchcord bent by 90° to enable the required deflection of optical signals between the optical axes of the waveguide interfaces (parallel to the backplane) and the daughtercard mounted connector plugs (orthogonal to the backplane). As shown in Figure 5-18, the short dual fibre jumpers provide a fully fibre populated 2x12 MT ferrule for vertical connection to the daughtercard plug, which fans out into two separate rows of 12 fibres, each terminated to the lowest row of a 6 row MT ferrule. This arrangement allows one vertical daughtercard connector plug to connect to two horizontal dual adjacent (Figure 5-18b) or opposite (Figure 5-18c) waveguide interfaces thus increasing interconnect density.

This technique of optical right angle deflection is robust, reliable and mechanically decouples the fragile glass waveguide interface from direct stresses and strains inherent to the board to board connections.

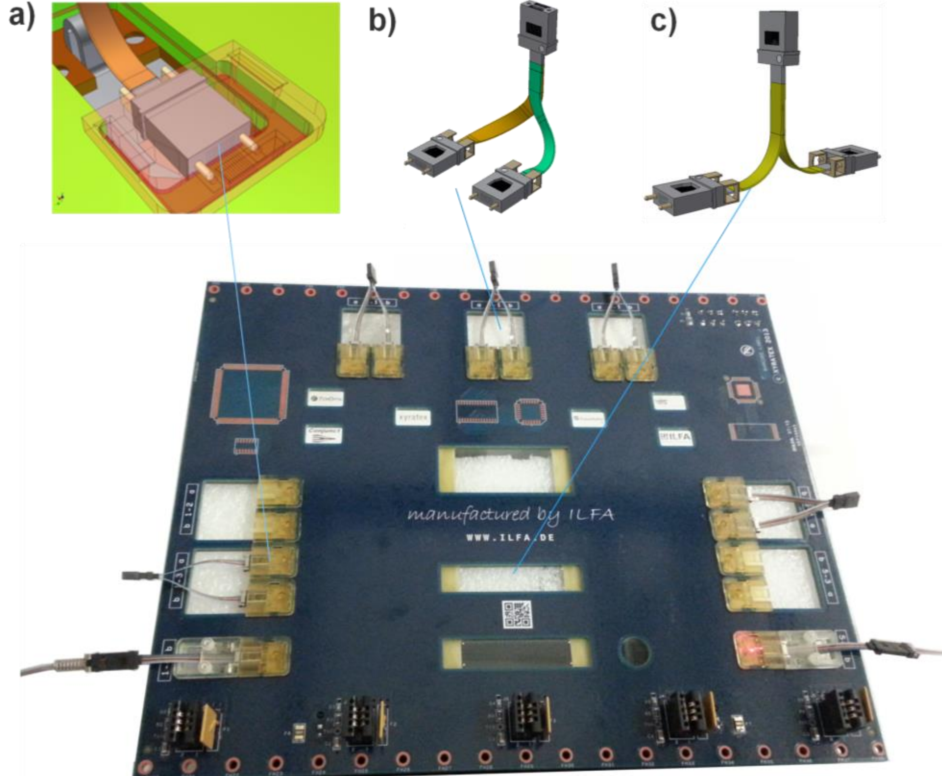


**Figure 5-18: OPCB backplane type SEPPLANE 1 with ferrule receptacles and jumpers populated, a) schematic view of ferrule receptacle populated with MT terminated fibre ribbon, b) schematic view of dual edge board fibre jumper (connecting to 2 adjacent interfaces), c) schematic view of dual mid-board fibre jumper (connecting to 2 opposite interfaces)**

Four different (2 single and 2 dual jumper) backplane receptacle variants were developed to accommodate different board positions and connector configurations.

### 5.5.3 Board assembly of fibre ferrule receptacles

Fraunhofer IZM carried out the assembly of the MT fibre ferrule receptacles onto the waveguide array interfaces using a proprietary active alignment and assembly routine to decrease misalignment and optimise coupling efficiency between the fibre-optic MT patchcord and 12 channel waveguide array. Their assembly process relied on pick-and-place assembly equipment supplied by ficonTEC with three translational and two rotational axes and a positional accuracy better than ±1 µm. A vacuum tool was designed to hold the MT ferrule receptacle with an MT terminated fibre patchcord in situ. The ficonTEC equipment included an adhesive dispensing system and UV curing as well as top and bottom vision control cameras for alignment purpose. For the assembly task, the equipment was supplemented with a separate 3-axis translation stage to position a second fibre-optic MT

patchcord at the launch facet of the waveguide group. For that, a two channel laser source and photodetector with an operating wavelength of 1310 nm were connected to the two MM fibre-optic MT patchcords for in-situ insertion loss measurement. During the assembly process only 1310 nm light was used to carry out insertion loss measurements.

The semi-automated assembly routine consisted of the following process steps:

a) Backplane suspension (Figure 5-19a).

Electro-optical backplane is placed on the working stage within the ficonTec assembly rig.

b) Launch fibre ferrule alignment (Figure 5-19b).

An MT terminated 12 fibre ribbon is actively aligned by a 3-axis translation stage over the selected waveguide group. Active alignment of the entire 12 channel fibre array to the waveguide array is achieved by adjusting the positions of the two illuminated outer fibres in the launch ferrule (fibre 1 and fibre 12) such as to maximise light coupled into the corresponding outer waveguides in the group, and monitoring light exiting those waveguides from the other waveguide interface with a vision control camera. These serve as reference channels.

c) Detection fibre and ferrule receptacle alignment (Figure 5-19c).

In order to align the ferrule receptacle for assembly, it is populated with a 6 row MT ferrule terminated with a 12 fibre ribbon on the lowest row and positioned into the PCB clearance area near the waveguide interface. Pre-positioning of the ferrule onto the waveguide interface is achieved by scanning fiducial marks on the glass panel and calculating the corresponding position. An active alignment routine adjusts the position of the fibre array and attached ferrule receptacle to the waveguide array until the insertion loss through both the illuminated reference channels as measured through the corresponding butt-coupled fibres is minimised.

d) Ferrule receptacle assembly on detection side (Figure 5-19d).

Once the ferrule receptacle is in precise position, UV-curable adhesive is applied between the glass surface and receptacle at key positions and, through UV exposure, the ferrule receptacle is rigidly bonded in place onto the electro-optical backplane.

e) Ferrule receptacle assembly on launch side (Figure 5-19e)

With a ferrule receptacle now permanently in place over one waveguide interface, the process steps c) and d) are repeated to align and assemble the launch ferrule receptacle onto the first waveguide interface. Upon completion, both waveguide interfaces bounding a given waveguide group have been successfully connectorised.

This assembly process was repeated to fully connectorise all waveguide groups on the electro-optical backplanes.

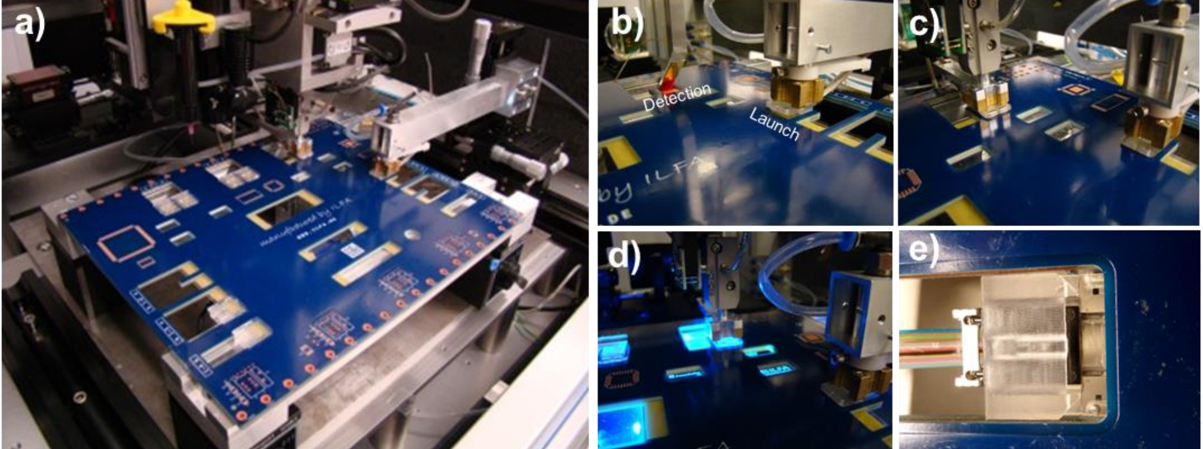


**Figure 5-19: Ferrule receptacle assembly process steps: a) backplane suspended in ficonTEC assembly rig, b) active launch fibre alignment with 3-axis translation stage, c) active detection fibre and ferrule receptacle alignment with 5-axis translation stage, d) MT ferrule receptacle glued in place, e) fibre patch cord plugged into assembled MT multi-fibre ferrule**

### 5.5.4 Assembly of connector plugs and receptacles

The connector plug housings were passively attached into compliant cut-outs on the connecting edge of the daughtercard, with the metal side clips on the housing providing a limited float of the plug relative to the boards. Figure 5-20 shows a fully assembled “SEPPLANE1” OPCB backplane with 16 MT ferrule receptacles bonded (Figure 5-20a) and 9 backplane receptacles assembled (Figure 5-20b).

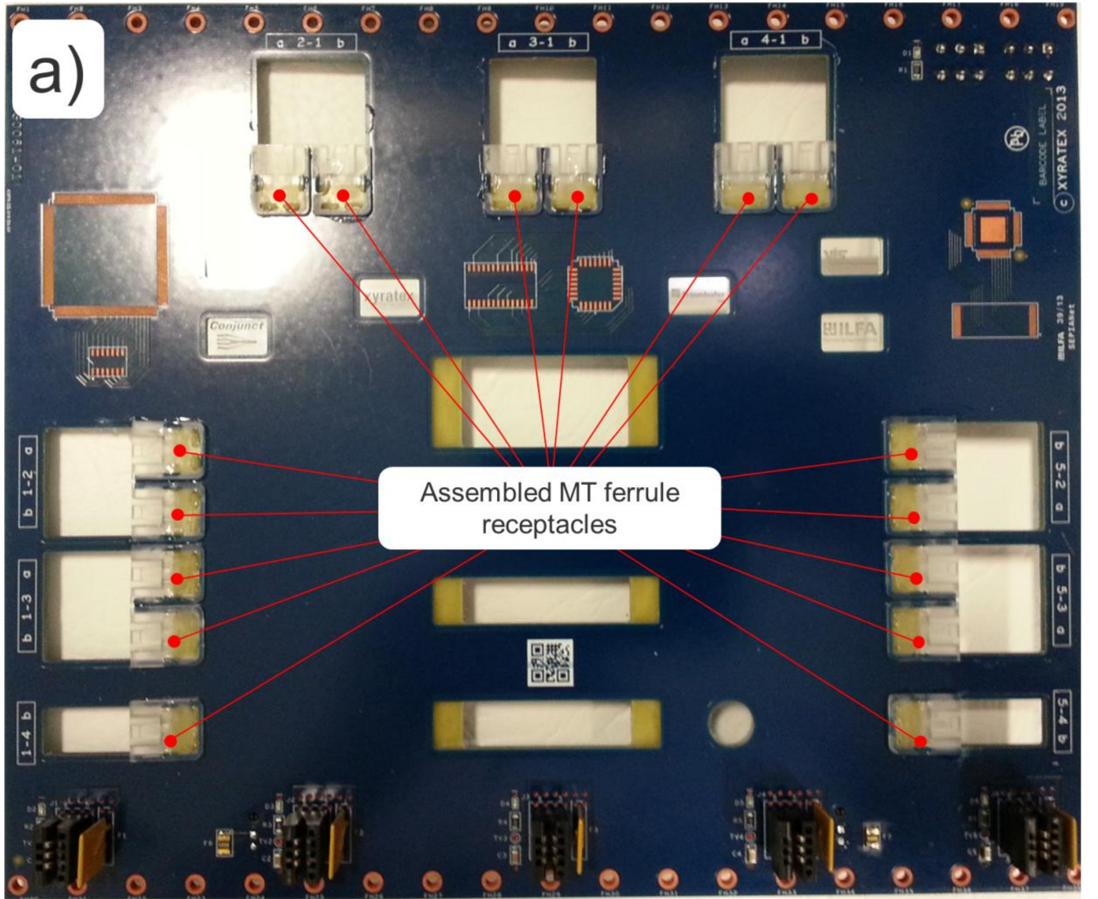


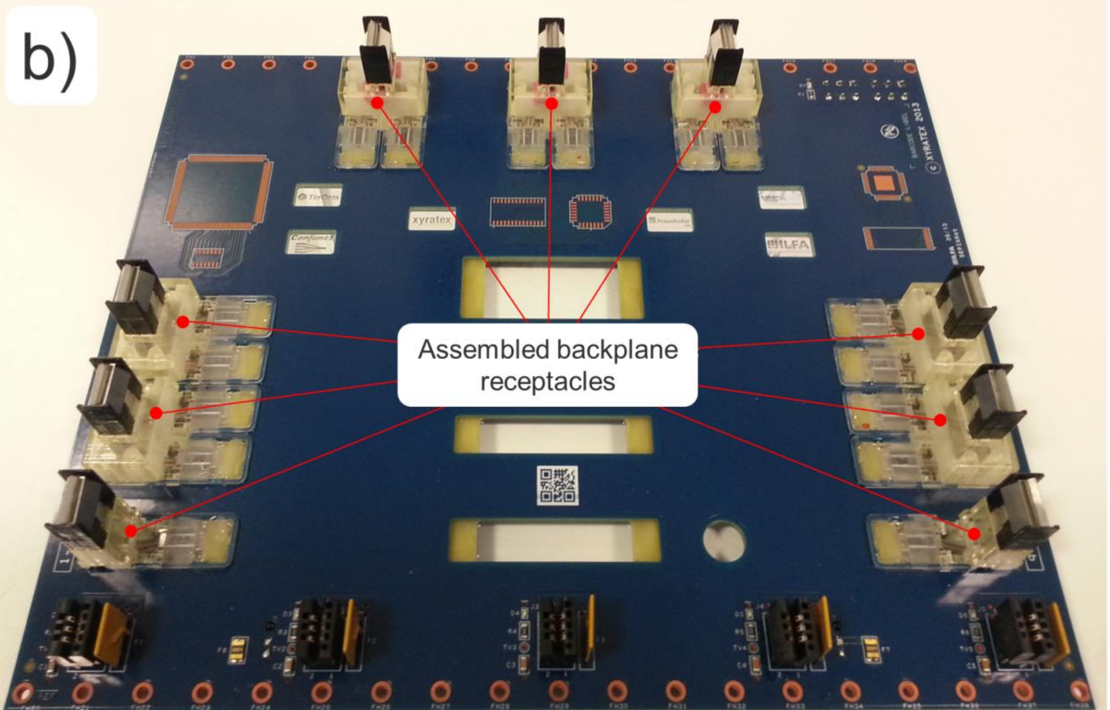


**Figure 5-20: a) SEPPLANE1 after assembly of all MT ferrule receptacles, b) SEPPLANE1 after assembly of all backplane receptacles**

The 6×12 MT ferrule ends of the dual fibre jumpers shown in Figure 5-18 were then passively connected into the MT ferrule receptacles and the vertical 2×12 MT ferrule end mounted into the backplane receptacle head section. A spring clip was provided to ensure that the MT ferrules were held in place against the glass waveguide interface. The backplane receptacle housings were then fastened to the backplane with screws. The connector plug and receptacle MT ferrule to MT ferrule engagement mechanism was tested and validated independently on a mechanical test rig showing no deterioration of insertion loss or misengagements over 100 mating cycles.

### 5.5.5 SEPIANet platform design

The author prepared the design requirements specification for the SEPIANet platform, which comprised:

- 7U (311.15 mm) high, 84 HP (426.72 mm) sub-rack chassis with an integrated fan tray and power supply
- Electro-optical backplane, type (SEPPLANE1 or SEPPLANE2) populated with 1) electronic backplane receptacles to provide power and low-speed data to the daughtercards, 2) optical ferrule receptacles for direct cable attach and 3) backplane receptacles to enable board-to-board pluggability
- 5 test cards based on the Euro-card form factor. Three variants of test card were designed supporting an extensive variety of test data interfaces, optical transceiver subassemblies and optical connector plugs. Each test card can accommodate 4 edge connector plugs.
- All optical fibre jumpers used were based on multimode, 50/125 µm OM3 type fibre

Two test and measurement platforms, SEPDEM1 and SEPDEM2 were developed and assembled to respectively allow full optical characterisation of the SEPPLANE1 and SEPPLANE2 type backplane and pluggable interconnect systems.

**Figure 5-21: Two test and measurement platforms SEPDEM1 and SEPDEM2**

Overall, three backplanes were fabricated and characterised: two of type SEPPLANE1 (designated SEPPLANE1a, SEPPLANE1b) and one of type SEPPLANE2 (designated SEPPLANE2a) were fabricated and characterised in the test and measurement platforms.

#### *5.5.5.1. Test daughtercards*

The author developed the design requirements specification for five different varieties of test daughtercard. The designs were rendered by Xyratex electronic designer Paul Stevens.

The five different varieties of test daughtercard were designated **SEPTEST1**, **SEPTEST1A**, **SEPTEST1C**, **SEPTEST2** and **SEPTEST3**.

The test daughtercards were designed to support a variety of different data interfaces on the front end to allow different data streams to be injected into the system including 10 GbE LAN (10.3 Gb/s) and 12 G SAS (12 Gb/s), as well as user specific test data loads spanning a range of possible data rates up to 25 Gb/s. The front end data interfaces include Mini-SAS HD cable ports, Quad Small Form-factor Pluggable (QSFP) ports and high frequency SMP connectors.

Each of these data interfaces could accommodate 4 duplex channels and can thus each be connected to a dedicated midboard optical engine permanently mounted to the board.

An LED backlight card was also designed to provide back-illumination for the six partner logos when the system was powered up.

**SEPTEST1** (Figure 5-22a) housed two types of mezzanine cards, **SEPTEST1A** and **SEPTEST1C** (Figure 5-22b). Each mezzanine card was designed to house a Conjunct optical engine and alignment brace with a data interface to convey test data electronically to and from the engine. The data interface port on SEPTEST1A was a mini-SAS HD connector allowing SAS data to be conveyed to the engine for electro-optical conversion. The data interface ports on SEPTEST1C were high speed (40 GHz) SMP connectors allowing 28 Gb/s test data to be conveyed to the engine for electro-optical conversion.

**SEPTEST2** (Figure 5-22c) was designed to house two Avago McLink midboard transceivers (12G, 2Tx + 2Rx) and two finisar BOA (12G, 12 Tx + 12 Rx) midboard transceivers. These transceivers were supplied by three data interface port types: 5 Mini-SAS HD ports (supporting 20 duplex channels), 1 QSFP port (supporting 4 duplex channels) and 16 SMP ports (supporting 4 duplex channels).

**SEPTEST3 was designed to** house one high speed Finisar BOA (25G, 12Tx + 12Rx) to which external electronic test data could be conveyed via 48 SMP ports for electro-optical conversion.

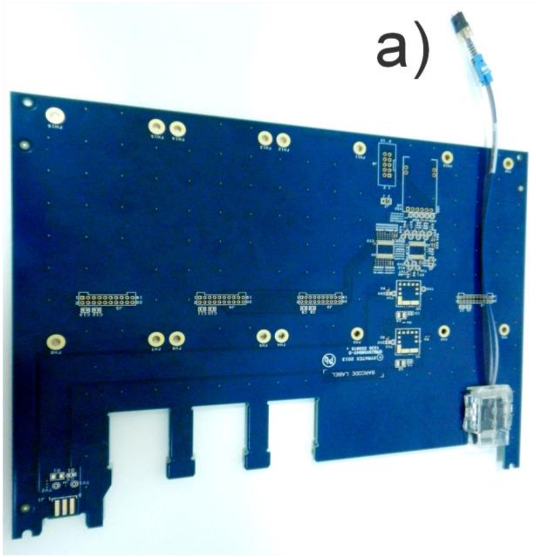


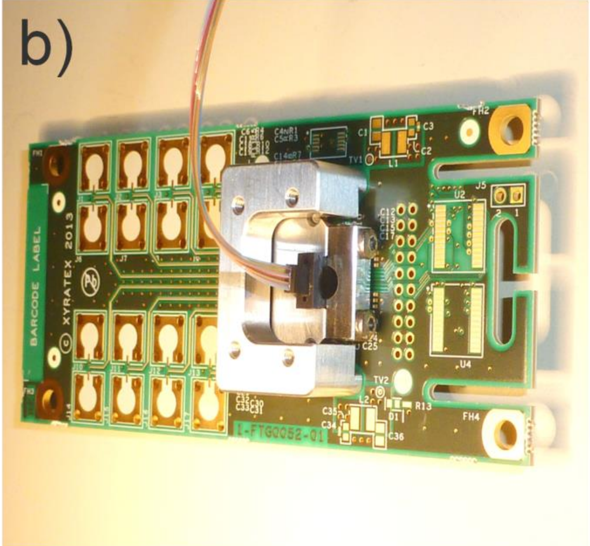


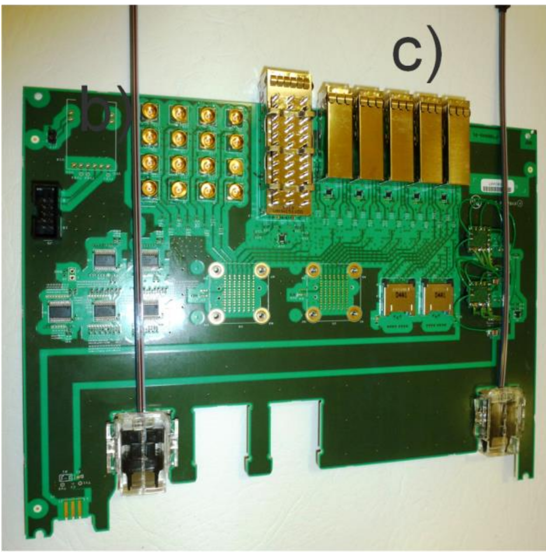


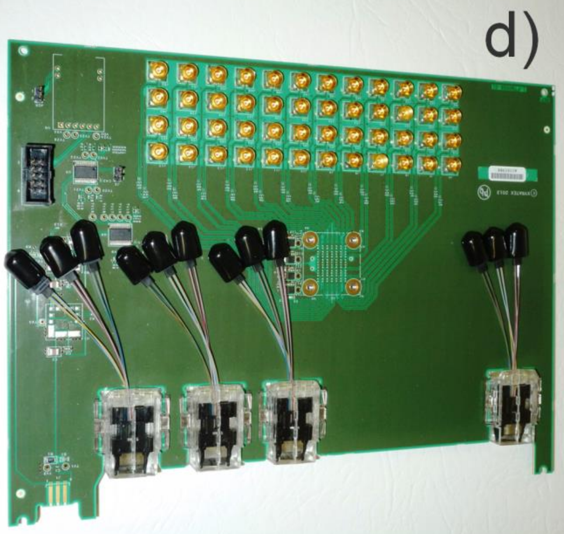


**Figure 5-22: SEPIANet test cards: a) SEPTEST1 test card with vacant slots for mezzanine cards and populated with 1 edge connector plug, b) SEPTEST1C mezzanine card populated with alignment brace and MT patchcord attached, c) SEPTEST2 populated with 2 edge connector plugs, d) SEPTEST3 fully populated with 4 edge connector plugs**

The purpose for designing this variety was to also accommodate other 3rd party optical engines thus providing alternative optical sources on the test cards driven by the same data test pattern. This was part of an extensive risk mitigation strategy to decouple the risk of a failure or non-delivery of one or more partner transceiver deliverables and enable partial demonstration of partner deliverables within the demonstration platform.

SEPTEST1A and SEPTEST1C could not be completed due to delays by one of the SEPIANet partners to deliver functional optical engines. SEPTEST3 could not be completed due to the high speed midboard optical transceiver assembly not being delivered.

Limited testing with on-board engines was only carried out on SEPTEST2, however it was determined that the most useful testing would be based on the use of external test data being conveyed directly to the edge connectors of any of the test cards, as opposed to via an on-board transceiver. The reason for this is that it allowed for a truer characterisation of the board-to-board pluggable performance of the connector and eliminated the variable performance attributed to the midboard transceiver.

## *5.6 Optical interconnect test and measurement*

The author developed a comprehensive test and measurement regime for the two platforms SEPDEM1 and SEPDEM2 and carried out measurements jointly with Kai Wang at Xyratex. In all cases, optical test data was conveyed over the complete waveguide, connector and patchcord link between the measurement points 1 and 2 and vice versa shown in Figure 5-23. The measurement points are easily accessible 2 x12 MT interfaces, to which external fibre-optic test cables can be connected by the tester without the need to reach into the enclosure. They are also chosen to emulate a mid-board transceiver that would typically be connected with a fibre-optic jumper to the backplane connector.

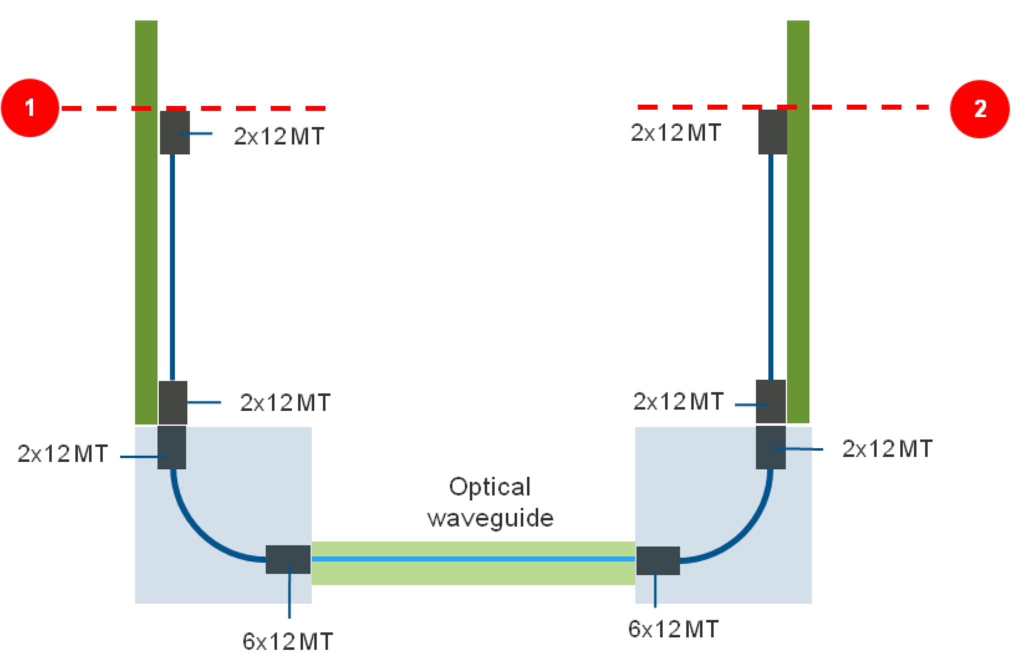


**Figure 5-23: Schematic view of complete optical connector and waveguide link under test**

### 5.6.1 Bidirectional insertion loss measurements

Insertion loss measurements were carried out at 850 nm and 1310 nm for all functional waveguides in the system. The chosen optical sources were an 850 nm Class 1 VCSEL and a 1310 nm Class 1 DFB laser. Eight insertion loss measurements were taken for each waveguide: 1) 850 nm and 1310 nm, 2) received power as measured by a large area photodetector and the integrated photo-receiver of a Tektronix CSA8000B Communications Signal Analyser (CSA) each calibrated to either 850 nm or 1310 nm depending on the source, 3) measurements in both directions along the same waveguide, in order to evaluate optical link reciprocity.

Figure 5-24 shows the results of insertion loss profiles measured on waveguide group G4 on SEPPLANE1. This is an example of one of the best performing waveguide groups tested.

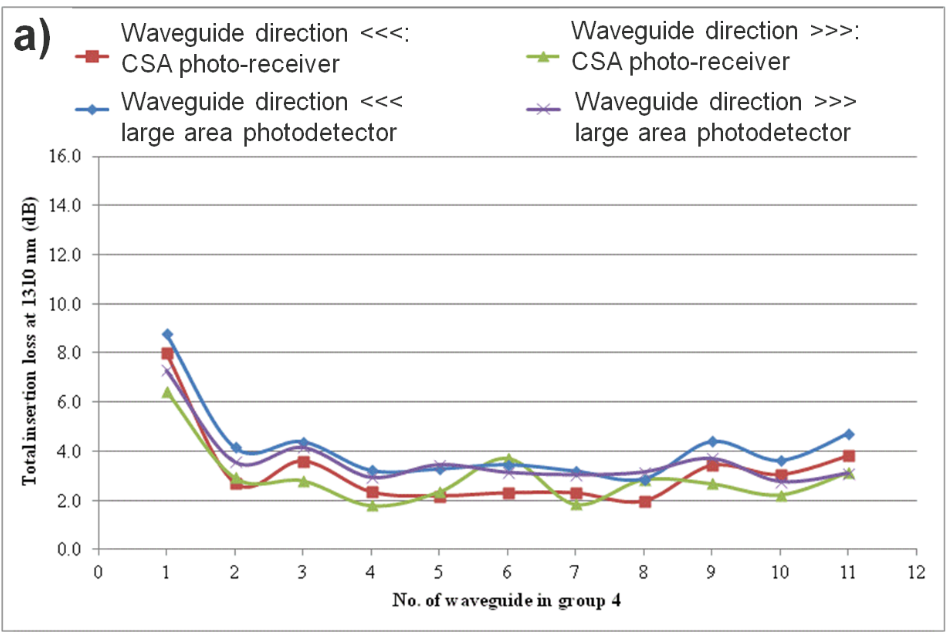


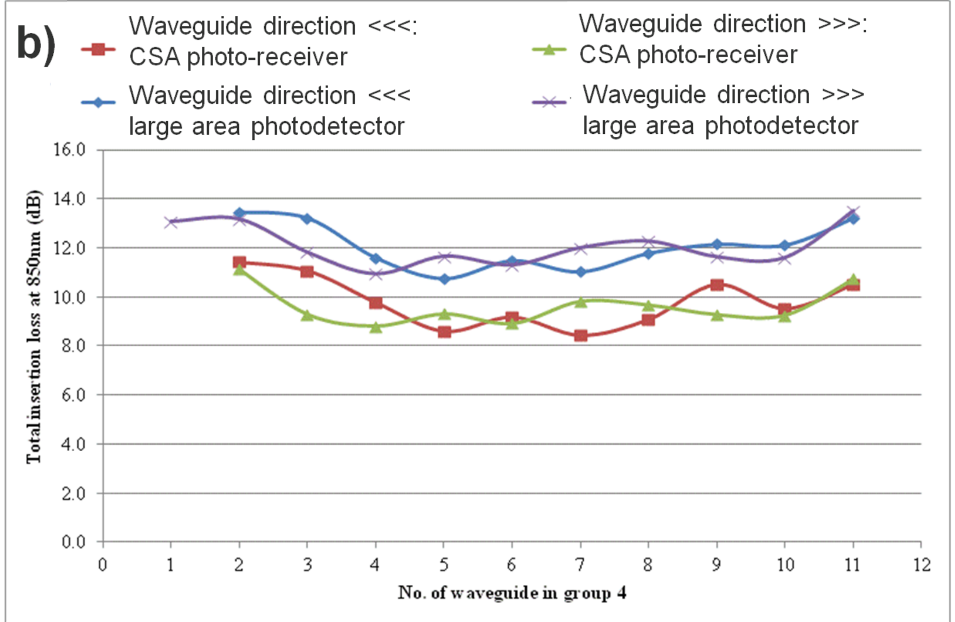


**Figure 5-24: Bidirectional insertion loss on SEPPLANE1 - waveguide group 4 in both directions as recorded directly using a large area photodetector and CSA: a) at 1310 nm, b) at 850 nm. Propagation direction >>> signifies propagation from left to right along a given waveguide as viewed from the top of the backplane, while <<< indicates right to left propagation**

Table 5-2 shows a summary of insertion loss measurement results and variability for the full optical connector and waveguide link on all waveguide groups on the 3 backplanes under test. Bidirectional measurements were taken into account, but only insertion loss measured directly with the photodetector were considered. As expected, insertion losses at 850 nm are substantially higher than those measured at 1310 nm, which is in compliance with the direct waveguide measurement results reported in Section 5.3.2. Upon inspection, it was determined that the variability in results was most likely due to sporadic glue contamination on some waveguide interfaces.

**Table 5-2: Summary of insertion loss as measured on all backplanes with 1310 nm and 850 nm optical test signals modulated at 10.3 Gb/s. The standard deviation of each measurement set is also shown to provide an indication of the measurement spread and variability of results due to contamination**

| **Backplane** | **Average insertion loss / dB** | **Standard deviation** |
|---|---|---|
| **1310 nm** | | |
| SEPPLANE1a (prior to interfaces being cleaned) | **7.8** | **3.4** |
| SEPPLANE1a (after interfaces have been cleaned) | **3.9** | **0.9** |
| SEPPLANE1b | 5.4 | 2.2 |
| SEPPLANE2a | 4 | 1.1 |
| **850 nm** | | |
| SEPPLANE1a (prior to interfaces being cleaned) | 11.1 | 1.8 |
| SEPPLANE1a (after interfaces have been cleaned) | 11.3 | 1.5 |
| SEPPLANE1b | 12.4 | 1 |
| SEPPLANE2a | 10.7 | 2.2 |

### 5.6.2 Insertion loss measurements through both daisy chained demonstration platforms

In order to also demonstrate extended passive optical connectivity between multiple electro-optical backplane systems, both demonstration platforms SEPDEM1a and SEPDEM2a were daisy-chained together with passive fibre-optic connections (Figure 5-25) and some of the best performing waveguide groups evaluated in series through the pluggable connector interfaces.

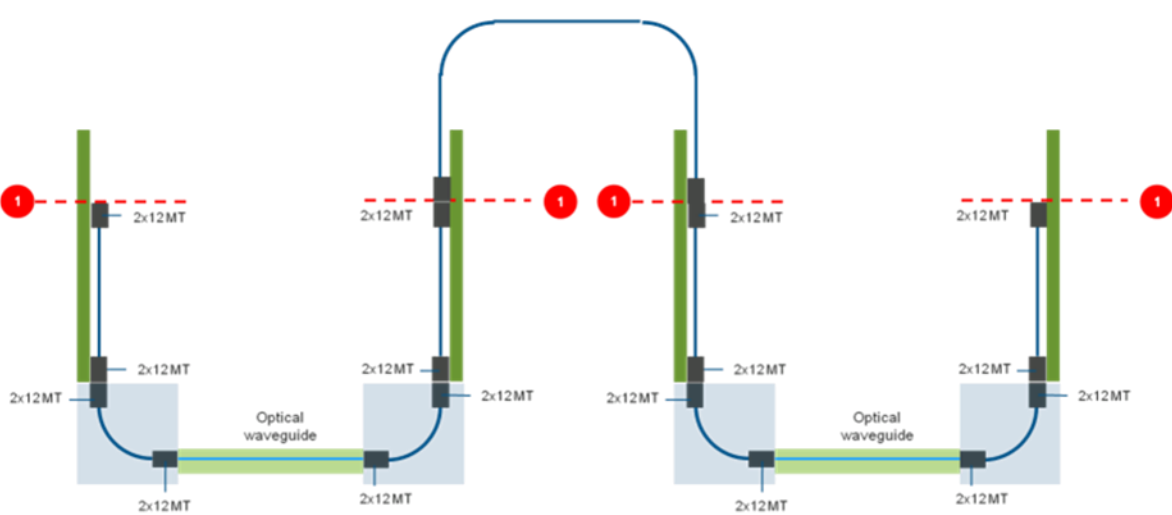


**Figure 5-25: Schematic view of complete daisy chained link under test through both SEPDEM1b and SEPDEM2a**

Propagation direction is defined by first identifying the waveguide group in question (and if appropriate, the waveguide in said group), then identifying the connector position into which light is coupled into the waveguide, followed by a hyphen and finally identifying the connector position from which the light is extracted for measurement. Connector positions are defined by the daughtercard slot number, which can take the values 1 – 5, followed by the connector row designation, which can take the values A-D as shown in Figure 5-4. For example G12 3D – 1D means measurement on waveguide group 12 launching into connector position 3D and receiving out of connector position 1D.

Figure 5-26 shows the bidirectional insertion loss and jitter measurements at 1310 nm as measured through daisy-chained link over both demonstration platforms, through SEPPLANE1b G2 in propagation direction 2A - 1B and through SEPPLANE2a G12 in propagation direction 3D-1D. The average overall insertion loss measured with a 1310 nm 10.3 Gb/s optical signal was 11.9 dB and average peak to peak jitter was 87 ps.

Figure 5-27 shows eye diagrams for waveguides 1 to 11 in this group as measured by the CSA. When the amplitude of the optical signals received by the CSA receiver becomes too low for the receive circuit to properly resolve them, then the eye diagrams begin to deteriorate. The quality of the eye diagrams in Figure 5-27 starts to deteriorate from 9 onwards, which is consistent with the increasing loss shown in Figure 5-26 a. An eye diagram for waveguide 12 could not be extracted due to excessive loss.

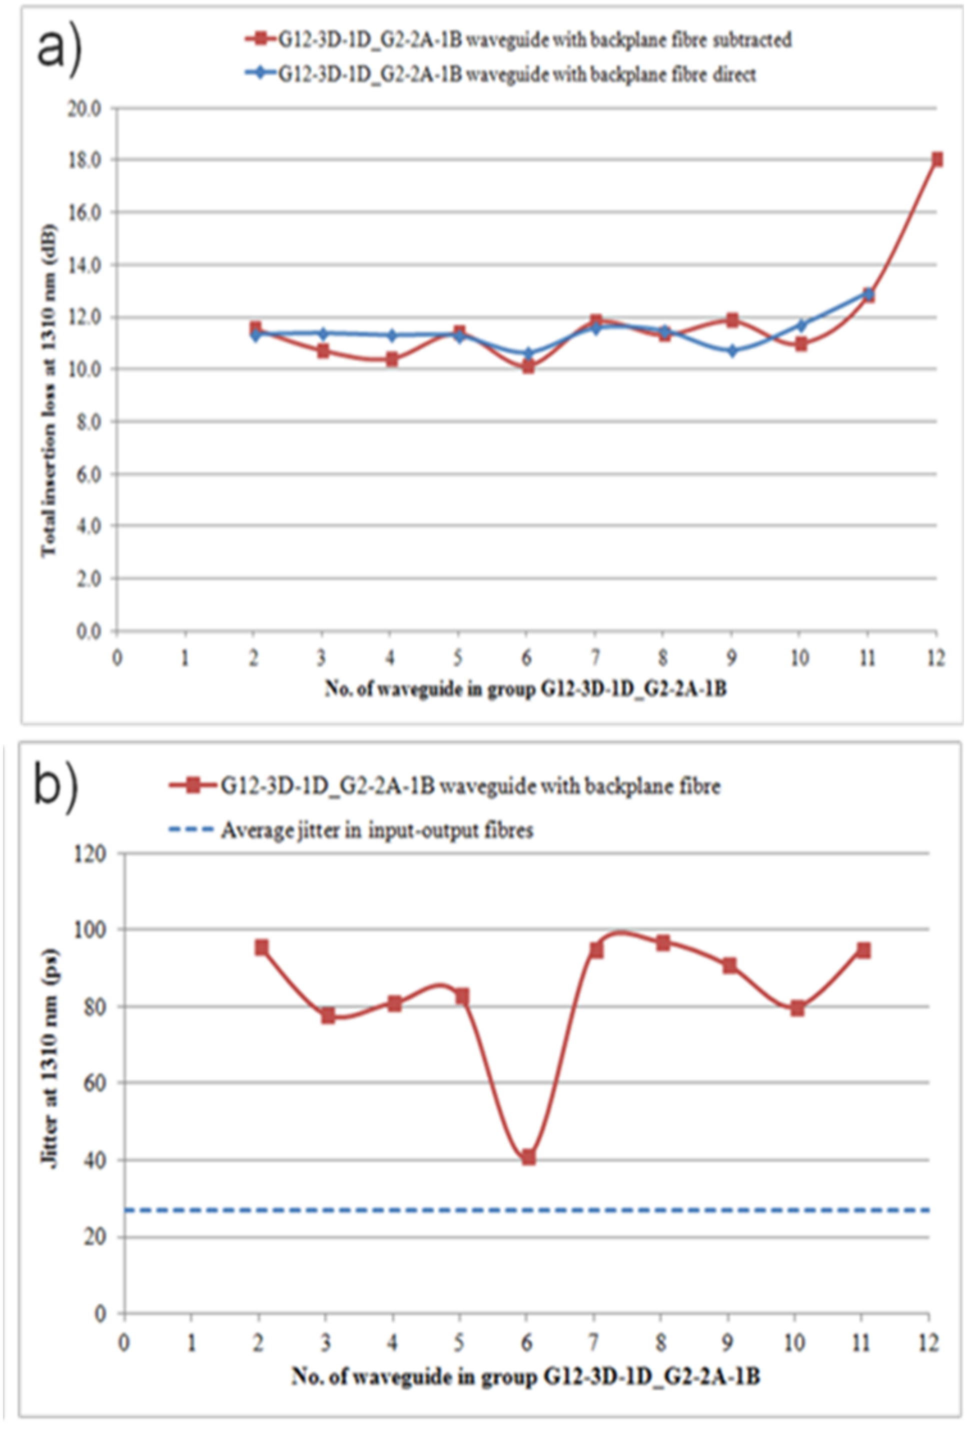


**Figure 5-26: a) Bidirectional insertion loss and b) jitter measurements at 1310 nm as measured through daisy-chained link over both demonstration platforms, through SEPPLANE1b G2 in propagation direction 2A -1B and through SEPPLANE2a G12 in propagation direction 3D-1D**

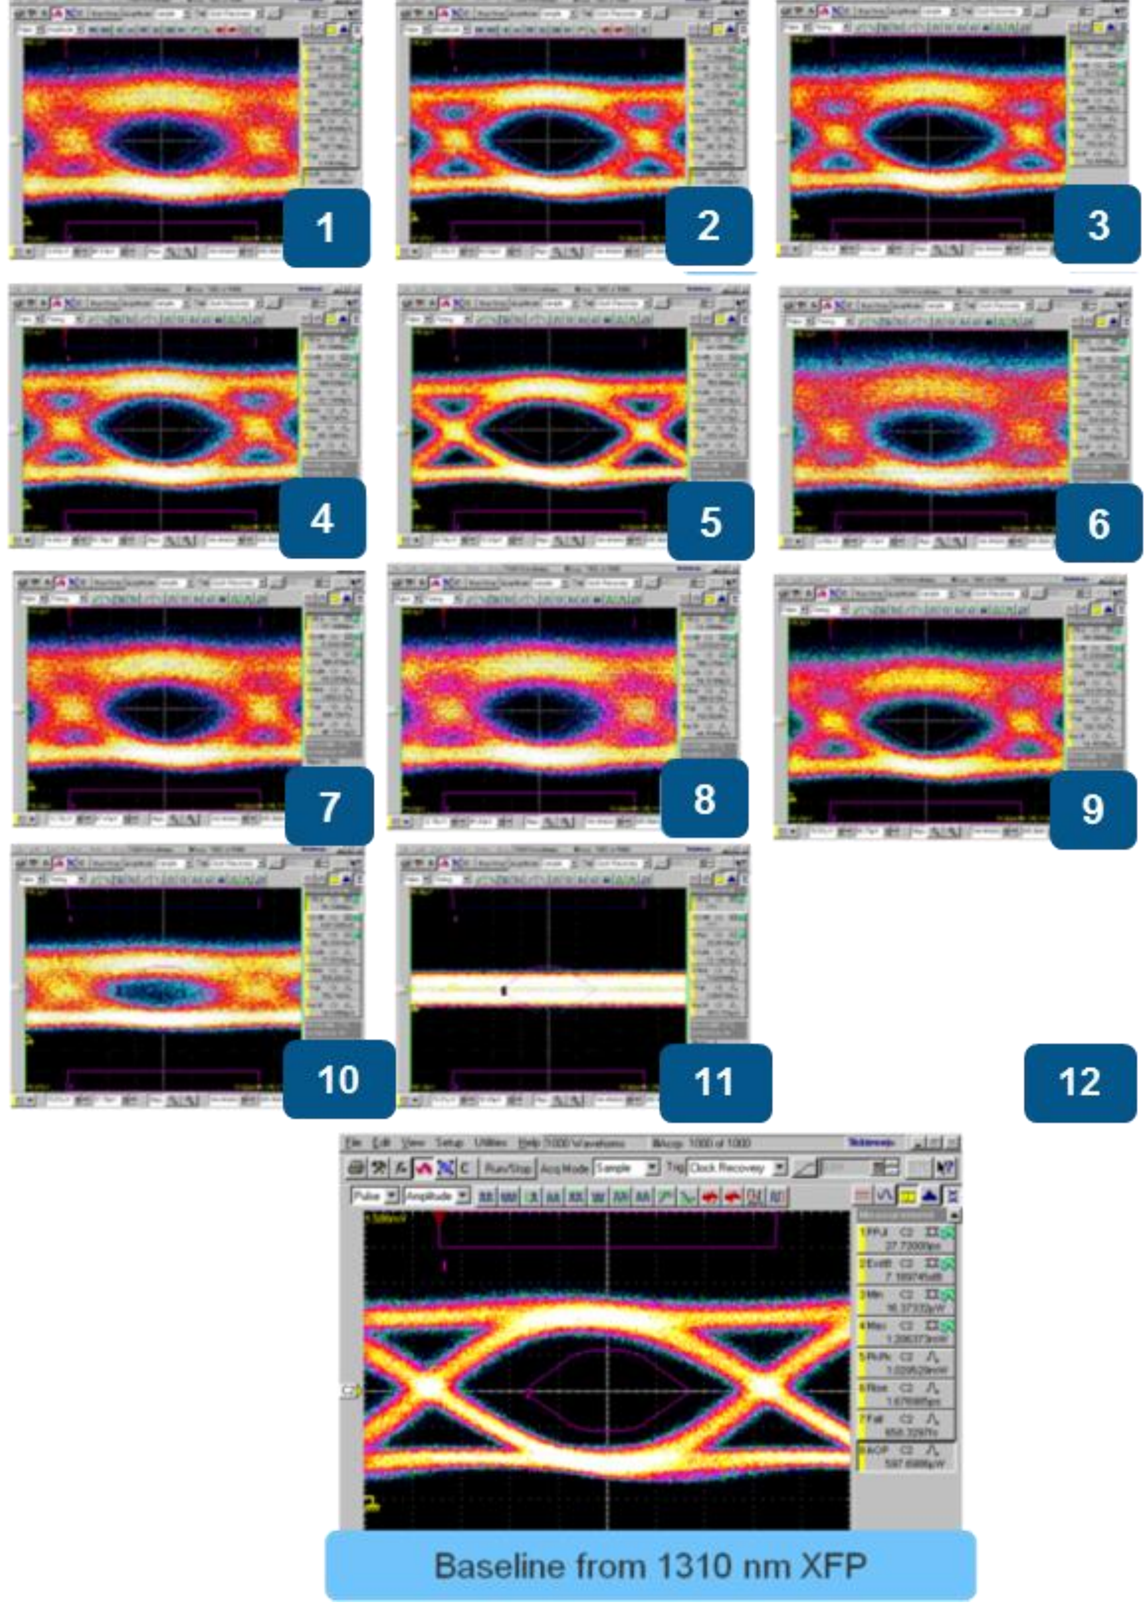


**Figure 5-27: Eye diagrams measured through SEPPLANE1b G2 in propagation direction 2A-1B and through SEPPLANE2a G12 in propagation direction 3D-1D**

### 5.6.3 Correlation between contamination and non-reciprocal insertion loss profiles

In a number of cases, non-reciprocal insertion loss results were shown to coincide with occasional glue contamination on the waveguide interfaces. This contamination occurred when, in some cases, excessive glue was applied to the ferrule receptacles as part of the assembly process, and some glue migrated over the waveguide interface. The ferrule receptacle assembly process has since been improved to eliminate this type of problem. Further characterisation was undertaken to evaluate the correlation between non-reciprocal and asymmetric insertion loss profiles on given waveguide groups and the distribution of glue contamination over the interfaces bounding those groups.

Figure 5-28 shows a comparison of two waveguide interfaces bounding waveguide group G5 on SEPPLANE1b before and after the removal of glue contamination from the interfaces.

As measurements were taken in both directions for each waveguide group, the non-reciprocality of waveguide performance could be assessed depending on which direction the optical signals were conveyed through the link. A strong correlation is shown between the insertion loss profiles for optical signals conveyed through a waveguide in a given direction and the pattern of glue contamination on the waveguide interface, into which the light is launched. In this case insertion loss profiles prior to cleaning were highly non-reciprocal and asymmetric, matching the asymmetric distribution of glue over the corresponding launch interfaces. After cleaning, however, the insertion loss profiles became reciprocal and a substantial overall reduction in insertion loss and variability over the waveguide group is observed.

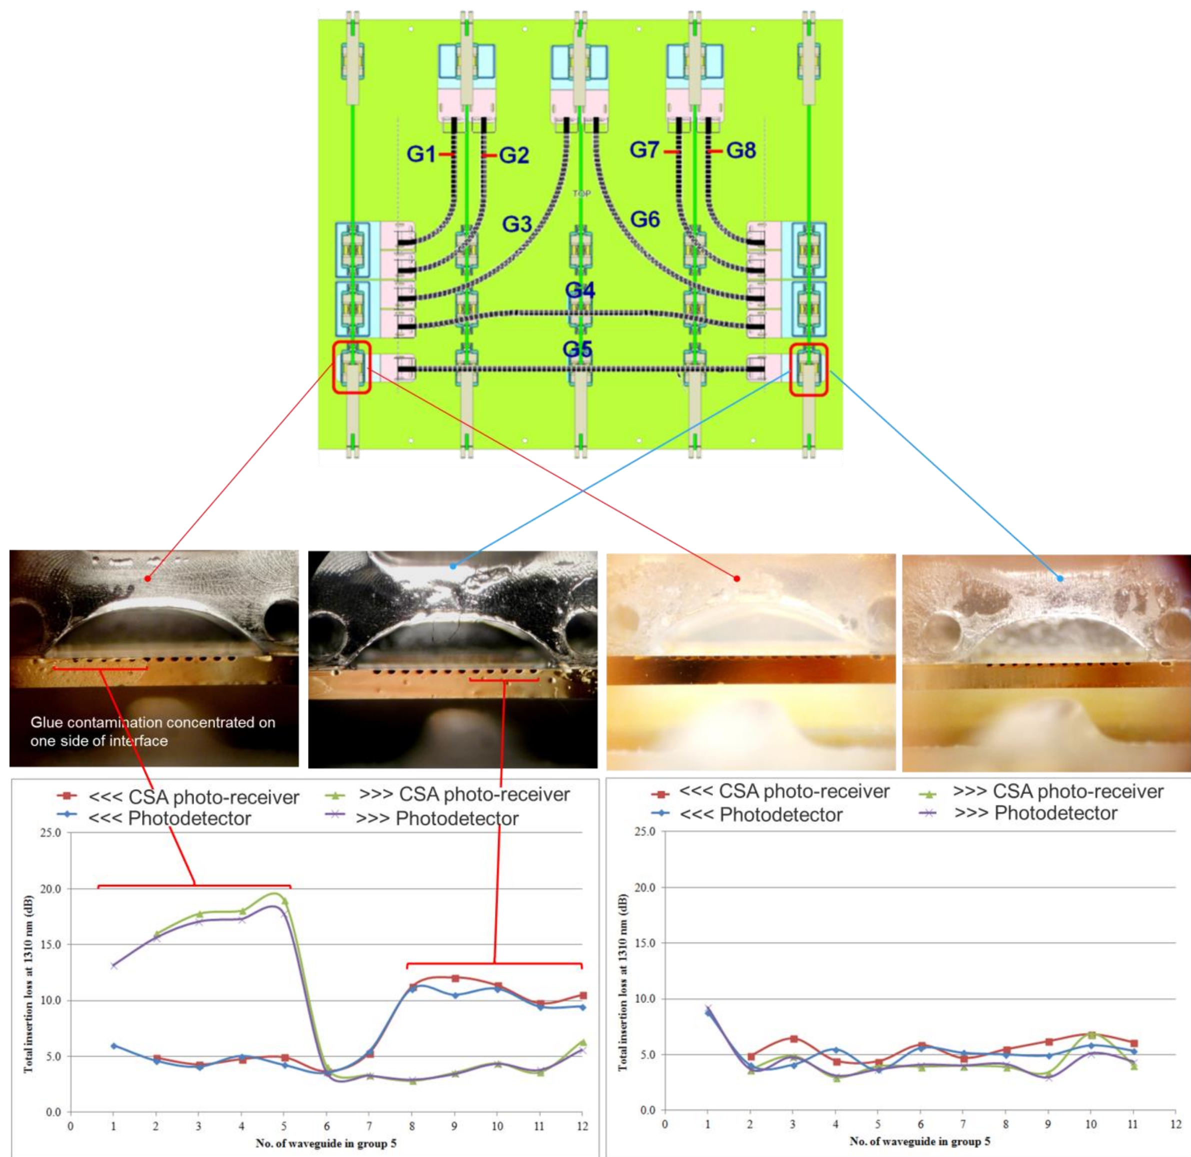


**Figure 5-28: Correlation between glue contamination pattern on waveguide interfaces and bidirectional insertion loss profile before and after removal of glue contamination**

### 5.6.4 Bidirectional signal integrity measurements at 10.3 Gb/s

A 10.3125 Gb/s PRBS $2^{31}$-1 data pattern was generated by an Anritsu MP1800A pattern generator and BERT system and bidirectional signal integrity measurements were carried out on all functional waveguides at 850 nm and 1310 nm.

Figure 5-29 shows the eye diagrams measured at 1310 nm for all 12 waveguides in one of the best performing groups, group 2 of SEPPLANE1b

Table 5-3 shows a summary of total jitter for the full optical connector and waveguide link on all waveguide groups on the 3 backplanes under test as measured on the CSA.

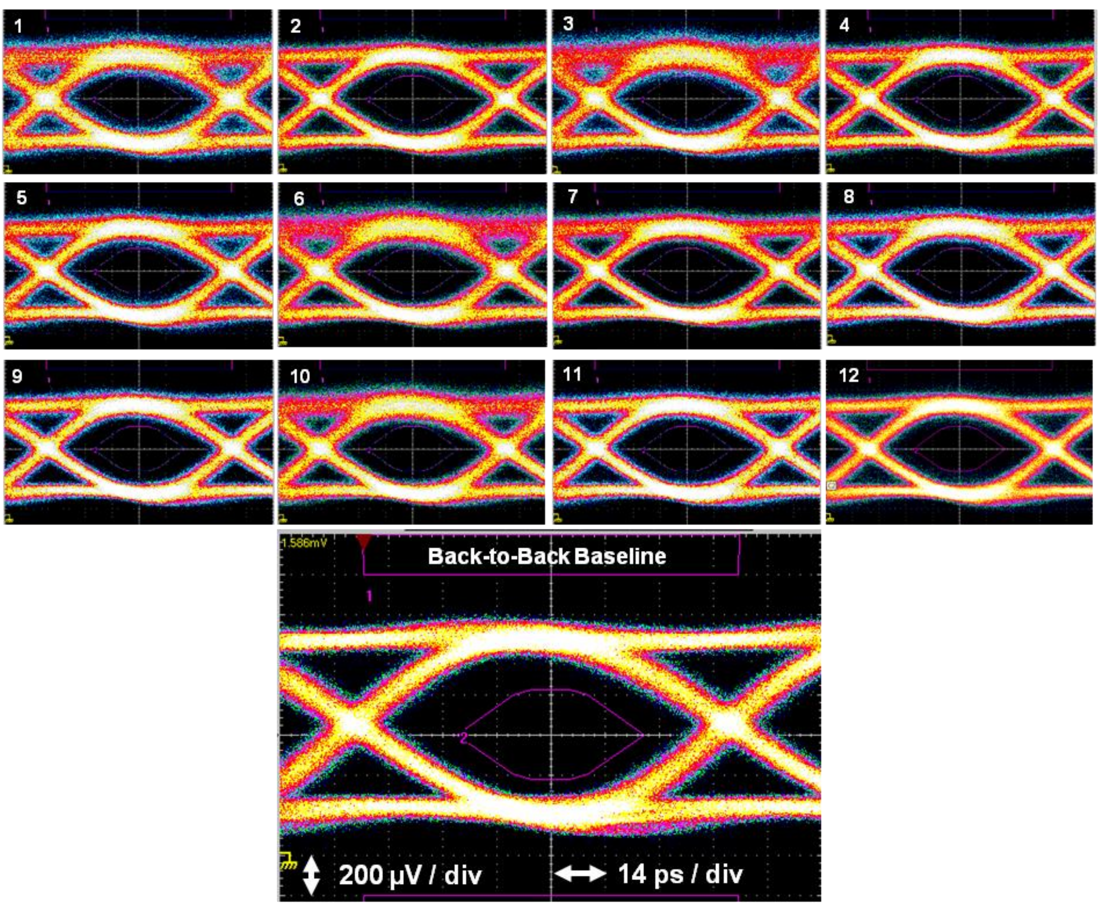


**Figure 5-29: Eye diagrams for 10.3 Gb/s with PRBS $2^{31}$-1 test signal conveyed at 1310 nm over all 12 waveguides of group 2 of SEPPLANE1b. A B2B eye diagram is included at the bottom**

**Table 5-3: Summary of peak to peak jitter as measured on all backplanes with 1310 nm and 850 nm optical test signals modulated at 10.3 Gb/s. The standard deviation of each measurement set is also shown to provide an indication of the measurement spread and variability of results due to contamination**

| Backplane | Average peak to peak jitter / ps | Standard deviation |
|---|---|---|
| **1310 nm** | | |
| Back-to-back reference measurement | 25 | N/A |
| SEPPLANE1a (prior to interfaces being cleaned) | 47.8 | 23.2 |
| SEPPLANE1a (after interfaces have been cleaned) | 35.9 | 21.1 |
| SEPPLANE1b | 27.88 | 8.3 |
| SEPPLANE2a | 38.8 | 10.28 |
| **850 nm** | | |
| SEPPLANE1a (prior to interfaces being cleaned) | 96.1 | 9.8 |
| SEPPLANE1a (after interfaces have been cleaned) | 95.4 | 2.36 |
| SEPPLANE1b | 84.3 | 21.5 |
| SEPPLANE2a | 91.2 | 20.1 |

Table 5-2 shows that SEPPLANE1a (after interfaces have been cleaned) exhibits the lowest insertion loss at 1310 nm, while SEPPLANE2a exhibits the lowest insertion loss at 850 nm. Table 5-3 shows that SEPPLANE1b incurs the lowest jitter values at both 1310 nm and 850 nm.

The variation in performance between the different SEPPLANE backplanes was most likely due to both the incremental differences in fabrication processes employed and the variability in the quality of assembly of the connector receptacles.

### 5.6.5 Bit Error Rate characterisation of on-board 850 nm transceiver subassemblies

Two high power 850 nm optical transceiver subassemblies were mounted on two test daughtercards connected to each other across a SEPPLANE1 backplane. An electrical 10.3 Gb/s test pattern was generated by the Anritsu BERT to drive the on-board transceiver subassembly on the launch daughtercard, which conveyed the corresponding modulated 850 nm optical signal through the connector and waveguide link under test to the

transceiver subassembly on the receiving daughtercard. The received signal was then extracted electronically and conveyed back to the BERT for analysis.

Bit error free operation of BER $<10^{-13}$ for a 10.3 Gb/s test data stream was measured consistently over multiple 2 hour long test periods.

### 5.6.6 Bit Error Rate characterisation up to 32 Gb/s

An Anritsu pattern generator and Photline ModBox-850nm-28Gbps-NRZ optical modulation unit were used to produce 850 nm PRBS $2^{31}$-1 optical test signals at data rates of 15 Gb/s, 20 Gb/s, 28 Gb/s and 32 Gb/s. The optical test signals were conveyed in a loop back configuration through selected connector and Group 4 waveguide links on a SEPPLANE1 type backplane and back to the Photline ModBox receiver and Anritsu BERT. The waveguides alone had a length of 134.65 mm. A bit error rate better than $10^{-12}$ was measured in each case. The results are summarised in Figure 5-30.

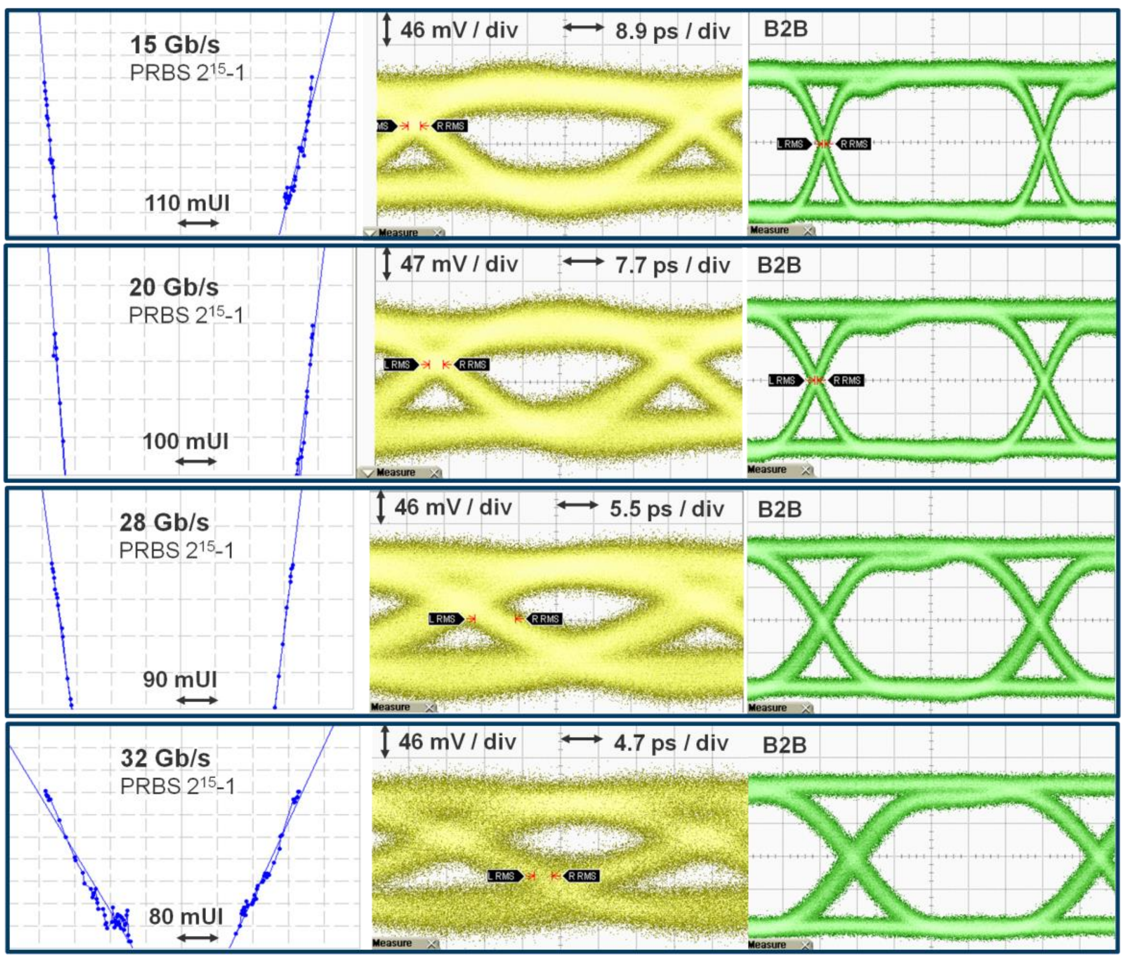


**Figure 5-30: BER bathtub curves and eye diagrams for 850 nm optical test signals conveyed through full optical connector and waveguide link and corresponding B2B profiles. The increased degradation observed on the 32 Gb/s profile is partly due to a reduced power budget and insufficient light being received to resolve a clean eye**

## *5.7 Summary*

In this chapter, the design and the development of a complete technology eco-system around embedded planar glass waveguide based OPCBs was detailed. The work was carried out in collaboration with European partners as part of the SEPIANet project. A suite of passive optical connectors was designed and developed to allow parallel optical fibre cables to connect directly to OPCB integrated glass waveguides, and a corresponding passive dense board-to-board connection system to allow daughtercards to plug into an OPCB backplane. These efforts culminated in the first successful demonstration of a fully integrated planar glass waveguide OPCB backplane and pluggable connector platform driven by system embedded and external 850 nm and 1310 nm optical transceiver technologies and validated for both fibre-to-board and board-to-board optical connectivity at data rates of up to 32 Gb/s.

# 6 INTELLECTUAL PROPERTY

This chapter provides a summary of the intellectual property portfolio developed by the author during the course of this programme in the field of optics and photonics. The portfolio consists of 19 patents and patent applications spanning optical printed circuit boards, optical waveguides, optical connectors, optical assembly and system embedded optical interconnects.

The purpose of this chapter is to provide a Value Proposition for this intellectual property portfolio and to underline and reinforce the competitiveness of the inventions with respect to the current state of the art in embedded optical interconnect technologies.

## *6.1 Introduction*

The premise of this thesis is to explore how the projected performance bottleneck in data centre systems can be mitigated by incorporating electro-optical printed circuit board (OPCB) and optical interconnect technology solutions onto the backplane, midplane [71]–[76] and the peripheral line cards (controllers, servers or disk drives).

A number of key technology inhibitors to the commercial proliferation of embedded optical interconnect were however identified including:

- Lack of viable connector solutions and methods of optically connecting line-cards to optical PCBs

- Need for low cost repeatable high-precision assembly of optical components
- Need for low-cost methods of waveguide and optical PCB fabrication
- Need for tighter in-plane routing of waveguides i.e. smaller bend radii to overcome routing constraints
- Modal dispersion in planar optical waveguides currently too high to accommodate data rates >25 Gb/s over more than half a metre
- Requirement for solutions to increase optical channel bandwidth density, optical link budget and maximum link lengths
- Optical losses too high for longer wavelengths and signal propagation over longer waveguide distances

During course of this thesis the author has developed technology solutions, know-how and an intellectual property portfolio to help resolve these challenges.

## 6.2 *Optical intellectual property summary*

Table 6-1 contains all granted and pending patent applications that were developed and filed between 2005 and 2010. In summary there are 16 patent families with 19 separate patents filed of which 18 have been granted and 1 is still pending.

**Table 6-1: Summary of filed patents**

| Patent family | Patent number | Description | Official Title | Inventors | Priority Date | Countries | Patent / App No | Filing / Grant Date | Status |
|---|---|---|---|---|---|---|---|---|---|
| 1 | P1 | Optical transceiver connector module (Active pluggable optical backplane connector) | Optical connector, a communication system and a method of connecting a user circuit to an optical transceiver [170] | Richard Pitwon<br>Kenneth Hopkins | 01/06/2005 | US final | 7625134 | 01/12/2009 | **Issued** |
| | P2 | | | | | UK | GB2426831 | 25/04/2007 | **Issued** |
| 2 | P3 | High-precision component alignment to PCB (Method of high precision alignment to optical waveguides | Optical printed circuit board and manufacturing method [171] | Ian Johnson<br>Richard Pitwon<br>David Selviah<br>Ioannis Papakonstantinou | 15/07/2005 | US final | 7936953 | 03/05/2011 | **Issued** |
| 3 | P4 | Daughtercard interconnects | Optical circuit board, an optical backplane and an optical communication system [172] | Richard Pitwon | 31/03/2006 | US final | 7454097 | 18/11/2008 | **Issued** |
| 4 | P5 | Optical printed circuit board blank | Optical printed circuit board blank, a kit and a method of making an optical printed circuit board [173] | Richard Pitwon | 25/05/2006 | US final | 7422374 | 08/09/2008 | **Issued** |
| | P6 | | | | | Europe | 1860474 | 15/10/2014 | **Issued** |

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 5 | P7 | Crosstalk suppression (Method of optical crosstalk suppression in high density polymer waveguide optical PCBs) | Optical printed circuit board, a method of making an optical printed circuit board and an optical waveguide [174] | Richard Pitwon | 16/08/2006 | US final | 8731343 | 20/05/2014 | **Issued** |
| 6 | P8 | Multimode CWDM Multiplexer demultiplexer structures on polymer optical PCB and method of manufacturing same | Optical wavelength division multiplexed multiplexer/demultiplexer for an optical printed circuit board and a method of manufacturing the same [175] | Richard Pitwon<br>David Selviah<br>Ioannis Papakonstantinou | 27/09/2006 | US final | 7805033 | 28/09/2010 | **Issued** |
| 7 | P9 | Method of manufacturing multimode CWDM multiplexer / demultiplexer structures on polymer optical PCB and method of manufacturing same | Optical wavelength division multiplexed multiplexer/demultiplexer for an optical printed circuit board and a method of manufacturing the same [176] | Richard Pitwon<br>David Selviah<br>Ioannis Papakonstantinou | 27/09/2006 | US CIP | 8007965 | 30/08/2011 | **Issued** |
| 8 | P10 | Optical adapter module | Adapter for an optical printed circuit board, an optical printed circuit board and a method of connecting an adapter to an optical printed circuit board [177] | Richard Pitwon | 16/02/2007 | US final | 7490993 | 17/02/2009 | **Issued** |
| 9 | P11 | Fabricating a hybrid electro-optical printed circuit board with optical surface layers | Electro-optical printed circuit board, a blank and a method of making an electro-optical printed circuit board [178] | Richard Pitwon | 27/07/2007 | US final | 7899278 | 01/03/2011 | **Issued** |

| 10 | P12 | Lossless tapered waveguides | Method of making a waveguide and a waveguide made thereby [179] | Richard Pitwon | 13/02/2008 | US final | 9333717 | 10/05/2016 | **Issued** |
|---|---|---|---|---|---|---|---|---|---|
| 11 | P13 | Directly pluggable orthogonal in-plane optical PCB connector | An optical connector and a method of connecting a user circuit to an optical printed circuit board [180] | Richard Pitwon | 12/03/2010 | US final | 12/722908 | 12/03/2010 | Pending |
| 12 | P14 | Optical capstan to simplify routing of optical waveguides | An optical connector and a method of connecting an optical connector to an optical printed circuit board [181] | Richard Pitwon<br>Kenneth Hopkins<br>David Milward | 12/03/2010 | US final | 8306374 | 06/11/2012 | **Issued** |
| 13 | P15 | Optical PCB interconnect for storage devices | An interconnect for a data storage system [182] | Richard Pitwon<br>Kenneth Hopkins | 12/03/2010 | US final | 8861975 | 14/10/2014 | **Issued** |
| 14 | P16 | Modular interconnectable electro-optical backplane | A data storage system, a modular printed circuit board, a backplane and a backplane component [183] | Kenneth Hopkins<br>Richard Pitwon | 24/05/2010 | US final | 8417071 | 09/04/2013 | **Issued** |
| 15 | P17 | Optical amplification devices for polymer optical printed circuit boards | An amplification module for an optical printed circuit board and an optical printed circuit board [184] | Richard Pitwon | 01/09/2010 | US final | 9325146 | 26/04/2016 | **Issued** |
| | P18 | | | | | US CIP | 8891932 | 18/11/2014 | **Issued** |
| 16 | P19 | Design and manufacture of polymer optical waveguide amplification structures | An optical PCB and a method of making an optical PCB [185] | Richard Pitwon | 01/09/2010 | US final | 8488920 | 16/07/2013 | **Issued** |

## 6.3 *Areas of invention*

During the course of my optical research and development activities, it became apparent that there were a number of key challenges to the viability of system embedded optical interconnect, specifically optical backplanes / midplanes and the myriad technology solutions required to enable and support them. It was important for instance that daughterboards be able to plug optically to a backplane with embedded optical channels, however due to the inherently small size of optical waveguides, the task of aligning and assembling connector components to them with the required precision could not be overcome with conventional techniques. Patents 1 and 2 addressed this and have been successfully implemented and demonstrated as described in Chapter 3. Different variants of such optical PCB connector solutions, including right-angled, in-line and rotatable connectors, were developed to further enable and support future target applications in data centre environments (Patent families 3, 11, 12).

Polymer multimode waveguides, which underlie many of the target applications under consideration are fraught with limitations to their performance including high propagation and bend losses and modal dispersion. The author has devised innovative new waveguide concepts, structures and fabrication methods to overcome these intrinsic obstacles to their commercial deployment (Patent families 5, 14 and 15).

Electro-optical printed circuit boards (OPCBs) containing optical layers as well as electronic signal layers – are central to this research effort as they include optical backplanes / midplanes / motherboards and ultimately peripheral daughterboards. Innovations devised in this field include a method of creating reconfigurable optical PCBs with generic “optical sockets” (Patent family 3) and various methods of fabricating electro-optical printed circuit boards (Patents 4 and 9). Intellectual Property has also been generated to address how OPCBs can be scaled up to rack level data centre environments including how to connect many optical PCBs together to form “super-backplanes” (Patent family 13) and how to amplify optical signals propagating along a PCB embedded waveguide to increase the optical link budget and therefore maximum waveguide length (Patent families 15 and 16).

In addition the author devised more forward looking enhancements to further increase waveguide signal bandwidth including how to multiplex / demultiplex light of different wavelengths into / out of planar waveguides in order to increase channel bandwidth (Patent families 6 and 7).

## 6.4 *Optical intellectual property portfolio*

### 6.4.1 Patent 1: Optical transceiver connector module [170]

#### *6.4.1.1. Formal abstract on published patent*

This invention provides an optical connector for connecting a user circuit to an optical backplane, in use the connector being adapted for mounting on a user circuit. The connector comprises an active or passive photonic interface through which optical signals may be transmitted and received between a user circuit and a said optical backplane; a primary aligner for engagement with a corresponding aligner on a backplane to ensure alignment of the optical interface with the backplane; and a support for supporting the aligner and/or the optical interface on the connector. The support is selected to enable relative movement between a user circuit to which the connector is connected in use and the aligner and/or the optical interface. The support is preferably a flexible printed circuit board.

#### *6.4.1.2. Description*

As the backplane and its peripheral cards in a data communication system (Figure 6-1a) are connected in a mutually orthogonal way, the author invented and developed an in-plane pluggable connector technology and connection scheme [127] whereby the optical interfaces of optical transceiver modules housed on the mating edge of the peripheral cards can be butt-coupled to optical channels embedded on the midplane. This builds on the connection methodology demonstrated in the Storlite project [3]. According to this scheme the optical axis of the peripheral transceiver module is collinear with the embedded optical channels in the optical printed circuit board, thus eliminating the need for right-angled mirrors and minimising the number of boundaries incurring optical loss.

The patented optical transceiver connector module (**Figure 6-1**b) comprises a parallel optical transceiver circuit, a self-aligning optical interface and a connector mechanism. The transceiver circuit is constructed on a flexible material to enable the optical interface to mechanically float with respect to the line card, thus allowing the critical optical connection to remain relatively immune to displacements between line card and backplane. A manual connection mechanism (e.g. cam lever) controls the engagement and disengagement of the transceiver with the optical backplane.

What is claimed is a board to board optical connection system based on the use of optical component and alignment structures assembled on a flexible PCB material, in order to allow the active optical components in the daughterboard connector to be mechanically aligned with an optical interface (such as an optical waveguide arrangement) in a receptacle optical system (such as an optical PCB). This invention thus exploits the mechanical (as well as electrical) properties of flexible PCB to accommodate a free-floating active optical component system within a transceiver connector module, such that the module remains rigid with respect to the host board on which it is supported, whereas the active optical components and alignment features can be free-floating and mechanically translated into and out of registration with the optical interface on the receptacle optical system.

The implication of this invention is to make possible the implementation of optical (transceiver) connection systems between electronic daughter-boards and an optical PCB in a large rack system, without compromising the critical optical alignment requirement inherent to these systems.

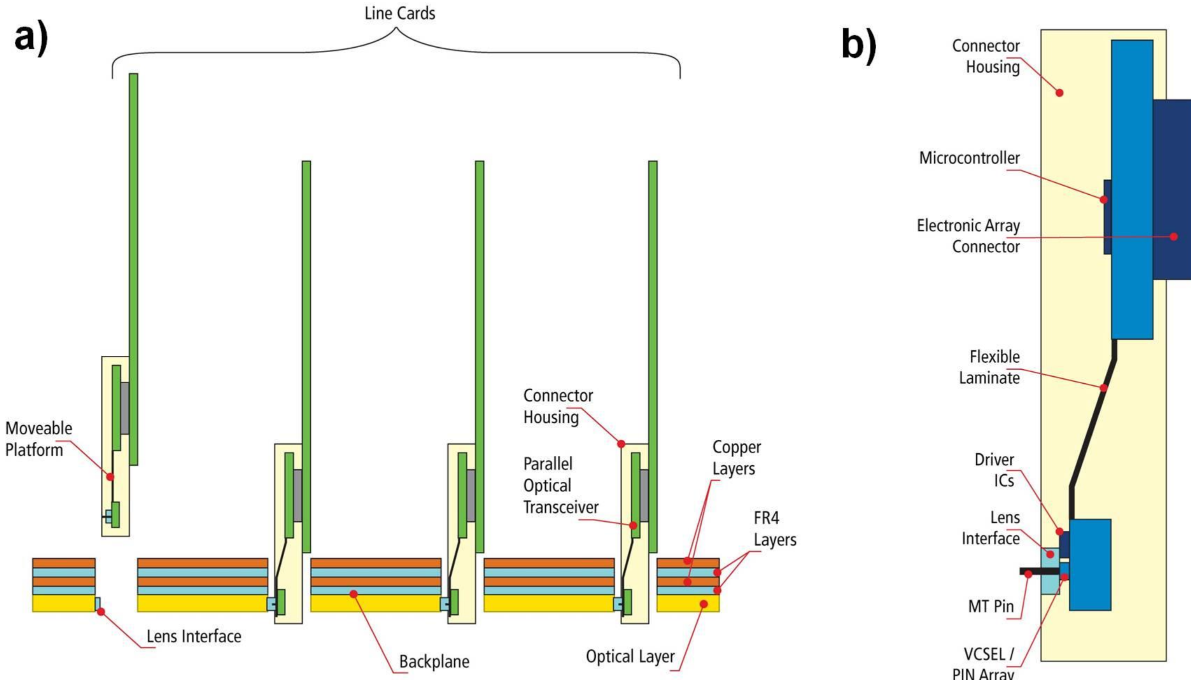


**Figure 6-1: (a) Electro-optical midplane connection scheme; active pluggable connectors housed on the edge of peripheral line cards engage with the embedded optical layer in the backplane PCB.**

**(b) Active optical connector comprising a parallel optical transceiver, connector housing and engagement mechanism.**

### *6.4.1.3. Implementation*

The author successfully developed and demonstrated a prototype active pluggable connector to allow optical connection between the peripheral line cards and the optical layer embedded in an electro-optical midplane. The connector comprised a parallel optical transceiver, connector housing and a pluggable engagement mechanism. Full details on this project are available from the following reference papers [2], [128], [129]. This work forms the basis of this thesis and is described in Chapter 3 and Chapter 4.

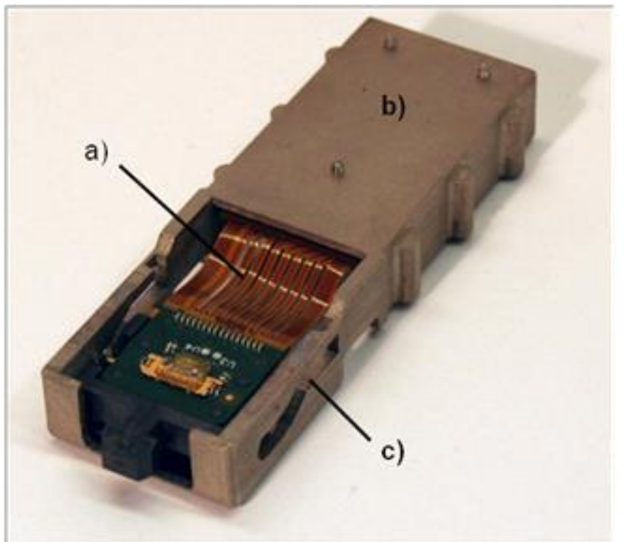


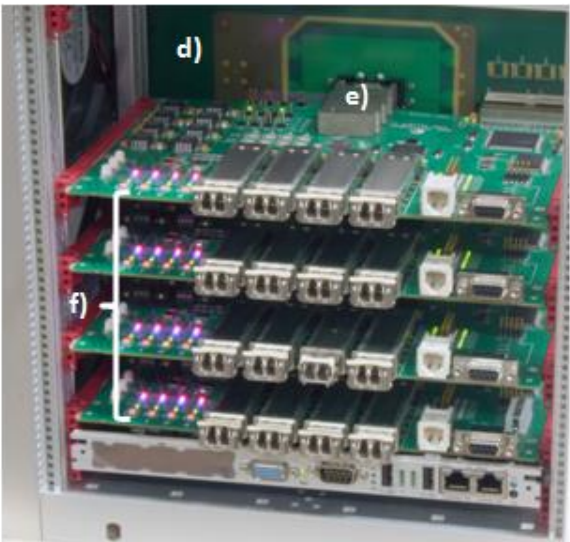


**Figure 6-2: a) Optical transceiver circuit mounted on flexi-rigid substrate, b) Connector module housing, c) Grooves to enable required movement of optical interface during mating process, d) Electro-optical midplane in FirstLight demonstration platform fully populated with all 4 test line cards and powered up, e) Connector module housed on line card and plugged into electro-optical backplane, f) 4 test line cards**

### 6.4.2 Patent 2: High-precision passive component alignment to PCB [171]

One crucial requirement for the commercial deployment of optical PCB (OPCB) technology is a low-cost technique for the high-yield assembly of optical interface components onto the optical layers. In order to enable high volume assembly it is preferable that such techniques be passive and repeatable.

#### *6.4.2.1. Formal abstract on published patent*

The invention provides a method of manufacturing an optical printed circuit board and an optical printed circuit board. The method comprises providing a support layer; on the support layer, providing an optical core layer; forming optical channels from the optical core layer and surrounding the optical channels with cladding thereby forming optical waveguides; and during said step of forming the optical channels, forming one or more alignment features, e.g. projections, on the optical printed circuit board.

#### *6.4.2.2. Description*

This invention concerns a method of self-alignment of components onto printed circuit boards (PCBs), which is very accurate with respect to other components on the PCB to very high precision. This is particularly relevant when dealing with on-board optical components or optical waveguides (such as one would find in an optical PCB) in which precise inter-component alignment is critical.

An optical PCB has optical waveguides to transmit light signals between components, as well as or instead of conventional copper conductors. It may usually consist of a base layer of, for example, glass fibre epoxy laminate material, such as FR4, as in conventional PCBs. In the area where optical waveguides are needed a lower optical cladding layer is first laid down. On top of this, a layer of optical core material is laid down. This has a higher refractive index than the cladding layer and will ultimately form the actual optical waveguides. By a conventional process of masking and etching most of the core layer is removed, leaving only salient strips of core material where the optical waveguides are needed. Finally, an upper cladding layer is laid down, so that the strips of core material are completely surrounded by cladding material, and hence are able to function as optical waveguides. These layers are illustrated in Figure 6-3.

Conventional multimode optical waveguide core layers are typically 50 – 70 μm thick; the waveguides are typically 50 μm across and can be fabricated to very high accuracy, currently of the order of 5 μm. High accuracy is obviously a key requirement of any optical waveguide structure.

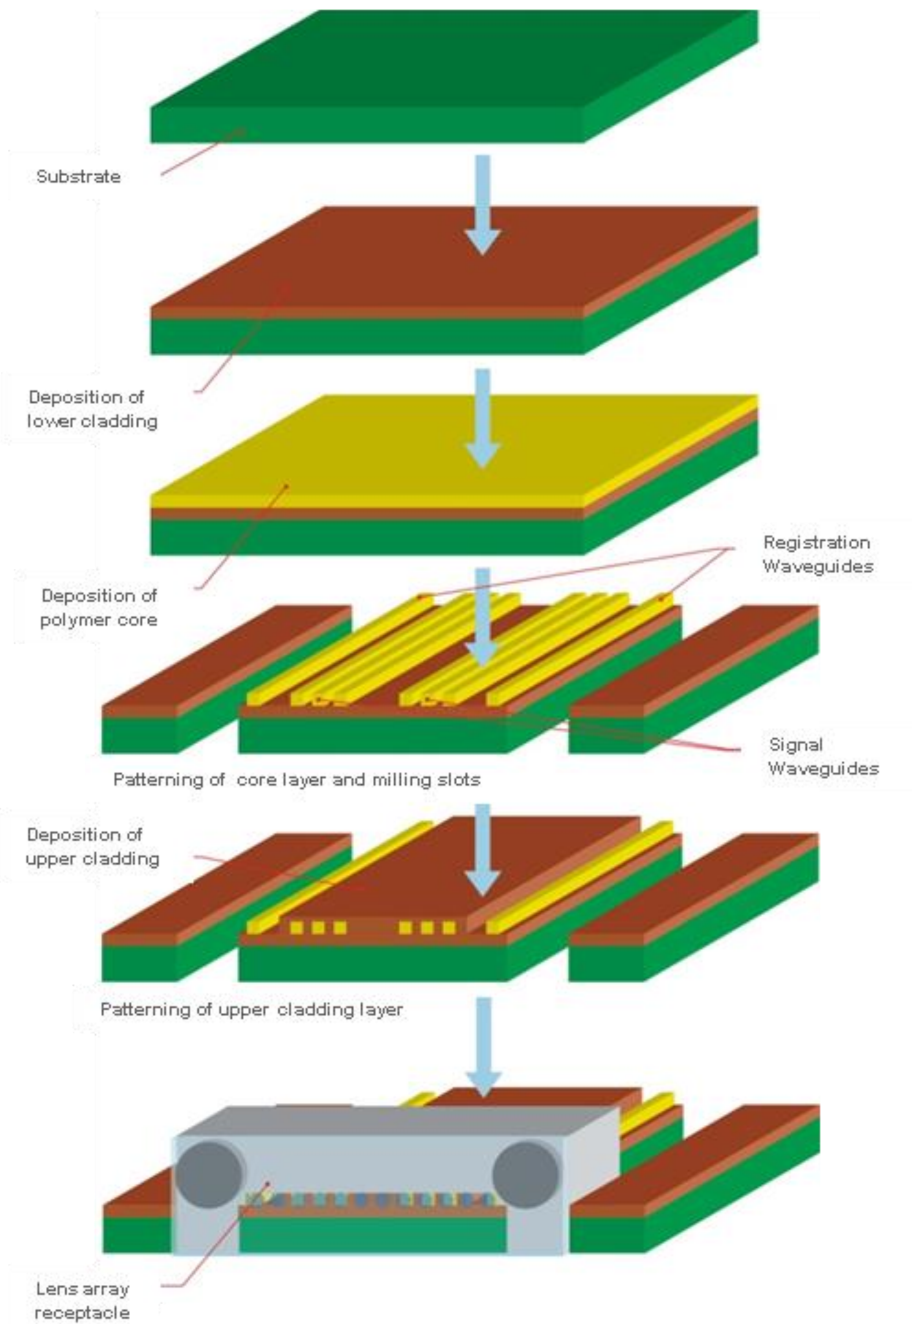


**Figure 6-3: Fabrication process for the passive alignment features on the optical layer**

The key premise of this invention is to construct mechanical registration stubs out of the optical core material, in the core material layer, in the same step by which the optical waveguides themselves are created. In effect, these registration stubs would be enlarged optical waveguides. The difference is that these registration stubs are not completely surrounded by cladding material; sections of core material remain exposed, particularly on their outer sides. Figure 6-4 shows that the sides of the registration stubs generally will be inclined (anisotropic) rather than

strictly orthogonal; in some situations this can actually facilitate the positioning of the component that aligns with the registration stubs.

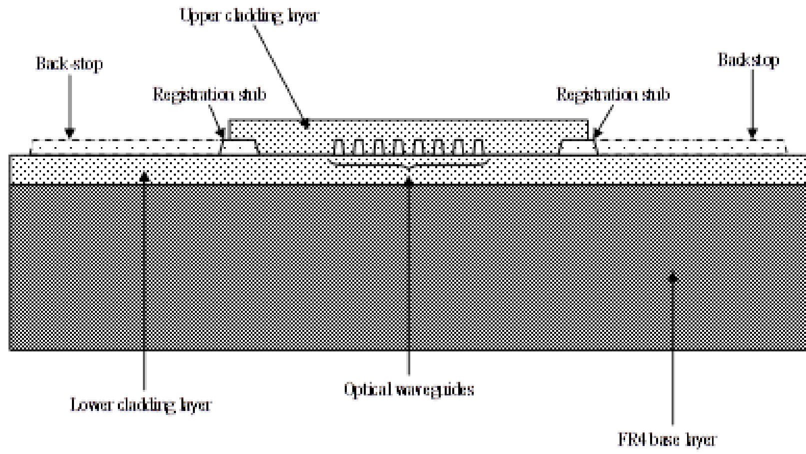


**Figure 6-4: Diagram from Pitwon precision alignment patent** [144]

Because the registration stubs are constructed in the same step as the optical waveguides, they are inherently aligned as accurately as possible relative to the waveguides. Therefore the exposed core material can be used as a highly accurate registration feature for the placement of components that will interface to the optical waveguides. The registration features can thus be positioned to an accuracy of around 5 µm, well within the currently required accuracy of 15 µm. They can be shaped so as to provide accurate registration in three dimensions:

- Registration stubs including back-stops account for lateral positional precision
- Top surface of the lower cladding layer accounts for vertical positional precision

A component such as a connector receptacle, 45° mirror or optical waveguide holder may be assembled on the PCB using the registration stubs for highly accurate positioning with respect to other components. The component can then be fixed in place in a conventional manner such as by means of the application of adhesive such as epoxy.

A proprietary fabrication technique and method of passively aligning and assembling parallel optical microlenses to embedded polymer waveguide arrays was successfully developed [24,30,31]. These form a critical part of the

pluggable in-plane connection interface between arbitrary external optical devices, either passive or active, and a PCB embedded optical circuit.

The complete fabrication process for the passive alignment features on the optical layer is outlined in [3]. The procedure involves the fabrication of passive mechanical registration features in the core layer during the same process step in which the waveguide cores themselves are patterned. Effectively these serve as additional “dummy” waveguides, which are positioned on either side of the signal waveguides and as a result their positional accuracy with respect to the signal waveguides is as high as those of the signal waveguides to each other. Instead of uniformly curing the upper cladding however, it must be selectively cured to ensure that the central signal waveguides are completely clad while the registration waveguides are not.

This clearance allows for direct mechanical registration of arbitrary components to the waveguides. In addition, the fabrication tolerances required to pattern the upper cladding for this purpose are far lower than those required to pattern the waveguides themselves. It is only important that the outer edges of the registration waveguides, which form the mechanical datum, be mechanically exposed. Preferably, the upper cladding should partially cover the registration waveguides in order to provide structural reinforcement and reduce the risk of the registration waveguides delaminating under the strain. However, this is not strictly necessary and as shown in **Figure 6-5** the registration waveguides in the FirstLight electro-optical midplane were left completely uncovered without any adverse effects. This technique can be implemented using most waveguide fabrication processes.

The positional tolerance of the mechanical registration features with respect to each other has been measured to be ±3 µm for lateral misalignment in-plane and ± 4 µm normal to the PCB plane [2], [156].

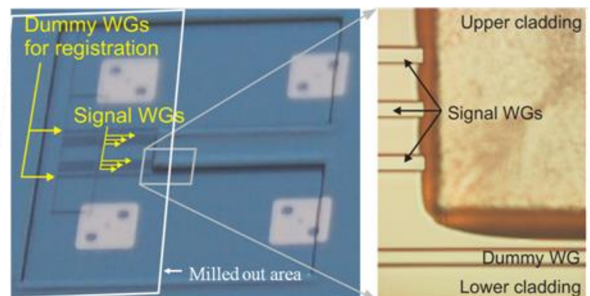


**Figure 6-5: Photographs of upper cladding opening on FirstLight electro-optical midplane showing waveguides for signal transmission as well as connector alignment features (dummy waveguides).**

### *6.4.2.3. Implementation*

As part of the Storlite project in 2005 and the FirstLight project in 2008, this invention was successfully demonstrated through the development and verification of connector receptacles, which could be passively aligned to very high accuracy to polymer waveguides. The custom receptacle included compliant structures to allow it to mechanically engage with the registration waveguides on the board and a recess to accommodate a standard high performance “mechanical transfer” connector (MT) compliant lens array.

Further details are available from following conference and academic papers [2], [3].

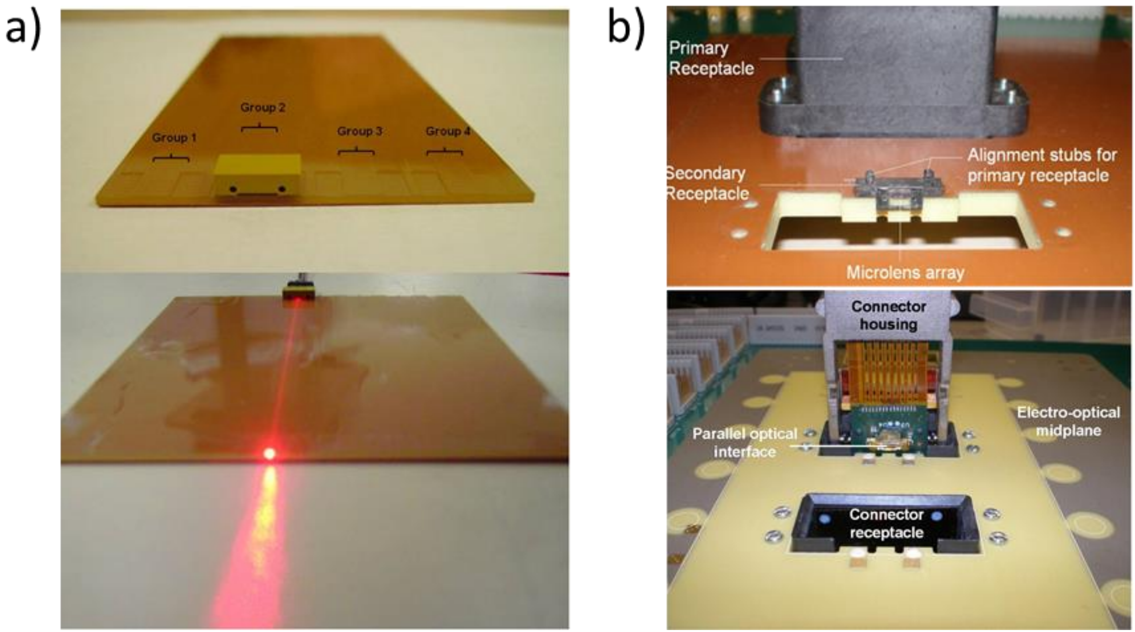


**Figure 6-6: a) Storlite waveguide receptacle, b) FirstLight waveguide receptacles**

### 6.4.3 Patent 3: Daughtercard optical waveguide interconnects [172]

#### *6.4.3.1. Formal abstract on published patent*

The invention provides an optical circuit board, the circuit board comprising: a rigid support layer; a flexible support layer formed on the rigid support layer; and, an optical layer formed on the flexible support layer. The optical layer includes one or more optical waveguides extending from a first area of the optical circuit board to an edge of the circuit board wherein the flexible support layer extends beyond the edges of the rigid support layer thereby defining a flexible passive optical connector for the circuit board.

#### *6.4.3.2. Description*

Most prior art concerning optical backplane systems does not take into account the need for continuous high precision alignment of pluggable optical interfaces on the daughtercard to compliant interfaces on the backplane. Little has been done to tackle the very significant problem of movement between daughtercards and backplane, which will occur in any real system as a result of vibrations, air flow variations and thermal and mechanical PCB deformation.

The proposed solution is based on the need for quasi free-floating optical interfaces on both the daughtercard and the backplane. The proposed method takes into account that electronic traces carrying high speed signals on the daughtercard should not be routed to a dedicated optical transceiver location on the edge of the card at the point of backplane interconnect. This is tolerable at current data-rates (12 Gb/s), but when data-rates increase this will become unmanageable. The method allows for optical transceivers to be placed at arbitrary locations on the daughtercard, close to the high-speed signal source to minimise trace lengths. The transceivers will then perform opto-electronic and electro-optic conversion somewhere on the daughtercard and direct the light into optical waveguides on dedicated optical layers inside the daughtercard. These waveguide layers will be deposited on flexible PCB material (e.g. Kapton polyimide), which itself is embedded within conventional rigid PCB material such as FR4. The waveguides will be routed to the edge of the daughtercard to sections where the flexible layer is exposed and not attached to FR4. This gives them the required mechanical flexibility **Figure 6-7**.

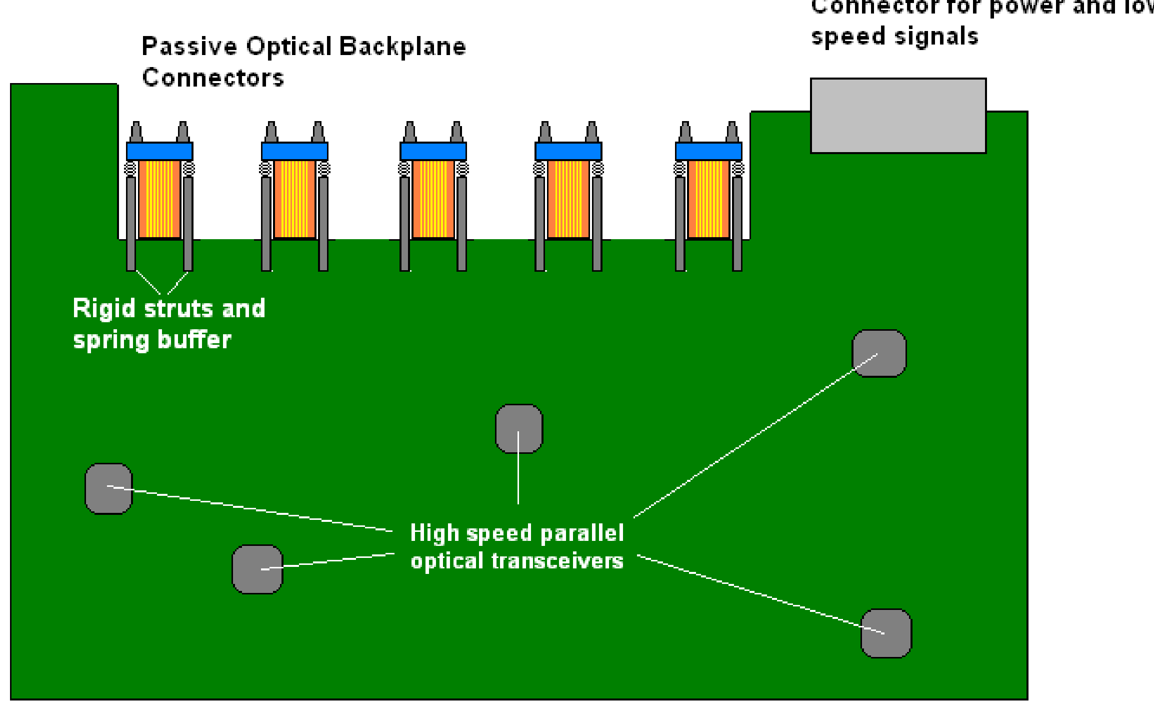


**Figure 6-7: Electro-optic daughtercard with embedded flexible optical layers**

These optical interface sections will be supplemented with passive high precision components which contain mechanical alignment structures and possible optical interface supplements such as micro-lens arrays. These would be high precision devices, but can be mass-produced at low-cost using injection moulds. These devices would be assembled and aligned to the waveguides at very low-cost using Xyratex's patented optical waveguide alignment method (Patent 2), in accordance with which they will contain registration features, which will self-align with the on-board registration features.

This patent specifically protects the concept of having embedded optical waveguides conveyed from an arbitrary midboard location on a PCB to a flexible edge board location.

#### *6.4.3.3. Implementation*

This invention was deployed on a joint demonstration system developed on a prototype data storage enclosure developed by the author at Xyratex in conjunction with Finisar, Vario-optics and Huber+Suhner in 2012 [186].

### 6.4.4 Patent 4: Reconfigurable optical printed circuit board [173]

#### *6.4.4.1. Formal abstract on published patent*

The present invention provides an optical printed circuit board blank, comprising, a support layer; one or more optical waveguides formed thereon; and at least one socket for receiving an optical component, the socket including one or more alignment features to ensure alignment of an input/output interface of the said optical component when arranged in the socket with an input/output interface of at least one waveguide; and a flexible optical connector arranged at an edge of the circuit board, the optical connector being optically connected to the, or each of the, sockets, to enable optical communication between an optical component located within one, or all of the, sockets and the optical connector.

#### *6.4.4.2. Description*

This patent is related to patent 3 in that it protects the concept of deploying active or passive optical devices such as e.g. transceivers or lens arrays at arbitrary locations in the middle of an electro-optical printed circuit board. Through the use of high precision structures of the type outlined in Patent 2, optical "sockets" can be created at midboard locations, which are connected to optical waveguides in the board. These sockets can accept active or passive optical components (**Figure 6-8**), which perform operations on the PCB embedded optical channels. These operations could include splitting, wavelength division multiplexing or optical to electrical signal conversion.

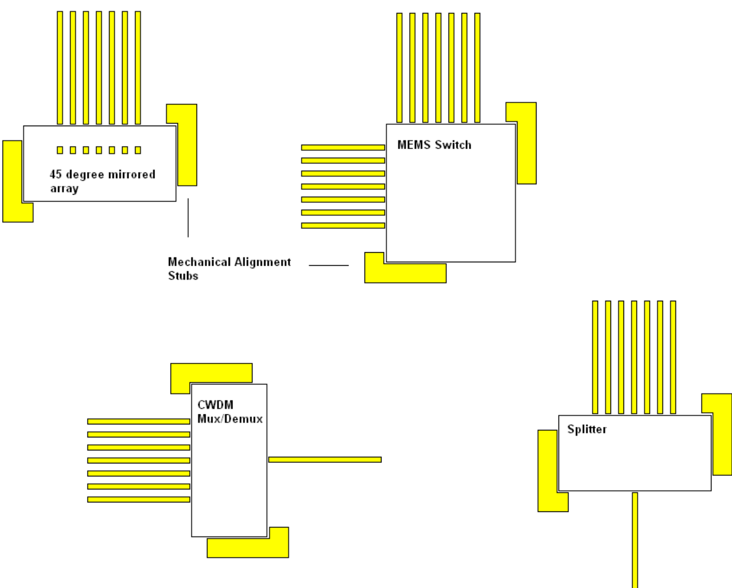


**Figure 6-8: Various passive and active optical components which are passively aligned into optical sockets and their associated channels through mechanical alignment features**

**Figure 6-9** shows an implementation of this with transceiver devices and out-of-plane deflection structures, whereby the region below these devices will be a section where the flexible PCB material and optical polymer is etched out. In this void can be inserted a passive optical device which caters for 45° deflection of the light into or out of the transceiver and waveguide. This passive optical device would be aligned to high precision with the waveguides using the Xyratex patented low-cost waveguide alignment method (Patent 2: High-precision passive component alignment to PCB) and assembled using index matching glue to minimise optical losses.

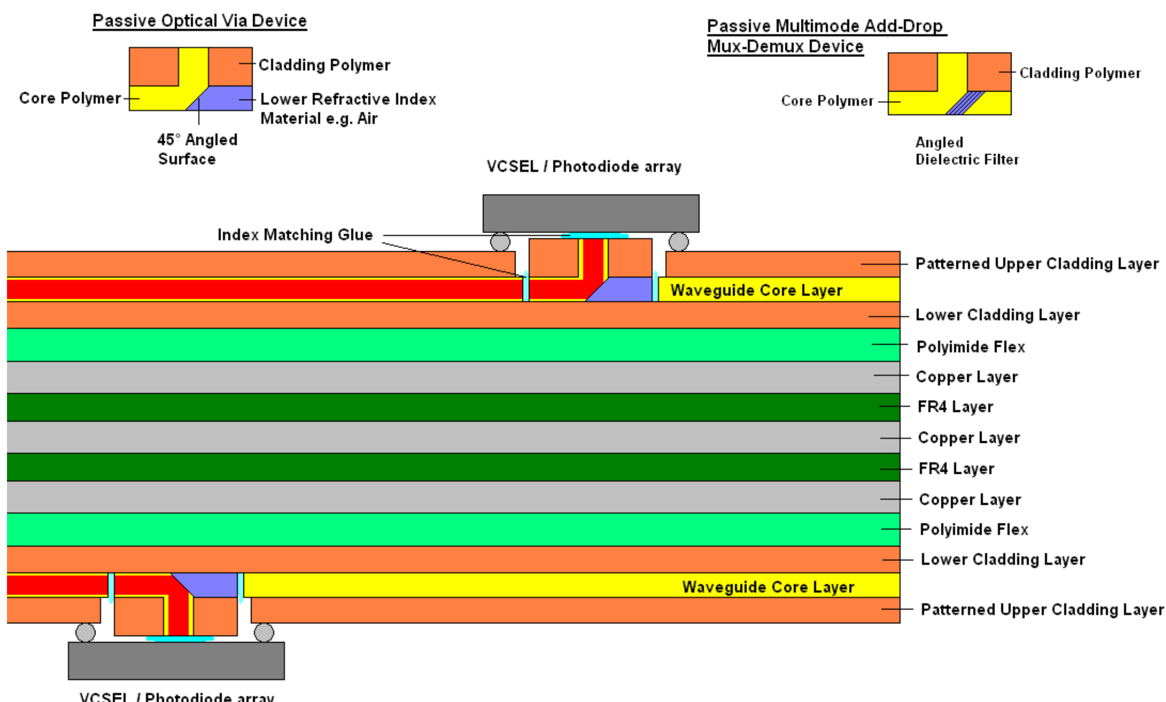


**Figure 6-9: Electro-optical printed circuit board with optical sockets housing optical transceiver devices and out-of-plane deflection structures**

### *6.4.4.3. Implementation*

In the SEPIANet project described in Chapter 5, transceiver modules are located in the middle of the test daughter board close to the high speed signal sources (the electrical RF connectors) to minimise electrical trace lengths. These modules were connected to backplane connector plug with fibre jumpers.

### 6.4.5 Patent 5: Crosstalk suppression on optical waveguides [174]

#### *6.4.5.1. Formal abstract on published patent*

The present invention provides an optical printed circuit board, comprising at least one optical waveguide for carrying optical signals on the optical printed circuit board; and a trench formed adjacent to the at least one optical waveguide, wherein the trench contains a light absorptive material to absorb light that strays from at least one waveguide.

#### *6.4.5.2. Description*

A method is proposed of implementing strong crosstalk suppression between the waveguides of an optical PCB, whereby trenches are deployed between waveguides however boundary reflections suppressed.

At every boundary between two materials of different refractive indices, an optical signal will be partially refracted and partially reflected. Therefore, if an unfilled trench is fabricated between the waveguides, some light will inevitably be reflected back. This has the effect of creating a secondary waveguide, which will give rise to greater optical jitter and noise (**Figure 6-10**).

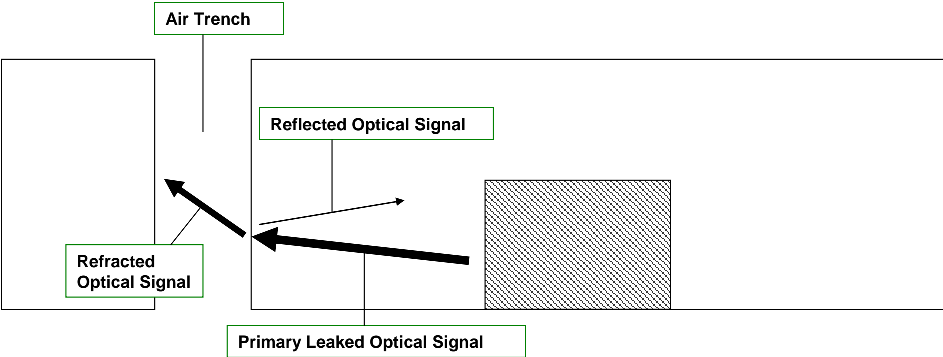


**Figure 6-10: Boundary effect on leaked light with air trench**

If the trench is filled with some material e.g. black ink, then, although any light which has penetrated the boundary will be absorbed, some light will again be reflected back into the waveguide if the black ink has a significantly different refractive index to the cladding.

Only if the trench is filled with a material with the same refractive index as the cladding will there be virtually no reflection. This is due to the fact that the signal ‘sees’ no boundary.

If the fill material in question is in turn doped with light absorbing impurities then the uninterrupted signal will eventually be absorbed.

The patent includes a very wide variety of proposed techniques and features to enable crosstalk suppression between waveguides and fabricate such structures including:

- Sparse particulate concentration in inter-waveguide trench
- Inverted Bell Shape Trench filled with doped cladding with a varied concentration profile to compensate for trench width variation
- Inverted Bell Shape Trench Fabrication by Laser Ablation
- Inverted Bell Shape Trench Fabrication by Laser Writing
- Inverted Bell Shape Trench Fabrication by Photolithography
- Gaussian inter-waveguide particulate concentration (Figure 6-11)
- Method of creating Gaussian inter-waveguide particulate concentration profile through dry deposition of particulates
- Method of creating Gaussian inter-waveguide particulate concentration profile through multi-nozzle polymer jetting

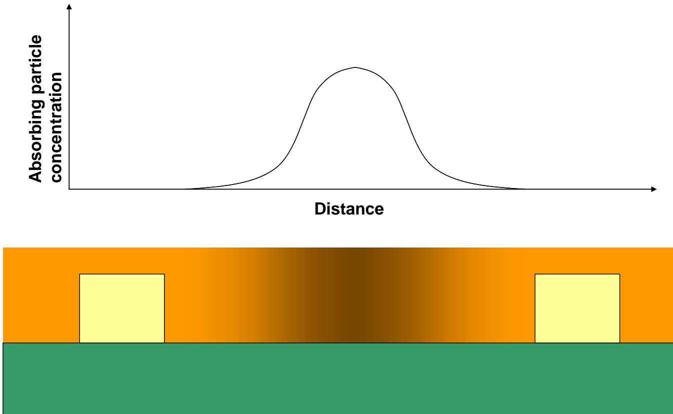


**Figure 6-11: Graded particulate concentration to reduce back-reflection**

## 6.4.6 Patent 6: Passive optical waveguide integrated CWDM device [175]

### *6.4.6.1. Formal abstract on published patent*

The invention provides an optical mux/demux for an optical printed circuit board. The mux/demux comprises: a first waveguide formed on a support layer for carrying a wavelength division multiplexed optical signal; a separator/combiner for separating the wavelength division multiplexed signal into component signals of corresponding wavelengths or for combining component signals into the said wavelength division multiplexed signal; and plural second waveguides, each for receiving or providing one or more of the said component signals, wherein the separator/combiner is at a predetermined location relative to the waveguides.

### *6.4.6.2. Description*

Prior art concerning optical wavelength division multiplexing and demultiplexing structures mostly applies to singlemode optical waveguides and requires expensive and highly precise fabrication processes. The foremost example of such a structure is the Arrayed Waveguide Grating (AWG), which is a passive wavelength multiplexer / demultiplexer which relies on the interference of multiple single mode optical signals to separate one wavelength-multiplexed-signal spatially, such that each wavelength-encoded signal is directed to a separate waveguide.

The application of optical PCB technology to Very Short Reach applications is still a very novel concept, but as "in the box" data-rates approach 10 Gb/s and beyond, it will most likely become a very attractive option. In addition to this, WDM structures (most likely CWDM) on the optical backplane would allow a significant increase in bandwidth by permitting multiple signals to be conveyed along single waveguides. One can foresee that interest in multimode WDM solutions on the backplane would increase with the predicted escalation of local bandwidth requirements.

This invention describes planar WDM structures, which form part of the optical layer and are as such fabricated in the same step as the optical waveguides. In **Figure 6-12** a prism structure is shown, however many different types of wavelength splitting structures are possible and are protected within this patent.

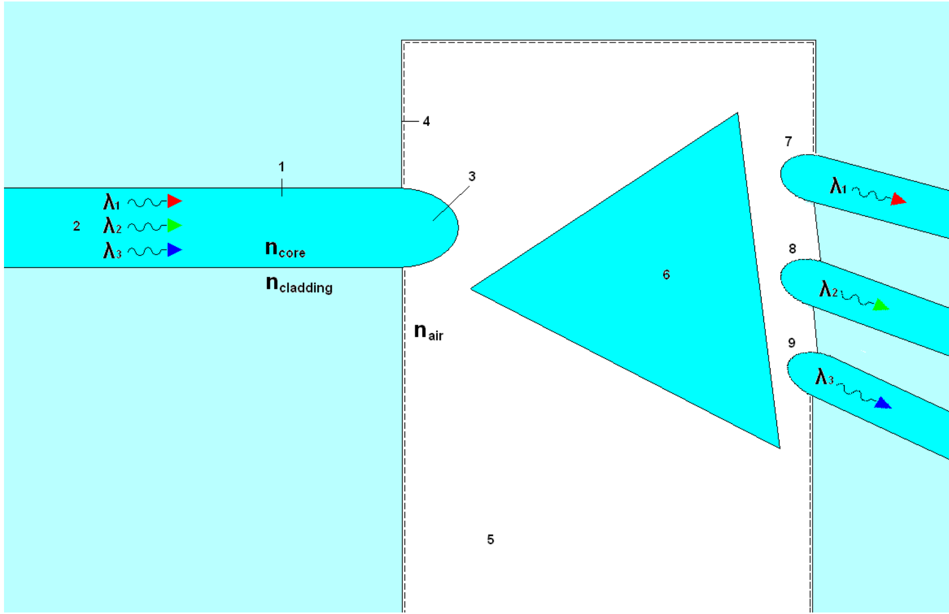


**Figure 6-12: Embedded prism style WDM multiplexer / demultiplexer structure**

### 6.4.7 Patent 7: Method of manufacturing multimode CWDM multiplexer / demultiplexer structures [176]

#### *6.4.7.1. Formal abstract on published patent*

The invention provides an optical mux/demux for an optical printed circuit board. The mux/demux comprises: a first waveguide formed on a support layer for carrying a wavelength division multiplexed optical signal; a separator/combiner for separating the wavelength division multiplexed signal into component signals of corresponding wavelengths or for combining component signals into the said wavelength division multiplexed signal; and plural second waveguides, each for receiving or providing one or more of the said component signals, wherein the separator/combiner is at a predetermined location relative to the waveguides.

#### *6.4.7.2. Description*

This patent is related to Patent 6: Passive optical waveguide integrated CWDM device and describes various methods of manufacturing the WDM structures described in that patent filing. These methods include creating the structures photolithographically **Figure 6-13** and by use of laser ablation to carve out the required structures.

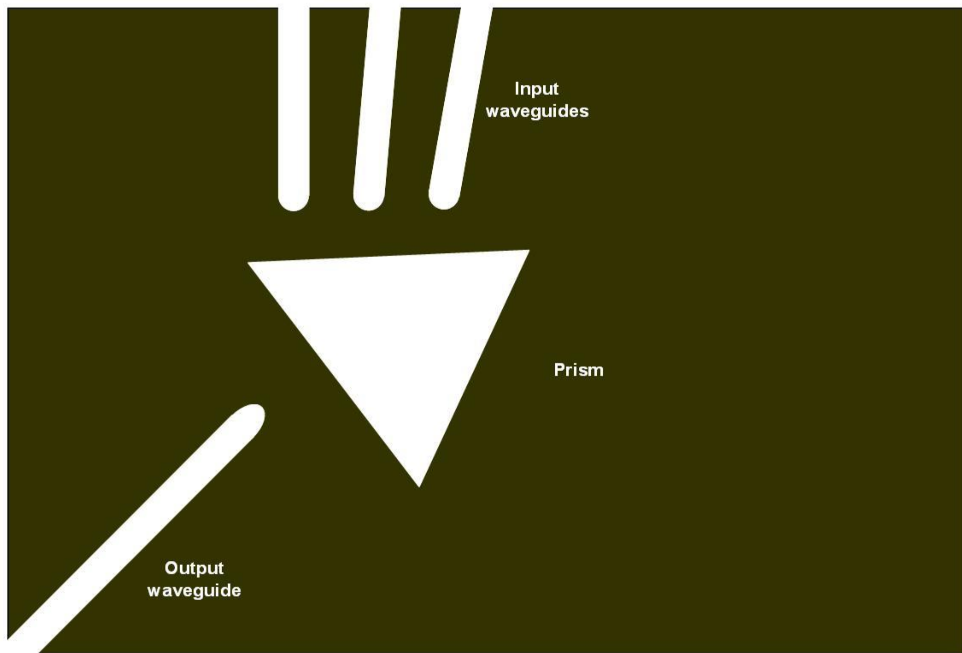


**Figure 6-13: Photolithographic mask to create WDM structures in photo-cureable polymer**

## 6.4.8 Patent 8: Optical adapter module [177]

### *6.4.8.1. Formal abstract on published patent*

The invention provides an adapter for an optical printed circuit board, the adapter comprising a socket for receiving a daughter card for connecting to a said optical printed circuit board; and a connector for engagement with the optical printed circuit board arranged such that when the connector engages with the optical printed circuit board an optical connection is established between the optical printed circuit board and the adapter.

### *6.4.8.2. Description*

The invention depicts a low-cost and reliable method of connecting conventional disk drives to an electro-optic backplane through the use of optical adapter modules, which are assembled and fixed to the backplane. The concept is naturally extendable to accommodate connection of any conventional line card to an electro-optical backplane.

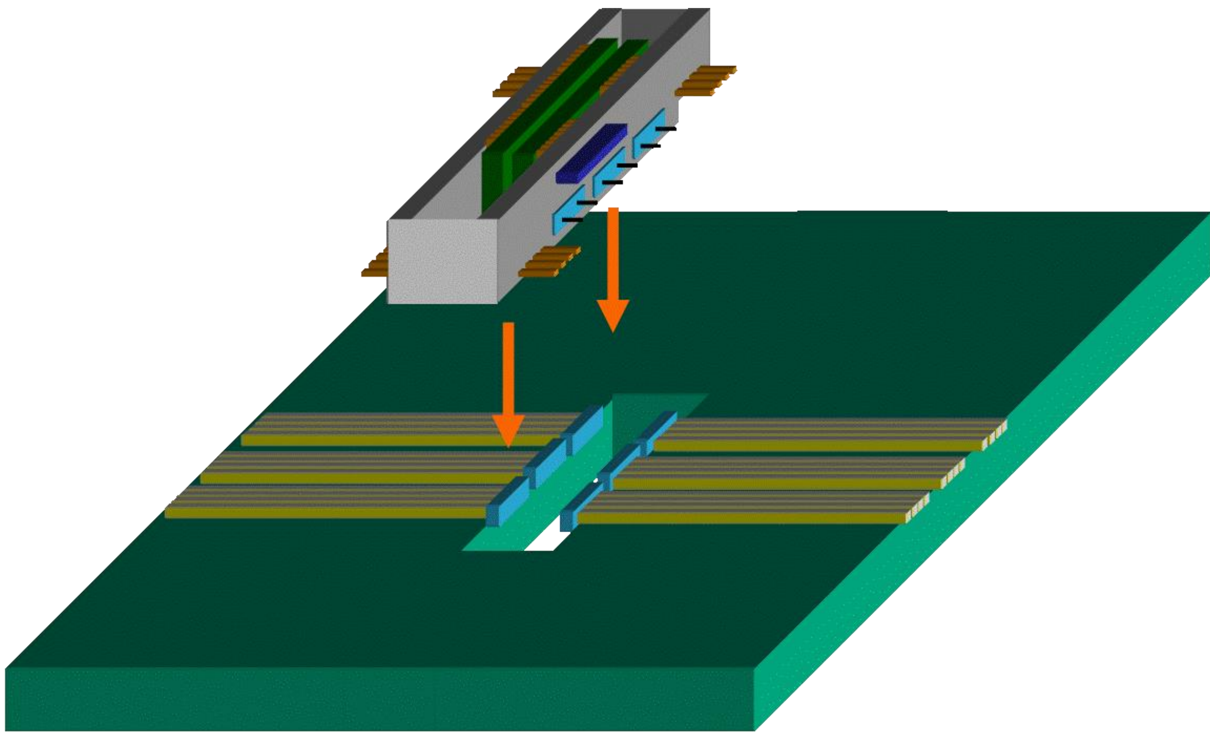

**Figure 6-14: Optical drive adapter module**

In the key implementation the conventional electrical contacts of the line-card will interface with compliant contacts in the adapter module as they would to a conventional backplane receptacle e.g. SAS / SATA backplane connectors.

Electro-optic conversion of designated high speed data lines is then performed by the appropriate circuitry in the adapter module (**Figure 6-14**) and the optical signals are launched into optical waveguides embedded in the substrate of the backplane via the proprietary parallel optical interface embedded in the module. The number of parallel optical interfaces per module can be extended without internal constraint. The only constraints will arise from the environment e.g. the real estate on the backplane available for the connector etc.

### 6.4.9 Patent 9: Fabricating a hybrid electro-optical PCB with optical surface layers [178]

#### *6.4.9.1. Formal abstract on published patent*

The invention provides a method of making an electro-optical printed circuit board, the method comprising: providing a support layer having thereon surface mounted electric components within a region of the support layer; forming one or more surface mounted optical components on the surface of the electro-optical printed circuit board; and during formation of the one or more surface mounted optical components shielding the region of the electro-optical printed circuit board where the surface mounted electric components are formed.

#### *6.4.9.2. Description*

The proposed invention comprises a series of methods of fabricating an electro-optic PCB containing an optical layer on at least one of its surfaces (**Figure 6-15**). The methods focus on the protection from contamination of conventional surface interface structures on the surface during the optical PCB fabrication process.

In one preferred example, the surface structures are protected using a photolithographic process of masking out the clear-out sections of the surface i.e. those sections of the surface explicitly reserved for conventional surface interface structures, when curing the core and cladding layers of the surface optical layer.

In another example, a protective lid, panel or sheet is used to shield the surface interface structures from any contamination during the optical layer fabrication process. The panel is then removed after the optical layer has been deposited.

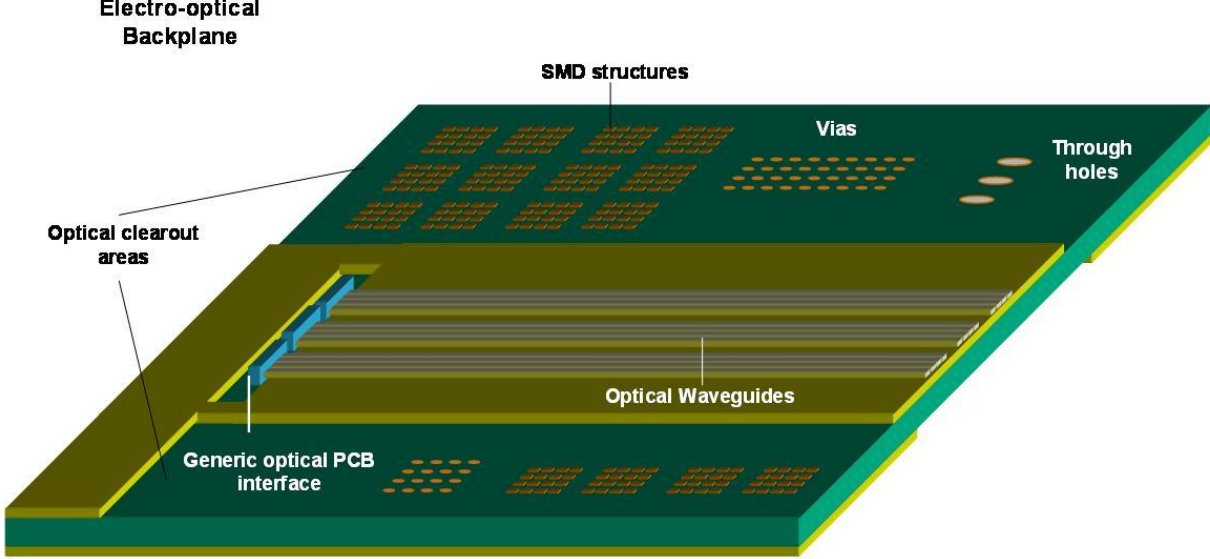


**Figure 6-15: Electro-optical PCB with partial optical layer on surface**

### *6.4.9.3. Implementation*

On the electro-optical midplane developed for the Candeo demonstration platform in 2008 only part of the top surface (shown in **Figure 6-16**) was devoted to optical polymer interconnect, while other areas devoted to electrical test holes, vias and connectors were kept clear of optical material.

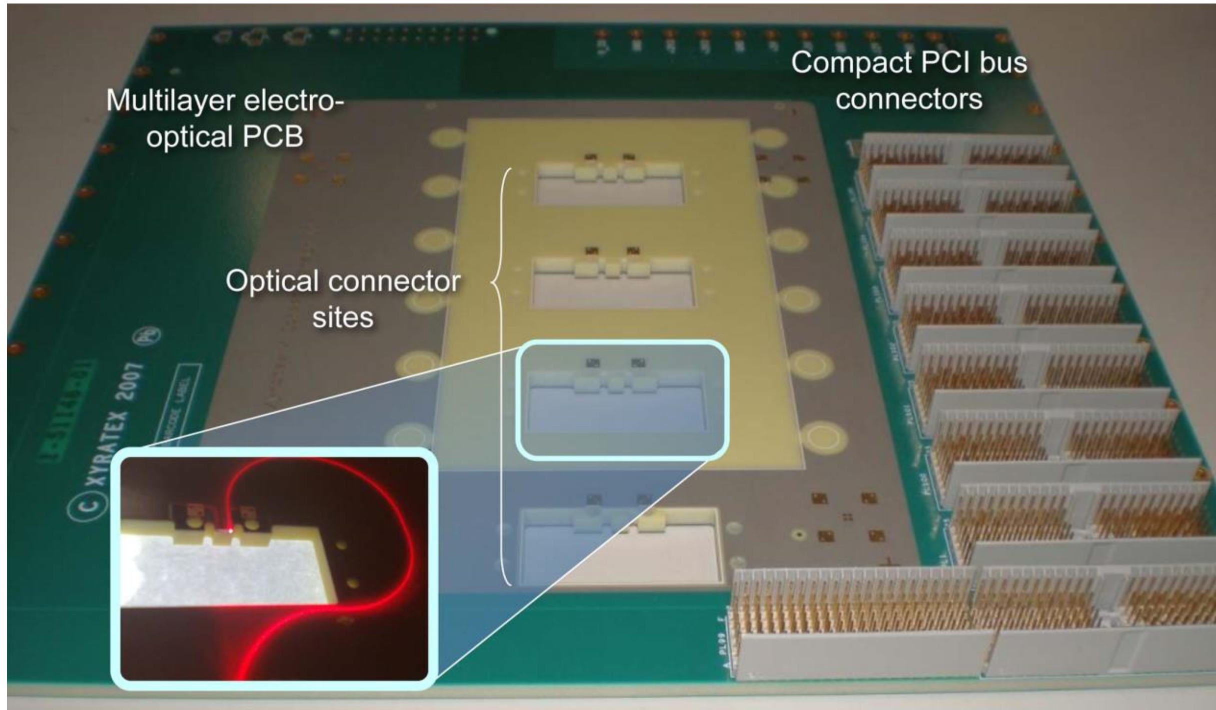


**Figure 6-16: FirstLight electro-optical midplane with optical connector slots milled out and electronic CompactPCI connectors populated. A close-up view of a connector aperture is shown with a single curved waveguide and its egress point illuminated with 635 nm visible light.**

### 6.4.10 Patent 10: Lossless tapered waveguides [179]

#### *6.4.10.1. Formal abstract on published patent*

A method of making a waveguide, the method including depositing discrete units of optical core material in a pattern of the waveguide, and controlling the refractive index of the discrete units such that the refractive index of the waveguide varies along its length.

#### *6.4.10.2. Description*

Optical waveguide tapers were investigated by UCL as part of the Storlite project (2005). The underlying idea is that optical waveguides should be wider at the ingress point to maximise misalignment tolerance, but as small as possible at the egress point to increase the probability of all light being captured by the receiving system (e.g. photodiode or passive optical element) given conventional lateral misalignment values.

This has however so far not been reliably achieved, due to the critical problem of modal mismatch along the taper giving rise to optical leakage of the propagating signal into the cladding region along the path from a larger core area to the smaller core area. Previous work on fibre tapers has shown that the choice of taper length and linearity of the taper can be modified to improve coupling efficiency [187][188].

However this invention theorises that the coupling efficiency along the physical waveguide taper may be further improved by varying the refractive index of the cladding material, core material or both along the taper, in order to generate a changing refractive index difference along the taper, such as to compensate for the NA change caused by the change in width (**Figure 6-17**).

This patent describes various concepts and fabrication techniques to enable such structures.

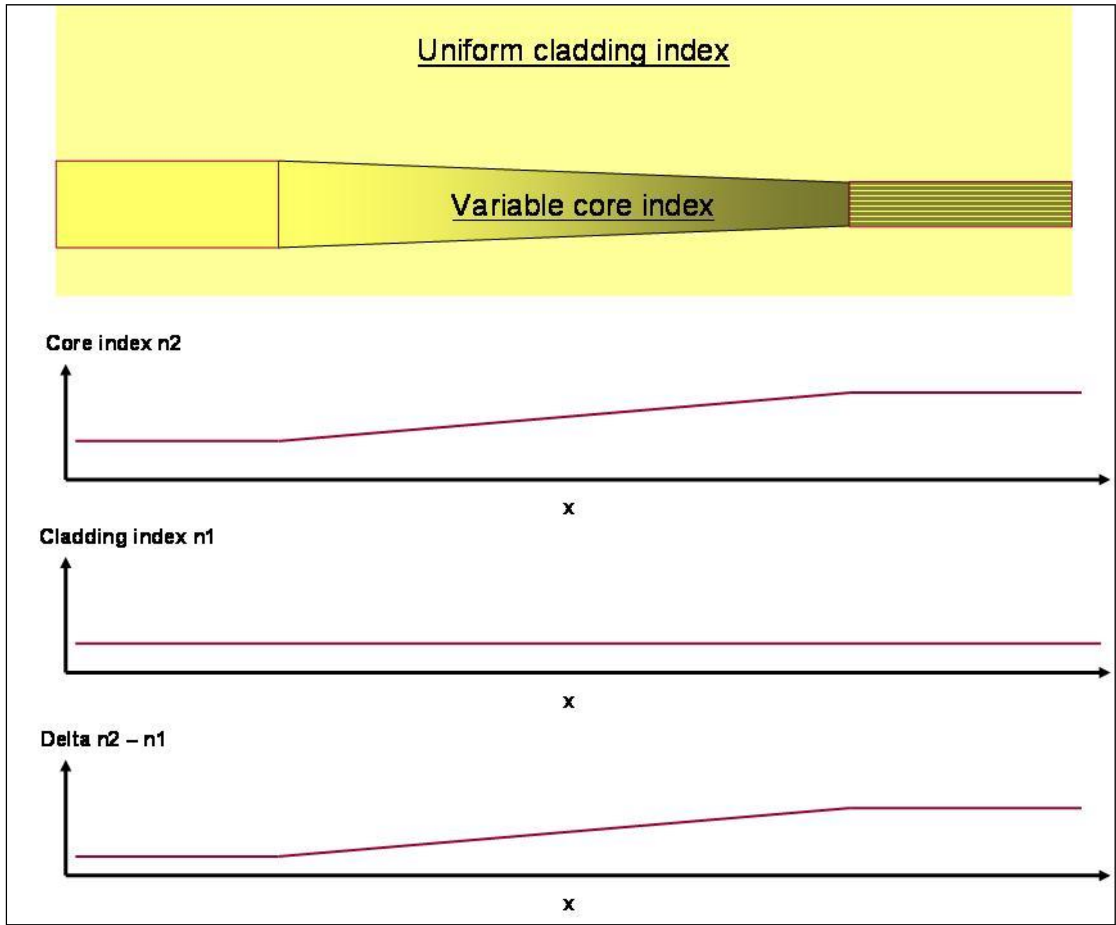


**Figure 6-17: Tapered waveguide with core index variation**

### 6.4.11 Patent 11: Directly pluggable orthogonal in-plane optical PCB connector [180]

#### *6.4.11.1. Formal abstract on published patent*

The invention provides an optical connector for connecting a user circuit to an optical backplane, in which the backplane has one or more waveguides on it for carrying optical signals, the connector comprising: a first optical interface provided on the user circuit for receiving optical signals from the backplane or for transmitting optical signals for passage along the one or more waveguides; a second optical interface provided on the backplane for receiving optical signals from the user circuit or for transmitting optical signals to the user circuit; alignment features provided on each of the user circuit and the backplane arranged to align the first and second optical interfaces such that upon insertion of the user circuit to the backplane, the optical interfaces are aligned in the direction of insertion.

#### *6.4.11.2. Description*

Currently, in order to provide a pluggable connection to an optical printed circuit board one must either use an out-of-plane or an in-plane optical interface. The out-of-plane optical interface would usually comprise an angled mirror to divert light at right-angles from the connector interface into the waveguides of the optical PCB. The main problem with this approach is the cost of the right-angled mirror and the additional optical loss incurred across the interface. The optical loss budget on an optical PCB is a critical issue and must be minimised wherever possible.

The in-plane optical interface up to now has required a connector mechanism to move the optical platform orthogonally with respect to the direction of insertion in order to stop the mechanical registration features from catching. It has not been possible to create a directly pluggable connector to an in-plane interface without such a mechanism. The advantage of an in-plane interface is that coupling components (e.g. micro-lens arrays) do not include mirror or other deflection structure and are thus cheaper and incur less optical loss.

This invention proposes a means of achieving an orthogonal connection to an in-plane optical interface, but without pins (or other salient registration features) so that the transceiver or optical interface can be plugged simply and directly without the need for a complex engagement mechanism to pull and push the optical interface (on the transceiver platform) in a direction orthogonal to the direction of insertion in order to stop the pins catching.

Specifically, registration features are proposed which are aligned in the direction of insertion (as opposed to orthogonal to the direction of insertion) such that an optical device on a line card may mate directly, without the

need for a secondary engagement step. In a preferred example these registration features constitute alignment rails on one element and compliant alignment grooves on the other element. The alignment grooves and/or rails will be tapered to ensure that the alignment features engage smoothly and with ample misalignment tolerance at the first point of contact. As an optical device on a line card engages with the PCB receptacle, the rails on the optical device slot into the grooves of the receptacle

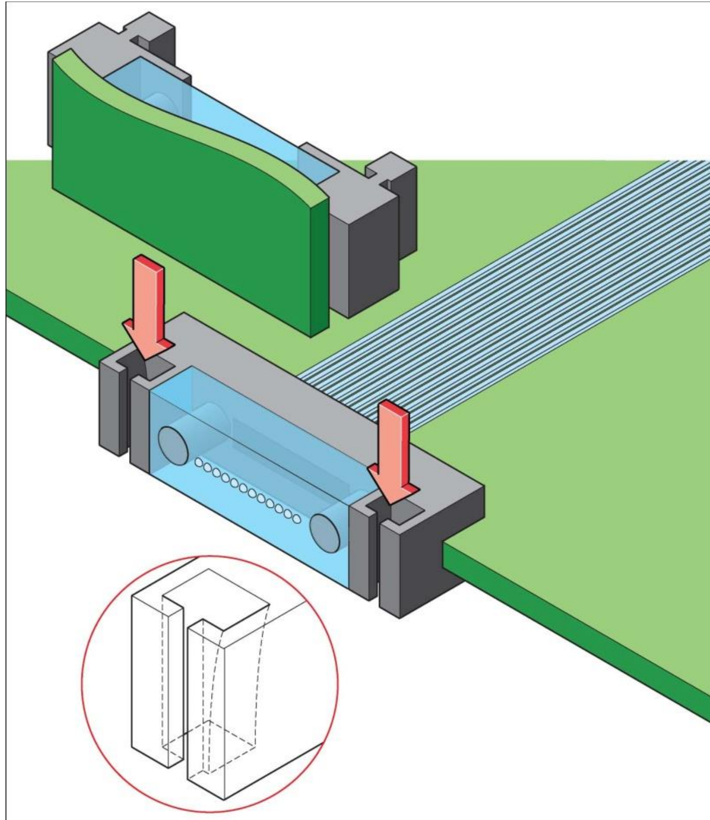

**Figure 6-18: Orthogonal optical connector with alignment rails and grooves. A free space lens coupling system allows physical contact between optical interfaces to be omitted**

### 6.4.12 Patent 12: Optical capstan to simplify routing of optical waveguides [181]

#### *6.4.12.1. Formal abstract on published patent*

The invention provides an optical printed circuit board connector, comprising: a housing having a major plane; an optical interface for connection in use to another optical interface on a device to which in use the optical printed circuit board connector is arranged to be connected, in which the optical interface on the connector is mounted such that it is twistable about a vertical axis in the major plane to vary the launch angle of light from the interface with respect to the housing.

#### *6.4.12.2. Description*

This invention puts forward an optical PCB connector type, which will remove the dependence of the waveguide launch angle on the standard orientation of the connector receptacle on the optical PCB. The part of the connector, which houses the optical interface will be allowed to rotate around an axis orthogonal to the plane of the optical PCB.

In a preferred example this will be enabled by arranging the optical interface on a flexible laminate material that will allow it to be twisted with respect to the body of the connector.

The invention targets both active and passive optical connectors and either in-plane or out-of-plane connections to the optical PCB:

- Active optical connectors are connectors whereby the active photonics (lasers, photodiodes) are located on the interface to the optical PCB usually with a lens to control the beam profile into and out of the waveguides. A number of previous Xyratex patents are based on active connectors.
- Passive optical connectors are connectors where there are no active photonics on the interface, but rather waveguides (e.g. fibres, polymer waveguides) that are carrying light from the active photonics located somewhere else. These waveguides could include optical fibres or waveguides on a flexible laminate which have been demonstrated by Swiss PCB company Varioprint.
- In-plane optical PCB connectors are those which inject light directly into the plane of the optical PCB i.e. directly into the embedded waveguides without the need for deflection of the light by mirrors.
- Out-of-plane optical PCB connectors are those whereby the light is launched in a direction orthogonal to the optical PCB and captured by deflection optics such as 45° mirrors, which turn the light at right angles and launch it into the embedded waveguides.

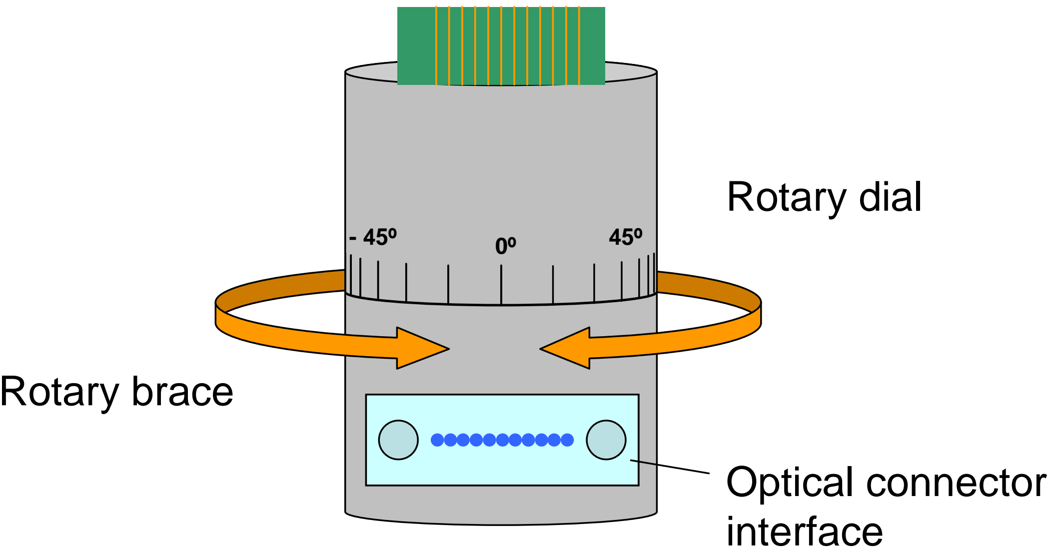


**Figure 6-19: Rotary optical brace with angle adjustable optical interface**

## 6.4.13 Patent 13: Directly pluggable orthogonal in-plane optical PCB connector [182]

### *6.4.13.1. Formal abstract on published patent*

The invention provides an interconnect for a data storage system, to enable optical communication between a data storage device and a backplane to which the data storage device in use is to be connected, wherein the interconnect includes an electrical power connection for providing power to a said data storage device; and, an optical engine for generating and receiving optical signals for transmission of data between the data storage device and the backplane.

### *6.4.13.2. Description*

This invention comprises a design solution for a low cost pluggable intermediary optical connection scheme between a storage device such as a hard disk drive (HDD), solid state device (SSD), optical storage device (such as a CD, DVD, HDDVD, Blu-Ray), holographic storage device or the emerging optical USB standard (USB 3.0) to an electro-optical PCB in a data storage system, computer or media player; in short any systems in which a storage device plugs into a mother board or interface card to a mother board.

The following description focuses on a preferred example of the invention applied to an array of disk drives in a data storage system.

A typical electronic SAS / SATA interface will include contacts for 2 high speed bidirectional channels and power. The invention would seek to replace the 2 high speed bidirectional gold contacts on the SAS / SATA interface with a parallel optical engine, which will provide two high speed bidirectional optical links. This would require a dual duplex channel “optical engine” i.e. a multi-chip module to convert 2 electronic transmit lines to 2 optical transmit lines (laser) and convert 2 optical receive lines to 2 electronic receive lines.

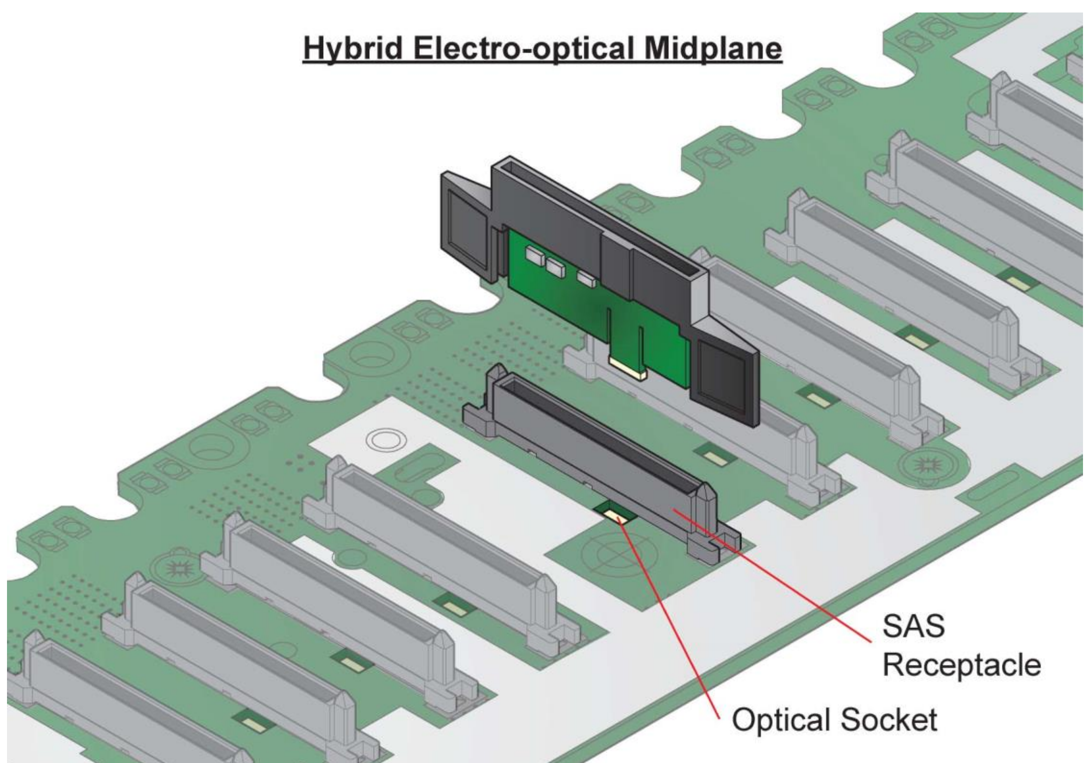


**Figure 6-20: Drive dongle with optical engine incorporated**

### *6.4.13.3. Implementation*

Variations of this invention have been realised in recent years by the author including an optically enabled data storage system demonstrator developed on the PhoxTroT project [88].

### 6.4.14 Patent 14: Modular interconnectable electro-optical backplane [183]

#### *6.4.14.1. Formal abstract on published patent*

The invention provides a component electro-optical printed circuit board backplane for assembly in a modular electro-optical backplane, the component electro-optical printed circuit board backplane comprising: optical channels for the propagation and transmission of optical data signals and electrical channels for the propagation and transmission of power and control signals; connectors for connection in-plane to at least one other component electro-optical printed circuit board; and, at least one socket for receiving, in use, a user circuit. The invention also provides a modular backplane made up of plural such component backplanes.

#### *6.4.14.2. Description*

The Xyratex roadmap puts forward proposals for very large scale integrated data centres, which will require large backplanes to support the interconnection of power, storage, controller and other devices across the rack. These backplanes may need to be as high as the racks and therefore potentially many metres in length. Reliable mass-production of printed circuit boards on this scale would pose considerable strain on the manufacturers. In most cases the equipment would simply not be available to build boards on this scale. If the equipment were available, the board size requirements would incur costly additional fabrication processes to cater for the added weight of the board, additional structural reinforcement and strongly impact yield and therefore the cost of the backplanes. An alternative solution would be that the complete backplane (henceforth “super-backplane”) be comprised of a number of smaller interconnectable printed circuit boards (henceforth “sub-backplanes”), each of which is of a standard size and configuration, which will not incur additional cost from the PCB fabricator. This way the super backplane can be constructed by plugging the sub-backplanes together in the correct order, and deconstructed by unplugging the sub-backplanes.

When the super-backplane is constructed the sub-backplanes will need to be electrically and electronically connected in order to convey power, control data and low-speed bus data across the entire super-backplane. In this invention it is proposed that high speed data is conveyed optically in order to provide a manageable and future-proof solution to convey data of ever increasing signal frequencies over the longer interconnect distances typical of rack-scale backplanes. Therefore each sub-backplane will require pluggable electrical and electronic connectors and receptacles as well as pluggable optical connectors (**Figure 6-21**).

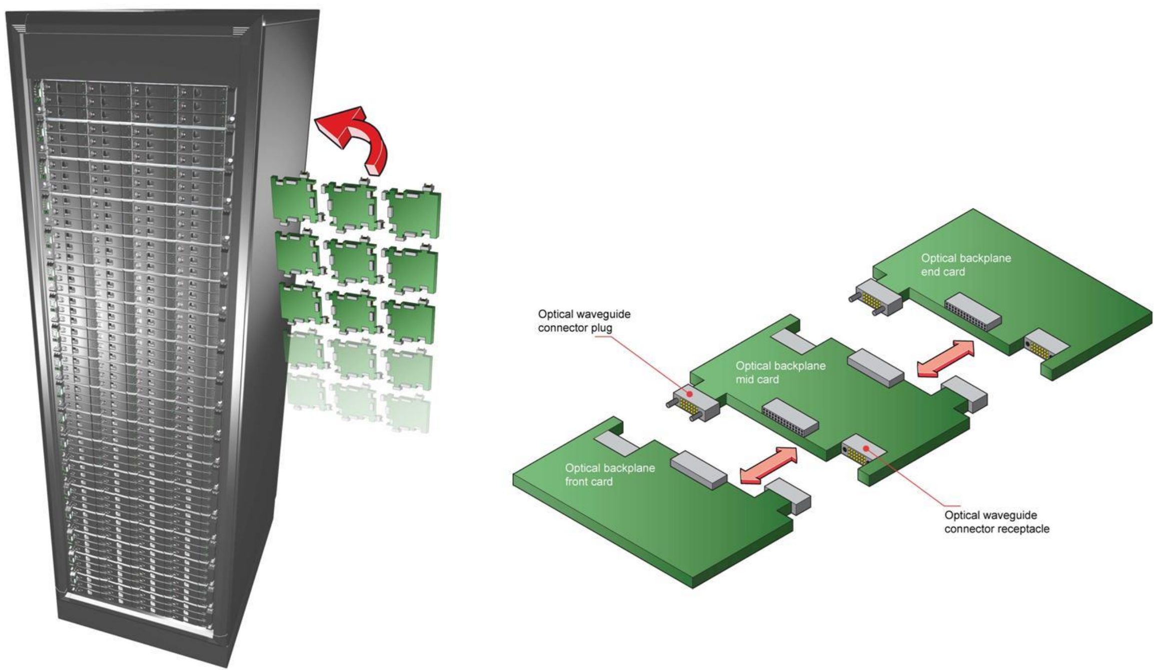


**Figure 6-21: Rack scale data centres with modular interconnectable electro-optical backplanes**

## 6.4.15 Patent 15: Optical amplification devices for polymer optical printed circuit boards [184]

### *6.4.15.1. Formal abstract on published patent*

The invention provides an amplification module for an optical printed circuit board, the optical printed circuit board comprising plural polymer waveguide sections from independent waveguides, each of the sections being doped with an amplifying dopant, wherein the plural waveguide sections are routed so as to pass through an amplification zone in which the plural polymer waveguide sections are arranged close or adjacent to one another, the amplification module comprising: a pump source comprising plural light sources arranged to provide independently controllable levels of pump radiation to each of the plural waveguide sections. In an embodiment, the amplification module also includes plural polymer waveguide sections corresponding to the plural polymer waveguides of the printed circuit board on which in use the amplification module is to be arranged, each of the sections being doped with an amplifying dopant.

### *6.4.15.2. Description*

The invention describes low cost devices to enable the on-board amplification of optical signals propagating along polymer waveguides embedded in an optical PCB.

These devices comprise controllable light sources to provide orthogonal pumping of the doped polymer optical waveguides in designated regions known as amplification nodes, where many waveguides are aggregated (See section 6.4.16).

The amplification module will be electrically connected to an electric layer of the electro-optical PCB through leads or stud bumps. If the optical layer is on a surface between the amplification module and the electrical contacts (e.g. the top surface) then clearance in the optical layer will be designed to allow access of the amplification module’s contacts to the contact lands on the electrical layer.

If the optical layer is not a surface layer but embedded in the board then clearance in the layers above the amplification node will be required to enable optical access by the amplification module (**Figure 6-22**).

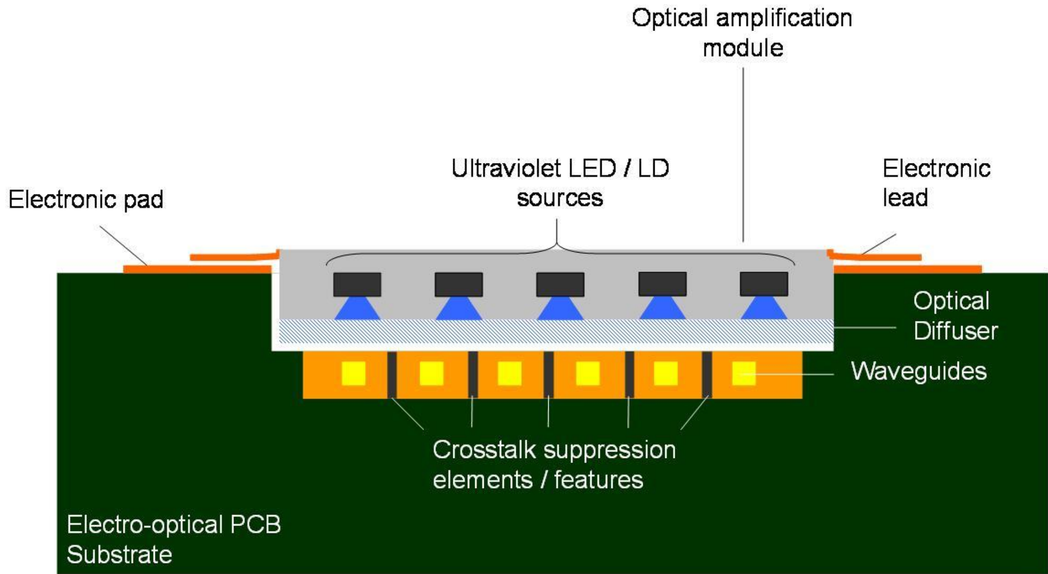


**Figure 6-22: Optical amplification module arranged over an embedded polymer waveguide layer**

### 6.4.16 Patent 16: Design and manufacture of polymer optical waveguide amplification structures [185]

#### *6.4.16.1. Formal abstract on published patent*

The invention provides an optical printed circuit board, comprising: plural polymer waveguide sections from independent waveguides, each of the sections being doped with an amplifying dopant; an optical pump source to pump the plural polymer waveguide sections, wherein the plural waveguide sections are arranged close or adjacent to one another such that the optical pump source is able to pump plural of the optical waveguide sections.

#### *6.4.16.2. Description*

The invention describes low cost design structures for polymer optical waveguides to enable the on-board amplification of optical signals on an optical PCB, which would be required for example to support wavelengths, to which the waveguide material is not suited. The invention also describes varied techniques of manufacturing polymer optical waveguides, which include active regions where the amplification can be carried out with an appropriate pump source. In order to minimise the number of pump sources the waveguide layout will be designed to create waveguide confluences on specified areas (amplification nodes) where the active regions of as many waveguides as possible can be pumped by the same source device.

Typical backplane PCBs in data storage applications would require hundreds of high speed channels. In the High Performance Computing (HPC) domain there could be thousands. If embedded polymer optical waveguides were used to convey high speed signals in these systems then the cost of separate amplification on each waveguide would be prohibitive. In order to minimise the cost of optical amplification on all high speed optical channels it is proposed that the waveguide layout be modified such that multiple waveguides (as many as is convenient to the designer) be brought into close proximity at a certain point, henceforth referred to as the amplification node. The segment of each waveguide in an amplification node will need to be composed of an active material e.g. a polymer matrix doped with an appropriate amplification material such as lanthanide complexes or dyes in order to allow it to be pumped or energised such as to amplify any optical signals conveyed along the waveguides. This is similar to the principle of Erbium Doped Fibre Amplifiers except optical polymers can sustain greater dopant concentrations allowing the doped regions to be smaller.

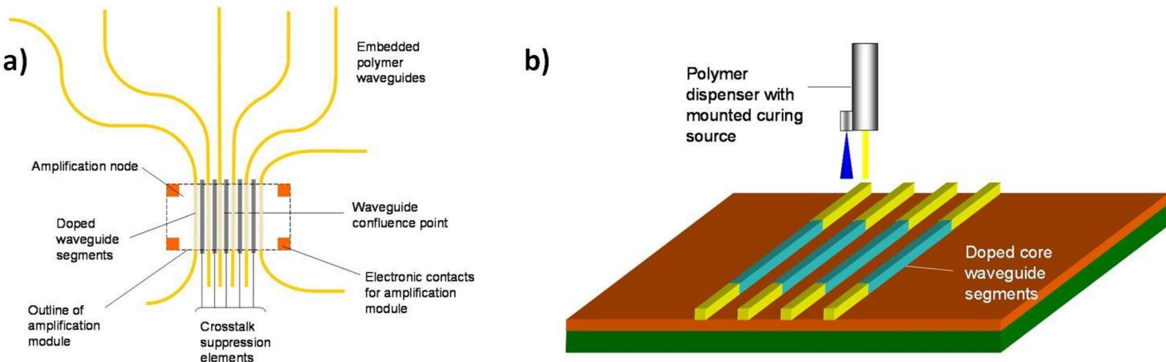


**Figure 6-23: a) Optical amplification node in which optical waveguides are brought together for collective pumping, b) Method of fabricating doped section of polymer waveguides**

## *6.5 Summary*

In this chapter, 16 patent families have been described, which seek to tackle various problems that have been identified over the course of the research and development activities, forming the basis of this thesis. The patents span diverse concepts on the theme of system embedded interconnect including connectors, assembly methods, waveguide structures and devices. Some patents, such as the active optical connector and high precision assembly method have been realised, while others set out more advanced concepts that, at the time of writing could not be tested.

# 7 CONCLUSION

The purpose of the research described in this thesis was to develop and characterise technology solutions to allow optically pluggable board-to-board and fibre-to-board connections to electro-optical printed circuit boards (OPCBs) and to investigate the commercial viability of deploying different classes of OPCB in future data centre systems through the development of appropriate demonstration platforms.

## *7.1 Summary*

### 7.1.1 Introduction to data centres and system embedded photonic eco-system

In Chapter 1 the author describes the architectures of data storage and switch systems, which form the building blocks of modern data centres and introduces modern communication protocols and the associated trends in data speeds.

In Chapter 2 the author then introduces the commercial technology eco-system, including midboard optical transceivers and optical connectors, which supports the migration of optical interconnect down through the hierarchal levels of the data centre and into system enclosures themselves. The author then explains limitations in PCB level electronic channels at increasing electronic signal modulation frequencies and where OPCB technology can provide benefits.

### 7.1.2 Pluggable active optical connectors for polymer waveguide interfaces

In Chapter 3 and Chapter 4, the author describes his research and development activities on the Storlite and FirstLight projects, on which he patented, designed, developed and successfully demonstrated two generations of active pluggable OPCB connector, which allowed peripheral devices to plug into and unplug from an electro-optical backplane with embedded multimode polymer waveguides. In addition, he has patented, co-developed with UCL and successfully implemented a technique to passively align optical devices to such waveguides with high precision.

The author designed two demonstration platforms to test and validate two generations of active pluggable optical connector and the measurement results have shown that a complex optical interconnect pattern of polymer waveguides can be deployed to convey high speed optical data across a densely populated board with a rack size of at least 6U.

### 7.1.3 Pluggable passive optical connectors for glass interfaces

In Chapter 5, the development and demonstration was reported of a complete technology eco-system (family of interdependent technologies) around embedded planar glass multimode waveguides appropriate for data centre environments. Proprietary methods were deployed by Fraunhofer IZM and ILFA GmbH to respectively pattern waveguides into thin glass foils and laminate the resulting waveguide panels into a PCB stack-up without incurring stresses in the embedded glass resulting in the successful fabrication of OPCBs. The author co-designed fibre-to-board connectors to allow parallel optical fibre cables to connect directly to OPCB integrated glass waveguides, and a corresponding board-to-board connection system to allow daughtercards to plug into an OPCB backplane.

These efforts culminated in the first successful demonstration of a fully integrated planar glass waveguide OPCB backplane and pluggable connector platform driven by system embedded and external 850 nm and 1310 nm optical transceiver technologies and validated for both fibre-to-board and board-to-board optical connectivity at data rates of up to 32 Gb/s.

In order for glass waveguide based OPCBs and their supporting technologies to be commercially viable for large scale manufacture and deployment in the near future, a number of challenges must still be overcome including glass layer number and wavelength dependence of performance i.e. how to improve transmissivity at 850 nm. These challenges continue to be addressed in research and development projects such as PhoxTroT.

### 7.1.4 Intellectual property

In Chapter 6, the author reported on 19 patents or patent applications within 16 patent families that he has generated to protect the novel technologies, the realisation of which form the basis of this thesis, or novel, untested concepts that may form the basis of future work.

## *7.2 Challenges and future work*

### 7.2.1 Higher density active connectors

The Storlite and FirstLight connectors were the first active pluggable connectors demonstrated to establish an optical connection to multimode waveguides in an OPCB. Though the form factor of the FirstLight connector prototype developed was large, consuming an impractical amount of space on both the edge of the plugging daughtercard and on the backplane, this prototype could be used to satisfy the interconnect requirements of a basic 24 drive storage array system in the following way. Due to the use of an MT compliant interface, 12 singlex channels and their drive electronics could be substituted for the 4 duplex channels deployed in the FirstLight prototype without changing the size of the optical interface. A maximum of four active connectors (2 singlex transmitter connectors and 2 singlex receiver modules) could fit on the edge of a standard data storage controller module in addition to the standard connectors required for power and low speed signals, allowing 48 duplex channels to be supported per controller and thus 96 duplex channels on the backplane. This in turn would be sufficient to provide full interconnection to 24 drives in an enclosure. In its current form, however, such a form factor would be too large to lend itself to the scalability required in modern data storage systems. Therefore, research and development efforts into active and passive pluggable connector schemes are currently underway to significantly increase the density of optical links in new edge and data centre OPCB connectors. The photonic subcomponents assembled in the FirstLight connector were operational up to 11.6 Gb/s and therefore the connector could be deployed into data storage systems to carry SAS 2 data operating at 6 Gb/s. Furthermore subcomponents, including VCSEL driver arrays, VCSEL arrays, photodiode arrays and trans-impedance amplifiers, are now available with similar chip sizes as those deployed in the FirstLight connector, which can accommodate serial data rates of over 24 Gb/s, therefore the connector could be easily upgraded to accommodate SAS 3 operating at 12 Gb/s per lane, PCI Express Gen 4 operating at 16 Gb/s per lane and SAS 4 expected to operate at 24 Gb/s per lane.

In-plane bends will be an inextricable feature in future complex waveguide layouts on OPCBs, however, bend radii as large as the minimum bend radius of 17 mm deployed on the FirstLight backplane would be difficult to accommodate within a high density PCB layout, especially where high numbers of optical waveguides are involved. Research into novel waveguide structures allowing a reduction in bend loss and corresponding reduction in minimum bend radius should be pursued to address this issue [154].

The timescale of commercial deployment of OPCB technologies will be gated by the comparative cost between implementing high speed interconnects electronically or optically. In particular, the commoditisation and associated cost reduction of optical transceiver subassembly technologies will play a key role in determining whether a cost transition point can be reached during the lifetime of, for instance, the 12 Gb/s SAS protocol or whether it will have to wait for the transition to 24 Gb/s SAS predicted in 2018.

### 7.2.2 Standardisation of polymer waveguide connectorisation

The International Electrotechnical Commission (IEC), is a not-for-profit, non-governmental organization, founded in 1906, which develops International Standards and operates conformity assessment systems in the fields of electro-technology [189]. In 2006, an IEC working group, IEC/TC86/JWG9 (Optical functionality for electronic assemblies), was established to prepare international standards and specifications for OPCBs, intended for use with opto-electronic assemblies [190]. As technology for OPCB is an interdisciplinary field between optical interconnection and electronic packaging, this group was established as a joint working group (JWG) of both TC86 (Fibre optics) and TC91 (Electronics assembly technology).

One important example of how this group is driving commercialization of OPCB technology is the development of a standard for an MT compliant optical connector, which can be assembled onto planar polymer waveguide structures. a and b show the construction and assembly of the polymer waveguide MT (PMT) connector. The PMT connector is comprised of a ferrule, cover and boot section, the external dimensions of which, once assembled match those of existing optical fibre based MT ferrules. The optical waveguide attachment region should have a maximum width of 3 mm and minimum length of 16 mm to fit within the ferrule and boot section. The waveguide strip should contain a centred group of 12 waveguides with a centre-to-centre pitch of 250 µm in direct accordance with the existing standards for single row MT fibre interfaces. The PMT ferrule is aligned to the optical waveguide interface on either a flexible or rigid optical waveguide strip such that the relative alignment of the waveguide core centres to the centres of the MT pin slots on the ferrule are in compliance with existing MT interface standards. The connector is typically fixed with quick-drying glue after which the

connector facet must be polished to ensure the waveguide end facets are in the same plane as the ferrule connecting surface. The waveguide end facets will typically be polished with a polishing cloth using techniques similar to those described in section 3.4.7.4.

Figure 7-1**Error! Reference source not found.**c shows an example of a commercial PMT connector for multimode polymer waveguides manufactured by Japanese connector company Hakusan Mfg. Co. Ltd to allow direct connection to standard optical fibre MT connectors. This PMT connector was developed according to the requirements of the JPCA (Japan Electronics Packaging and Circuits Association) standard, JPCA-PE03-01-07S.

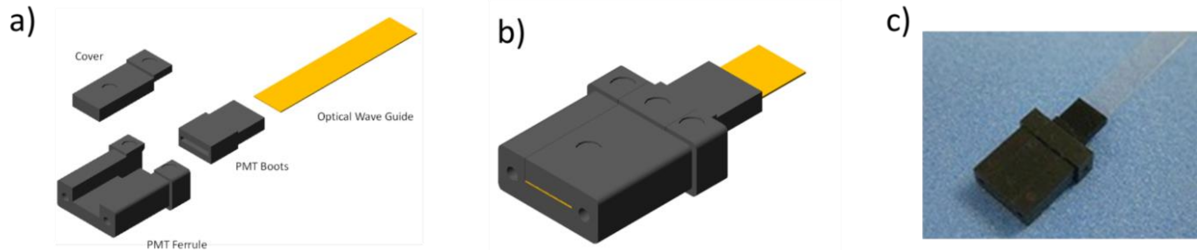


**Figure 7-1: IEC polymer MT (PMT) ferrule standard for the termination of optical waveguides a) PMT connector components, b) PMT connector structure after assembly, c) photo of commercial PMT assembled onto a polymer waveguide strip**

### 7.2.3 Connectorisation of single mode polymer waveguides

The passive precision alignment technique invented by the author has been taken forward by Huber+Suhner and Vario-optics as the basis to passively assemble connector ferrules to singlemode waveguides [145]. a and b shows the author's work reported in this thesis to align connector elements to multimode polymer waveguides using the exposed sections of the outer polymer waveguide cores to provide mechanical references. c and d show the same passive alignment concept deployed by Huber+Suhner to align connector elements to singlemode polymer waveguides.

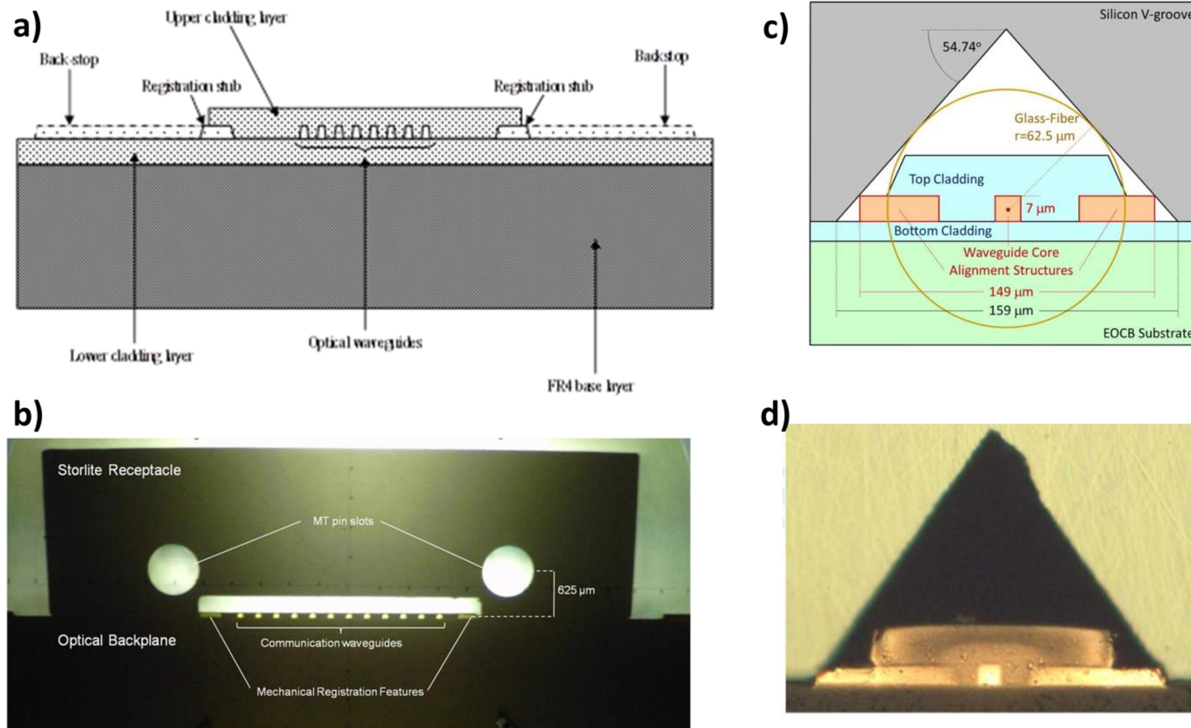


**Figure 7-2: Firstlight precision alignment concept deployed in singlemode polymer waveguide connectorisation – a) precision alignment concept diagram from Pitwon patent** [171]**, b) shadowgraph photo of precision alignment of an MT compliant Storlite receptacle passively aligned to an array of multimode waveguides, c) concept diagram for silicon V-groove to singlemode polymer waveguide** [145]**, d) photo of silicon V-groove passively aligned to singlemode polymer waveguide**

### 7.2.4 Method of futher improving robustness of passive alignment technique

In order to accommodate the possible variation in exact position of the mechanical registration guides on the MT mount due to fabrication tolerances and the possible variation in position of the waveguide registration features due to etch tolerances, a small amount of clearance between the waveguide registration features and MT mount registration guides needs to be designed in. Given the very tight tolerances involved, this clearance will typically fall within an acceptable lateral misalignment tolerance margin for multimode step-index waveguides, but could be problematic for singlemode waveguides.

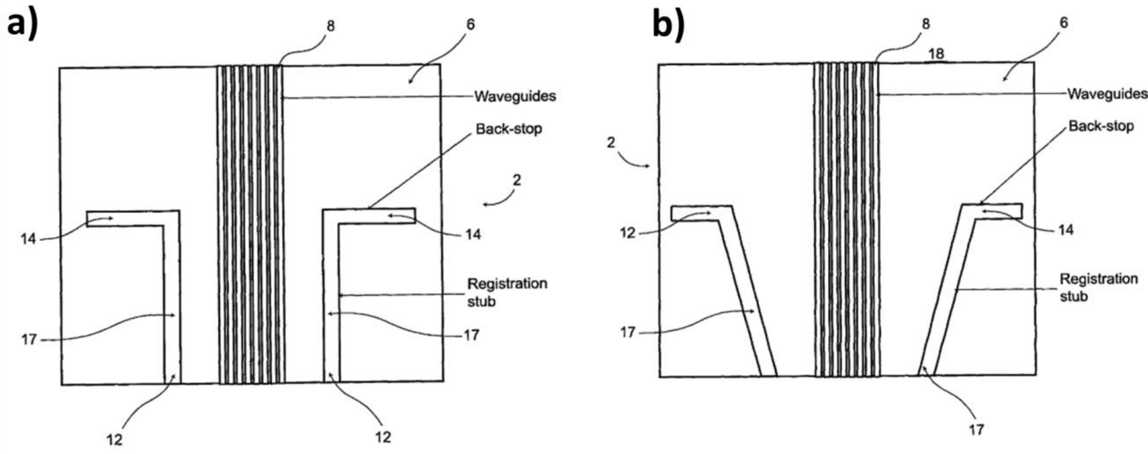


**Figure 7-3: Excerpts from published passive alignment patent showing slanted registration guides a) straight passive alignment features will require clearance, b) slanted alignment features will require clearance however will allow MT mount to accurately self-centre by squeezing the registration guides**

Figure 7-3 shows two images from the author's published patent on precision alignment [171]. Figure 7-3 a) shows straight alignment guides, similar to those deployed in the work reported in this thesis. However, as shown in Figure 7-3 b), the patent also proposes slanted registration features. This will allow the registration rails on the MT mount to deliberately slot over and start to squeeze the slanted waveguide registration features, causing them to very accurately self-centre the mount with respect to the waveguides. The alignment uncertainty in this case will be shifted to the z-axis (parallel to waveguide optical axis), denoting the distance between the MT mount surface and the waveguide facet, which are normally required to be flush. Reported work by Papakonstantinou [191] has shown that z-axis misalignment tolerances are far higher than lateral (x-y) misalignment tolerances, which means that a trade-off can be explored between waveguide registration slant angle and possible z-axis misalignment in order to further improve the robustness of the passive alignment technique in future.

### 7.2.5 Multilayer glass waveguide couplers

Regarding the deployment of multiple waveguide layers in an OPCB, the fabrication of separate waveguide layers on both the top and bottom surfaces of a single 300μm glass panel has been achieved [161], whereby the waveguide layers were separated by a vertical pitch of 250 μm. In order to couple to such a double waveguide panel, the ferrule receptacle and glass panel around the waveguide interface would need to be modified to enable 2 rows of a 6 x 12 MT ferrule to couple to the two waveguide rows simultaneously. The integration of multiple glass layers in a single OPCB is more challenging as the tolerances, which can be maintained between separate glass panels in a PCB stack-up, exceed those required between waveguide layers. If, however, separate waveguide panels in a given OPCB are not required to be accurately aligned with respect to each other, then this becomes more viable.

### 7.2.6 High yield alignment and assembly processes on glass OPCBs

The misalignment tolerance of fibre-to-waveguide coupling based on the multimode launch conditions was characterised in previous work by Fraunhofer IZM [83]. Due to the elliptical waveguide core cross-section, the misalignment tolerance is greater in the horizontal axis (parallel to waveguide layer) than in the vertical direction (perpendicular to waveguide layer), with a 1 dB misalignment tolerance of ±15 μm in the horizontal axis and ±8 μm in the vertical axis due to asymmetric flattened cross-sectional refractive index profile. In order to maximise coupling efficiency with minimal risk, an active alignment process was used for the assembly of the ferrule receptacles onto the glass waveguide interfaces in the work reported by Xyratex and Fraunhofer IZM [4]. In future, however, in order to be commercially viable, the alignment and assembly process of components onto the glass waveguide layer would need be optimised for high yield deployment. One viable approach would be to adapt high-speed pick-and-place processes to be used in combination with a semi-active alignment routine in order to provide faster and more reliable assembly. This process will be based on the use of alignment fiducials on the glass for lateral horizontal alignment and the glass surface as a mechanical stop in the vertical axis during assembly. The positional accuracy of the fiducials with respect to the waveguide positions is very high as the fiducial and waveguide diffusion structures are patterned in the same process step on the aluminium mask.

### 7.2.7 Wavelength dependent performance

The glass waveguide technologies presented perform best at 1310 nm with comparable results for 1550 nm light, which will be presented in future work. Losses at 850 nm are substantially higher, due to the formation of silver

ion clusters in the glass matrix, which induce strong intrinsic scattering at 850 nm. This effect can be mitigated by changing the glass composition and improving the waveguide process. Efforts are underway to investigate different glasses with improved fabrication processes, with the goal to provide comparable performance at the key wavelengths around 850 nm, 1310 nm and 1550 nm. The author will continue research into comparative performance of both polymer and glass waveguides.

## 7.3 Publications

During the course of the PhD programme the author wrote the following papers, which were published in journals and conference proceedings.

### 7.3.1 Journal papers

1. R. Pitwon, L. Brusberg, H. Schroeder, S. Whalley, K. Wang, A. Miller, P. Stevens, A. Worrall, A. Messina, and A. Cole, "Pluggable Electro-Optical Circuit Board Interconnect Based on Embedded Graded-Index Planar Glass Waveguides," *J. Light. Technol.*, vol. 33, no. 4, pp. 741–754, 2015.
2. K. Schmidtke, F. Flens, A. Worrall, R. Pitwon, F. Betschon, T. Lamprecht, and R. Krraehenbuhl, "960 Gb/s Optical Backplane Ecosystem Using Embedded Polymer Waveguides and Demonstration in a 12G SAS Storage Array," *Journal of Lightwave Technology*, vol. 31, no. 24. pp. 3970–3975, 2013.
3. R. Pitwon, K. Wang, J. Graham-Jones, I. Papakonstantinou, H. Baghsiahi, B. J. Offrein, R. Dangel, D. Milward, and D. R. Selviah, "FirstLight: Pluggable optical interconnect technologies for polymeric electro-optical printed circuit boards in data centers," *J. Light. Technol.*, vol. 30, no. 21, pp. 3316–3329, 2012.
4. R. Pitwon, K. Hopkins, D. Milward, M. Muggeridge, D. R. Selviah, and K. Wang, "Passive assembly of parallel optical devices onto polymer-based optical printed circuit boards," *Circuit World*, vol. 36, no. 4, pp. 3–11, 2010.
5. D. R. Selviah, A. C. Walker, D. A. Hutt, K. Wang, A. McCarthy, E. A. Fernandez, I. Papakonstantinou, H. Baghsiahi, H. Suyal, M. Taghizadeh, P. Conway, J. Chappell, S. S. Zakariyah, D. Milward, R. Pitwon, K. Hopkins, M. Muggeridge, J. Rygate, J. Calver, W. Kandulski, D. J. Deshazer, K. Hueston, D. J. Ives, R. Ferguson, S. Harris, G. Hinde, M. Cole, H. White, N. Suyal, H. U. Rehman,

and C. Bryson, “Integrated optical and electronic interconnect PCB manufacturing research,” *Circuit World*, vol. 36, no. 2, pp. 5–19, 2010.

6. I. Papakonstantinou, D. R. Selviah, R. Pitwon, and D. Milward, “Low-Cost , Precision , Self-Alignment Technique for Coupling Laser and Photodiode Arrays to Polymer Waveguide Arrays on Multilayer PCBs,” *Adv. Packag. IEEE Trans.*, vol. 31, no. 3, pp. 502–511, Aug. 2008.

7. R. Pitwon, K. Hopkins, and D. Milward, “An optical backplane connection system with pluggable active board interfaces,” *Circuit World*, vol. 33, no. 4, pp. 20–25, 2007.

### 7.3.2 Conference papers

1. H. Schröder, L. Brusberg, R. Pitwon, S. Whalley, K. Wang, A. Miller, C. Herbst, D. Weber, and K.-D. Lang, “Electro-optical backplane demonstrator with integrated multimode gradient-index thin glass waveguide panel,” in *Proc. SPIE*, 2015, vol. 9368, p. 93680U–93680U–13.

2. L. Brusberg, H. Schroeder, R. Pitwon, S. Whalley, A. Miller, C. Herbst, J. Roeder, D. Weber, and K. D. Lang, “Electro-optical backplane demonstrator with gradient-index multimode glass waveguides for board-to-board interconnection,” *2014 IEEE 64th Electronic Components and Technology Conference (ECTC)*. pp. 1033–1041, 2014.

3. R. Pitwon, A. Worrall, P. Stevens, A. Miller, K. Wang, and K. Schmidtke, “Demonstration of fully enabled data center subsystem with embedded optical interconnect,” *Opt. Interconnects Xiv*, vol. 8991, no. 0, p. 899110, 2014.

4. L. Brusberg, H. Schroder, R. Pitwon, S. Whalley, C. Herbst, A. Miller, M. Neitz, J. Roder, and K. D. Lang, “Optical backplane for board-to-board interconnection based on a glass panel gradient-index multimode waveguide technology,” *Electron. Components Technol. Conf.*, pp. 260–267, 2013.

5. R. Pitwon, H. Schröder, L. Brusberg, J. Graham-Jones, and K. Wang, “Embedded planar glass waveguide optical interconnect for data centre applications,” *Optoelectron. INTERCONNECTS XIII, Proc. SPIE*, vol. 8630, no. 0, p. 86300Z, 2013.

6. R. M. Dorward, K. Symington, L. Brusberg, J. R. Kropp, A. Miller, R. Pitwon, and S. Whalley, “Market drivers and architectural requirements for backplane inter-connect capacities in next generation PON head-end equipment in the access network,” *Int. Conf. Transparent Opt. Networks*, vol. 7, pp. 1–4, 2013.

7. K. Schmidtke, F. Flens, A. Worrall, R. Pitwon, F. Betschon, T. Lamprecht, and R. Kraehenbuehl, “960 Gb/s optical backplane using embedded polymer waveguides and demonstration in a 12G SAS storage array,” *2013 Optical Interconnects Conference*. pp. 29–30, 2013.

8. R. Pitwon, K. Wang, and A. Worrall, “Embedded Photonics Interconnect Eco-system for Data Center Applications,” *Int. Symp. Microelectron.*, vol. 2013, no. 1, pp. 361–366, 2013.

9. R. Pitwon, C. Smith, K. Wang, J. Graham-Jones, D. R. Selviah, M. Halter, and A. Worrall, “Polymer optical waveguides with reduced in-plane bend loss for electro-optical PCBs,” *Proc. SPIE*, vol. 8264, no. 0, p. 82640Z–82640Z–10, 2012.

10. R. Pitwon, K. Hopkins, K. Wang, D. R. Selviah, H. Baghsiahi, B. J. Offrein, R. Dangel, F. Horst, M. Halter, and M. Gmür, “Design and implementation of an electro-optical backplane with pluggable in-plane connectors,” in *Proceedings of SPIE*, 2010, vol. 44, no. 0, p. 76070J–76070J–12.

11. R. Pitwon, K. Hopkins, D. Milward, and M. Muggeridge, “Embedded optical interconnect technology in data storage systems,” *Proc. SPIE*, vol. 7716. p. 77161D–77161D–13, 2010.

12. R. Pitwon, “Pluggable Connector Technologies for Polymeric Electro-optical backplanes,” in *VII. ITG Workshop Photonische Aufbau- und Verbindungstechnik*, 2009.

13. Papakonstantinou, D. R. Selviah, K. Wang, R. Pitwon, K. Hopkins, and D. Milward, “Optical 8-channel 10 Gb/s MT pluggable connector alignment technology for precision coupling of laser and photodiode arrays to polymer waveguide arrays for optical board-to-board interconnects,” in *Electronic Components and Technology Conference*, 2008, no. June 2008, pp. 1769–1775.

14. R. Pitwon, K. Hopkins, and D. Milward, “An Optical Backplane Connection System with Pluggable Board Interfaces,” in *Optical Communication Systems and Networks (509)*, 2006.

### 7.3.3 Articles

In addition, the author wrote the following online articles.

1. *Pitwon, R.,* “The Light Alternative”, IEC e-tech September / October 2014 issue

2. *Pitwon, R., “*Embedded optical interconnect for use in data-storage systems*“,* SPIE Newsroom. DOI: 10.1117/2.1200912.002528

## 7.4 *Invited talks and presentations*

During the course of the PhD programme, the author delivered presentations and gave a number of invited talks at conferences and symposia. He was also invited to give lectures at universities, summer schools and the IET on his research activities.

1. Pitwon, R., "Evolution of optical interconnect technologies and architectures for exascale data centres", 3rd Symposium on Optical Interconnect in Data Centres, ECOC 2015, 29th September 2015
2. Pitwon, R., "Photonic Interconnect Technologies for Data Centre Environments", ECOC workshop on Optical Technologies for the Exascale Cloud Datacenter Era, ECOC 2015, 27th September 2015
3. Pitwon, R., "Migration of Optical Interconnect into Sub-TOR Data Centre Subsystems", ITG Workshop, Berlin, 20th May 2015
4. Pitwon, R., "The Route to Commercial Adoption of Polymer Waveguide Interconnect for Data Communication", 56th Optical Packaging Technology Conference, Japan, 7th November 2014
5. Pitwon, R., Immonen, M., "HDPuG Optical Interconnect Project" ,EPIC Technology Workshop on Photonics Integration Circuits Packaging Standardization, 19th June, 2014
6. Pitwon, R., "Standardisation of System Embedded Optical Interconnect" ,EPIC Technology Workshop on Photonics Integration Circuits Packaging Standardization, 19th June, 2014
7. Pitwon, R., "Optical Circuit Board Design", Lecture to Summer School on Optical Interconnects, Thessaloniki, 4th June, 2014
8. Pitwon, R., "Optical Interconnects in Data Storage Systems", Lecture to Summer School on Optical Interconnects, Thessaloniki, 4th June, 2014
9. Pitwon, R., "Route to Adoption of System Embedded Optical Interconnect in Data Centre Environments", 5th Annual SU2P Symposium, Glasgow, 1st April 2014
10. Pitwon, R., "International standardisation of optical circuit board technologies", Symposium on Optical Interconnect in Data Centers, Berlin, 19th March, 2014
11. Pitwon, R., "Electro-optically enabled data storage systems for exascale data centres", Symposium on Optical Interconnect in Data Centers, Berlin, 18th March, 2014
12. Pitwon, R., "Demonstration of fully-enabled data centre subsystem with embedded optical interconnect", OPTO Conference 8991, Photonics West 2014, San Francisco, 5th February 2014

13. Pitwon, R., "Introduction to PhoxTrot project", IEC / TC86 / JWG9 meeting, Charlotte, North Carolina, 31st October 2013

14. Pitwon, R., "Embedded Photonics Interconnect Ecosystem for Data Center Applications", IMAPS 46th Symposium, Florida, 1st October 2013

15. Pitwon, R. "System Embedded Photonic Interconnect Technologies for Data Centre Environments", European Cluster for Optical Interconnects (ECO) Workshop, London, 25th September 2013

16. Pitwon, R. "Migration of Embedded Optical Interconnect into Data Centre Systems", ECOC 2013, WS4, 22nd September 2013

17. Pitwon, R. "Optical Interconnect for Future Information Communication Technologies", JIEP (Japan Institute of Electronic Packaging) conference, Tokyo, 10th July 2013

18. Pitwon, R., "Optical Data Communications Interconnect for Future Information Communication Technologies", Xyratex seminar at Keio University, Tokyo, 3rd July 2013

19. Pitwon, R., "Migration of Embedded Electro-optical Interconnect Technologies into Data Centre Systems", CLEO-PR & OECC/PS 2013 Workshop, Kyoto, Japan, 28th June 2013

20. Pitwon, R. "Photonic Interconnects for Data Centers & HPC Systems: The EU FP7 PhoxTroT Approach", DataCentre Dynamics conference, Stockholm, 28th May 2013

21. Pitwon, R., "Optical Data Communications Interconnect Technologies for Future Information and Communication Technologies", Invited talk at ECO Cluster meeting, Brussels, 28th April 2013

22. Pitwon, R., "Standards Proposal for Method of Specifying Optical Waveguide Measurements", IEC / TC86 / JWG9 meeting, Kista, Sweden, 19th April, 2013

23. Pitwon, R., "Embedded Optical Interconnect Technologies for Exascale Data Centre Systems", IEEE CPMT Webinar "Electro-optical Printed Circuit Board and Interconnect Technologies and their Application to Data Center and HPC Systems", 13th February 2013

24. Pitwon, R. et al. "Embedded Planar Glass Waveguide Optical Interconnect for Data Centre Applications", OPTO Conference, Photonics West 2013

25. Pitwon R., "Technology Roadmap of Optical Interconnection", ECOC 2012 WS01 Plasmonics for Optical Interconnects, September 2012

26. R. Pitwon et al., "Polymer optical waveguides with reduced in-plane bend loss for electro-optical PCBs," OPTO Conference, Photonics West 2012

27. Pitwon R., “Optical Printed Circuit Boards”, IET Lecture, 13th January 2011

28. Selviah D., Pitwon R, “Fibre Optic Connectors – A Different View”, invited talk to SEAFOM (Subsea Fibre Optic Monitoring Group) conference, Dublin, 7th December 2010

29. Pitwon R., “Optical Backplanes in Data Storage Applications”, invited talk to Cambridge University Computer Lab and Centre for Advanced Photonics (CAPE), 16th November 2010

30. Selviah D., Pitwon R, “Optical printed circuit board and connector technology”, invited talk to 74th IEC International Electrotechnical Commission General Meeting, Seattle, 9th October 2010

31. Pitwon, R, “Optical Backplanes for Data Storage Applications”, invited talk to SEPNet (South East Photonics Network) meeting, Oclaro, Caswell, 1st July 2010

32. Pitwon, R., “Embedded optical interconnect technology in data storage systems“, Photonics Europe Conference 7716, SPIE paper 7716-31, 2010. Invited talk

33. Pitwon, R, “Design and implementation of an electro-optical backplane with pluggable in-plane connectors“, Photonics West Conference 2010, SPIE Paper Number 7607-1. Invited talk

34. Pitwon, R, “Embedded Optical Interconnect in the Data Storage Domain”, invited talk to SEPNet (South East Photonics Network) meeting, Havant, 30th September 2009

35. Pitwon, R, “Pluggable Connector Technologies for Polymeric Electro-optical backplanes“, invited talk to VII. ITG Workshop Photonische Aufbau- und Verbindungstechnik, 7/8 May 2009, Wernigerode, Germany

36. Selviah, D.R., Wang, K., Papakonstantinou, I, Yau, M., Yu, G. F., Fernández, A., Walker, A., McCarthy A., Suyal, A., Taghizadeh, M., Hutt, D., Conway, P., Zakariyah, S., Chappell, J., Hin, T.Y., Milward, D., Pitwon, R., “Integrated Electrical – Optical Substrate Manufacture – The OPCB Project” , SUMEEPnet Scientific and Technical Conference 2008, Henry Ford College, Loughborough, UK (19th March 2008)

37. Pitwon, R., “Design and Application of an Optical Backplane Connection System” , TecPreview talk to IEC DesignCon 2007 Conference, Santa Clara, California, USA, 30th January 2007

38. Pitwon, R., Hopkins, K., Milward, D., Selviah, D. R, Papakonstantinou, I., Wang, K. and Fernández, F. A. (2006). “High speed pluggable optical backplane connector” Fraunhofer IZM and VDI/VDE-IT International Symposium on Photonic Packaging: Electrical Optical Circuit Board and Optical Backplane, Munich, Germany, (16th November 2006)

39. Pitwon, R., “An Optical Backplane Connection System with Pluggable Board Interfaces”, Sixth IASTED International Multi-Conference on Wireless and Optical Communications, Optical Communication Systems and Networks, July 3-5, Banff, AB, Canada (2006)

40. Pitwon, R., “Pioneering Technologies for the Optical Era” , invited talk to Shadow Secretary of Education; David Willetts MP, Member of Parliament (3rd February, 2006)

41. Pitwon, R, “Verbindungstechnik für Optische Leiterplatten Anwendungen“ , invited talk to ITG-PAVT Meeting, Berlin (18th January 2006)

# 8 REFERENCES

[1] R. Pitwon, K. Wang, J. Graham-Jones, I. Papakonstantinou, H. Baghsiahi, B. J. Offrein, R. Dangel, D. Milward, and D. R. Selviah, “FirstLight: Pluggable optical interconnect technologies for polymeric electro-optical printed circuit boards in data centers,” *J. Light. Technol.*, vol. 30, no. 21, pp. 3316–3329, 2012.

[2] I. Papakonstantinou, D. R. Selviah, R. Pitwon, and D. Milward, “Low-Cost , Precision , Self-Alignment Technique for Coupling Laser and Photodiode Arrays to Polymer Waveguide Arrays on Multilayer PCBs,” *Adv. Packag. IEEE Trans.*, vol. 31, no. 3, pp. 502–511, Aug. 2008.

[3] R. Pitwon, K. Hopkins, D. Milward, M. Muggeridge, D. R. Selviah, and K. Wang, “Passive assembly of parallel optical devices onto polymer-based optical printed circuit boards,” *Circuit World*, vol. 36, no. 4, pp. 3–11, 2010.

[4] R. Pitwon, L. Brusberg, H. Schroeder, S. Whalley, K. Wang, A. Miller, P. Stevens, A. Worrall, A. Messina, and A. Cole, “Pluggable Electro-Optical Circuit Board Interconnect Based on Embedded Graded-Index Planar Glass Waveguides,” *J. Light. Technol.*, vol. 33, no. 4, pp. 741–754, 2015.

[5] Tolga Tekin, “PhoxTrot project.” [Online]. Available: http://www.phoxtrot.eu/. [Accessed: 16-Nov-2015].

[6] D. Reinsel, “Where in the World Is Storage: A Look at Byte Density Across the Globe,” *Ind. Dev. Model.*, vol. Oct, no. 243338, 2013.

[7] “Storage Bridge Bay ( SBB ) Specification,” *Storage Bridge Bay Working Group Inc.*, 2008. [Online]. Available: http://www.sbbwg.org.

[8] “Enterprise SSD Form Factor,” *SSD Form Factor Working Group*, 2011. [Online]. Available: http://www.ssdformfactor.org/.

[9] M. Al-Fares, A. Loukissas, and A. Vahdat, “A scalable, commodity data center network architecture,” *Proc. ACM SIGCOMM 2008 Conf. Data Commun. - SIGCOMM ’08*, vol. 38, no. 4, p. 63, Oct. 2008.

[10] “Moor Insight & Strategy,” *Intel’s Disaggregated Serv. Rack*, 2013.

[11] National Technical University of Athens, “Nephele Project.” [Online]. Available: http://www.nepheleproject.eu. [Accessed: 08-Dec-2015].

[12] T. Benson, A. Akella, and D. A. Maltz, “Network traffic characteristics of data centers in the wild,” *Proc. 10th Annu. Conf. Internet Meas. - IMC ’10*, p. 267, 2010.

[13] “IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks - Amendment 22: Equal Cost Multiple Path (ECMP),” *IEEE Std 802.1Qbp-2014 (Amendment to IEEE Std 802.1Q-2011)*. pp. 1–127, 2014.

[14] S. Han, N. Egi, A. Panda, S. Ratnasamy, G. Shi, and S. Shenker, “Network support for resource disaggregation in next-generation datacenters,” in *HotNets-XII Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*, 2013, pp. 1–7.

[15] Intel Corporation and D. Version, “Design Guide for Photonic Architecture,” no. January, 2013.

[16] Open Compute, “Open Compute Project,” 2013. [Online]. Available: www.opencompute.org.

[17] D. Grice, H. Brandt, C. Wright, P. McCarthy, A. Emerich, T. Schimke, C. Archer, J. Carey, P. Sanders, J. A. Fritzjunker, S. Lewis, and P. Germann, “Breaking the petaflops barrier,” *IBM J. Res. Dev.*, vol. 53, no. 5, p. 1:1-1:16, 2009.

[18] a. F. Benner, D. M. Kuchta, P. K. Pepeljugoski, R. a. Budd, G. Hougham, B. V. Fasano, K. Marston, H. Bagheri, E. J. Seminaro, H. X. H. Xu, D. Meadowcroft, M. H. Fields, L. McColloch, M. Robinson, F. W. Miller, R. Kaneshiro, R. Granger, D. Childers, and E. Childers, “Optics for high-performance servers and supercomputers,” *Opt. Fiber Commun. (OFC), collocated Natl. Fiber Opt. Eng. Conf. 2010 Conf.*, pp. 2–4, Mar. 2010.

[19] K. Chen, C. Hu, K. Zheng, A. V Vasilakos, and N. Technical, “Survey on Routing in Data Centers: Insights and Future Directions,” *IEEE Netw.*, vol. 25, no. 4, pp. 6–10, 2011.

[20] B. R. Rofoee, G. Zervas, Y. Yan, N. Amaya, Y. Qin, and D. Simeonidou, “Programmable on-chip and off-chip network architecture on demand for flexible optical intra-Datacenters.,” *Opt. Express*, vol. 21, no. 5, pp. 5475–80, Mar. 2013.

[21] T. Miyoshi, K. Oe, J. Tanaka, T. Yamamoto, and H. Yamashima, “New System Architecture for Next-Generation Green Data Centers: Mangrove,” *Fujitsu Sci. Tech. J.*, vol. 48, no. 2, pp. 184–191, 2012.

[22] J. Matsui, T. Yamamoto, K. Tanaka, T. Ikeuchi, S. Ide, S. Aoki, T. Aoki, T. Ishihara, M. Iwaya, K. Kamoto, K. Suematsu, and M. Shiino, “Optical Interconnect Architecture for Servers using High Bandwidth Optical Mid-plane - OSA Technical Digest,” in *Optical Fiber Communication Conference*, 2012, p. OW3J.6.

[23] AndyAtHP, "HP Networking showcases optical backplane innovation," *HP Networking*, 2011. [Online]. Available: http://h30507.www3.hp.com/t5/HP-Networking/HP-Networking-showcases-optical-backplane-innovation/ba-p/93551#.UaYa0UDVBYJ.

[24] Press Release, "Finisar and Xyratex Demonstrate 12G SAS Embedded Optical Interconnect Within Data Centre Subsystem at ECOC 2012," 2012. [Online]. Available: http://www.marketwatch.com/story/finisar-and-xyratex-demonstrate-12g-sas-embedded-optical-interconnect-within-data-centre-subsystem-at-ecoc-2012-2012-09-17.

[25] SCSI Trade Association, "Serial Attached SCSI Master Roadmap," 2014. [Online]. Available: http://www.scsita.org/library/2011/06/serial-attached-scsi-master-roadmap.html. [Accessed: 03-Mar-2016].

[26] International Committee for Information Technology Standards, "Technical Committee T10." [Online]. Available: http://www.t10.org/.

[27] Infiniband, "InfiniBand Roadmap," *InfiniBand® Trade Association (IBTA). [Online]. Available: http://www.infinibandta.org/content/pages.php?pg=technology_overview, Accessed: Apr 2016*, 2013. [Online]. Available: http://www.infinibandta.org/content/pages.php?pg=technology_overview.

[28] P. C. I. Express, "PCI Express," *Architecture*, 2004. [Online]. Available: http://www.pcisig.com/news_room/faqs/FAQ_PCI_Express_4.0/.

[29] D. Law, "IEEE 802.3 Industry Connections Ethernet Bandwidth Assessment," 2012.

[30] C. DeCusatis, "Optical Interconnect Networks for Data Communications," *J. Light. Technol.*, vol. 32, no. 4, pp. 544–552, 2014.

[31] R. Pitwon, "SEPIANet project," 2012. [Online]. Available: http://sepianet.eu/. [Accessed: 07-Mar-2016].

[32] "HDPuG Optoelectronics." [Online]. Available: http://hdpug.org/content/optoelectronics. [Accessed: 16-Nov-2015].

[33] SearchStorage.com, "Data storage trends 2011: Predictions of hot data storage technologies," *Searchstorage.Techtarget.Com*, 2011. [Online]. Available: http://searchstorage.techtarget.com/Data-storage-trends-2011-Predictions-of-hot-data-storage-technologies.

[34] C. Barrera and S. Wojtowecz, "Five storage trends for 2011," *Computer World UK*, 2011. .

[35] F. Gisin and G. Dudnikov, "State of the Art of Electrical High Speed Backplanes in Industry Today and

the Transition to Optical interconnects," *Biophotonics/Optical Interconnects VLSI Photonics/WBM Microcavities, 2004 Dig. LEOS Summer Top. Meet.*, vol. 1, pp. 7–8, 2004.

[36] W. Heirman, J. Dambre, I. Artundo, C. Debaes, H. Thienpont, D. Stroobandt, and J. Van Campenhout, "Predicting the performance of reconfigurable optical interconnects in distributed shared-memory systems," *Photonic Netw. Commun.*, vol. 15, no. 1, pp. 25–40, 2008.

[37] J. a Conway, S. Sahni, and T. Szkopek, "Plasmonic interconnects versus conventional interconnects: a comparison of latency, crosstalk and energy costs.," *Opt. Express*, vol. 15, no. 8, pp. 4474–4484, 2007.

[38] K. H. Koo, H. Cho, P. Kapur, and K. C. Saraswat, "Performance comparisons between carbon nanotubes, optical, and Cu for future high-performance on-chip interconnect applications," *IEEE Trans. Electron Devices*, vol. 54, no. 12, pp. 3206–3215, 2007.

[39] F. E. Doany, C. L. Schow, S. Member, B. G. Lee, R. a Budd, C. W. Baks, C. K. Tsang, J. U. Knickerbocker, R. Dangel, B. Chan, H. Lin, C. Carver, J. Huang, J. Berry, D. Bajkowski, F. Libsch, and J. a Kash, "Terabit / s-Class Optical PCB Links Incorporating Optical Transceivers," *Lightwave*, vol. 30, no. 4, pp. 560–571, 2012.

[40] C. DeCusatis, "Optical Networks," in *The Optical Communications Reference, 1st Edition*, 2009.

[41] C. DeCusatis, "Small form factor fiber optic connectors," in *The Optical Communications Reference, 1st Edition*, 1st ed., 2009.

[42] A. F. Rizky, N. Nishimura, Y. Nekado, T. Uemura, and H. Nasu, "Polymer waveguide-coupled 14-Gb/s x 12-channel parallel-optical modules mounted on optical PCB through Sn-Ag-Cu-solder reflow," *CPMT Symposium Japan (ICSJ), 2013 IEEE 3rd*. pp. 1–4, 2013.

[43] Lightwave, "FCI acquires MergeOptics assets," 2010. [Online]. Available: http://www.lightwaveonline.com/articles/2010/02/fci-acquires-mergeoptics-assets-84971362.html. [Accessed: 12-Feb-2016].

[44] Investorpoint.com, "Zarlink Sells Optical Products Group to Tyco Electronics," 2010. [Online]. Available: http://www.investorpoint.com/stock/TEL-TE+Connectivity+Ltd/news/30914496/. [Accessed: 12-Feb-2016].

[45] Photonics.com, "Molex Buys Luxtera's Active Optical Cable Biz." [Online]. Available: http://www.photonics.com/Article.aspx?AID=45762. [Accessed: 12-Feb-2016].

[46] Bloomberg Business, “Company Overview of Samtec Optical Group.” [Online]. Available: http://www.bloomberg.com/research/stocks/private/snapshot.asp?privcapId=78978221. [Accessed: 12-Feb-2016].

[47] Lightwave, “Volex acquires AppliedMicro technology for active optical applications,” 2013. [Online]. Available: http://www.lightwaveonline.com/articles/2013/01/volex-acquires-appliedmicro-technology-for-active-optical-applic.html. [Accessed: 12-Feb-2016].

[48] CIR Enlightended Analysis, “Active Optical Cable Markets 2015: Volume One: Data Center and High-Performance Computing Markets,” 2015.

[49] M. H. Fields, J. Foley, R. Kaneshiro, L. McColloch, D. Meadowcroft, F. W. Miller, S. Nassar, M. Robinson, and H. Xu, “Transceivers and optical engines for computer and datacenter interconnects,” *Optical Fiber Communication (OFC), collocated National Fiber Optic Engineers Conference, 2010 Conference on (OFC/NFOEC)*. pp. 1–2, 2010.

[50] O. Interconnect, “LEAP ® ON-BOARD TRANSCEIVER,” 2016. [Online]. Available: http://portal.fciconnect.com/Comergent/fci/documentation/datasheet/opticalinterconnect/oi_leap_obt.pdf. [Accessed: 28-Jan-2016].

[51] J. Duis and T. Hultermans, “The Cool Future of Optics ‘CoolBit,’” *Optical Communication (ECOC), 2014 European Conference on*. pp. 1–3, 2014.

[52] T. Yamamoto, “Optical Interconnect Technology for High-bandwidth Data Connection in Next-generation Servers,” *Fujitsu Sci. Tech. J.*, vol. 50, no. 1, pp. 117–122, 2014.

[53] Samtec, “FireFly™ Micro Flyover System.” [Online]. Available: http://www.samtec.com/cable-systems/active-optics/on-board-optical/firefly.aspx.

[54] Reflex Photonics, “120Gbps Surface Mount and Pluggable LightABLE™ optical engine,” 2015. [Online]. Available: http://www.reflexphotonics.com/products_3.html. [Accessed: 12-Jul-2015].

[55] Ultra Communications, “Rugged Vertical Connector(RVCON™).”

[56] M. G. H.W. Johnson, *High-Speed Digital Design : A Handbook of Black Magic*. 1993.

[57] Martin G. Howard W. Johnson, *High-speed Signal Propagation: Advanced Black Magic*. 2003.

[58] W. T. Beyene and C. Yuan, “An accurate transient analysis of high-speed package interconnects using convolution technique,” *Analog Integr. Circuits Signal Process.*, vol. 35, no. 2–3, pp. 107–120, 2003.

[59] M. M. Pajovic, J. Yu, Z. Potocnik, and A. Bhobe, “Gigahertz-Range Analysis of Impedance Profile and Cavity Resonances in Multilayered PCBs,” *IEEE Trans. Electromagn. Compat.*, vol. 52, no. 1, pp. 179–188, Feb. 2010.

[60] H. W. Johnson, “High-speed backplane connectors,” *IEEE Int. Symp. Electromagn. Compat.*, pp. 612–618, Aug. 2011.

[61] W.-T. Huang, C.-H. Lu, and D.-B. Lin, “Suppression of Crosstalk Using Serpentine Guard Trace Vias,” *Prog. Electromagn. Res.*, vol. 109, no. September, pp. 37–61, 2010.

[62] R. Ramzan, “Wideband Low-Reflection Transmission Lines for Bare Chip on Multilayer PCB,” *ETRI J.*, vol. 33, no. 3, pp. 335–343, Jun. 2011.

[63] H. F. Lee, C. Y. Chan, and C. S. Tang, “Embedding capacitors and resistors into printed circuit boards using a sequential lamination technique,” *J. Mater. Process. Technol.*, vol. 207, no. 1–3, pp. 72–88, Oct. 2008.

[64] D. Miller, “Optical interconnects to electronic chips.,” *Appl. Opt.*, vol. 49, no. 25, p. 70, 2010.

[65] D. A. B. Miller and H. M. Ozaktas, “Limit to the Bit-Rate Capacity of Electrical Interconnects from the Aspect Ratio of the System Architecture,” *J. Parallel Distrib. Comput.*, vol. 41, no. 1, pp. 42–52, Feb. 1997.

[66] D. A. B. Miller, “Rationale and challenges for optical interconnects to electronic chips,” *Proc. IEEE*, vol. 88, no. 6, pp. 728–749, 2000.

[67] D. Miller, “Optical interconnects,” in *Optical Fiber Communication Conference*, 2010, p. OThX1.

[68] G. Kshirsagar and M. H. Chowdhury, “Optical Interconnect Technology,” *IEEE Asis Pacific Conf. Circuits Syst.*, pp. 1426–1429, 2006.

[69] B. Carter, “Circuit Board Layout Techniques,” *Op Amps Everyone*, pp. 325–353, 2003.

[70] A. Alduino and M. Paniccia, “Interconnects: Wiring electronics with light,” *Nat Phot.*, vol. 1, no. 3, pp. 153–155, Mar. 2007.

[71] F. E. Doany, C. L. Schow, C. W. Baks, D. M. Kuchta, P. Pepeljugoski, L. Schares, R. Budd, F. Libsch, R. Dangel, F. Horst, B. J. Offrein, and J. a. Kash, “160 Gb/s bidirectional polymer-waveguide board-level optical interconnects using CMOS-cased transceivers,” *IEEE Trans. Adv. Packag.*, vol. 32, no. 2, pp. 345–359, 2009.

[72] L. Schares, J. a. Kash, F. E. Doany, C. L. Schow, C. Schuster, D. M. Kuchta, P. K. Pepeljugoski, J. M. Trewhella, C. W. Baks, R. a. John, L. Shan, Y. H. Kwark, R. a. Budd, P. Chiniwalla, F. R. Libsch, J. Rosner, C. K. Tsang, C. S. Patel, J. D. Schaub, R. Dangel, F. Horst, B. J. Offrein, D. Kucharski, D. Guckenberger, S. Hegde, H. Nyikal, C. K. Lin, A. Tandon, G. R. Trott, M. Nystrom, D. P. Bour, M. R. T. Tan, and D. W. Dolfi, "Terabus: Terabit/second-class card-level optical interconnect technologies," *IEEE J. Sel. Top. Quantum Electron.*, vol. 12, no. 5, pp. 1032–1043, 2006.

[73] R. Dangel, C. Berger, R. Beyeler, L. Dellmann, M. Gmür, R. Hamelin, F. Horst, T. Lamprecht, T. Morf, S. Oggioni, M. Spreafico, and B. J. Offrein, "Polymer-waveguide-based board-level optical interconnect technology for datacom applications," *IEEE Trans. Adv. Packag.*, vol. 31, no. 4, pp. 759–767, 2008.

[74] S. Uhlig, "Micro-Processing Applied to Optical Interconnects and High-Frequency Packaging," *Sci. Technol.*, no. 1011, pp. 1–163, 2006.

[75] P. Lukowicz, J. Jahns, R. Barbieri, P. Benabes, T. Bierhoff, A. Gauthier, M. Jarczynski, G. A. Russell, J. Schrage, W. Sullau, J. F. Snowdon, M. Wirz, and G. Troster, "Optoelectronic interconnection technology in the HOLMS system," in *IEEE Journal of Selected Topics in Quantum Electronics*, 2003, vol. 9, no. 2, pp. 624–635.

[76] N. Bamiedakis, J. Beals, R. V. Penty, I. H. White, J. V. DeGroot, and T. V. Clapp, "Cost-effective multimode polymer waveguides for high-speed on-board optical interconnects," *IEEE J. Quantum Electron.*, vol. 45, no. 4, pp. 415–424, 2009.

[77] C. Berger, M. Kossel, C. Menolfi, T. Morf, T. Toifl, and M. Schmatz, "High-density optical interconnects within large-scale systems," *Proc. SPIE*, vol. 4942, no. 2003, pp. 222–235, 2003.

[78] L. Schares, J. A. Kash, F. E. Doany, C. L. Schow, C. Schuster, D. M. Kuchta, P. K. Pepeljugoski, J. M. Trewhella, C. W. Baks, R. A. John, L. Shan, Y. H. Kwark, R. A. Budd, P. Chiniwalla, F. R. Libsch, J. Rosner, C. K. Tsang, C. S. Patel, J. D. Schaub, R. Dangel, F. Horst, B. J. Offrein, D. Kucharski, D. Guckenberger, S. Hegde, H. Nyikal, C.-K. Lin, A. Tandon, G. R. Trott, M. Nystrom, D. P. Bour, M. R. T. Tan, and D. W. Dolfi, "Terabus: Terabit/Second-Class Card-Level Optical Interconnect Technologies," *IEEE J. Sel. Top. Quantum Electron.*, vol. 12, no. 5, pp. 1032–1044, 2006.

[79] F. E. Doany, C. L. Schow, B. G. Lee, R. Budd, C. Baks, R. Dangel, R. John, F. Libsch, J. a. Kash, B. Chan, H. Lin, C. Carver, J. Huang, J. Berry, and D. Bajkowski, "Terabit/sec-class board-level optical interconnects through polymer waveguides using 24-channel bidirectional transceiver modules," *Proc. -*

*Electron. Components Technol. Conf.*, pp. 790–797, 2011.

[80] K. Matsuo, R. Yamazaki, and T. Ishigure, "Propagation mode analysis for polymer optical waveguide with graded index rectangular core," *2010 IEEE Photinic Soc. 23rd Annu. Meet.*, vol. 2, no. 1, pp. 16–17, 2010.

[81] T. Ishigure, "Graded-index core polymer optical waveguide for high-bandwidth-density optical printed circuit boards: fabrication and characterization," *Opt. Interconnects Xiv*, vol. 8991, p. 899102, Mar. 2014.

[82] G. V. S. G. Van Steenberge, N. H. N. Hendrickx, E. B. E. Bosman, J. V. E. J. Van Erps, H. T. H. Thienpont, and P. V. D. P. Van Daele, "Laser ablation of parallel optical interconnect waveguides," *IEEE Photonics Technol. Lett.*, vol. 18, no. 9, pp. 2006–2008, 2006.

[83] L. Brusberg, H. Schroder, R. Pitwon, S. Whalley, C. Herbst, A. Miller, M. Neitz, J. Roder, and K. D. Lang, "Optical backplane for board-to-board interconnection based on a glass panel gradient-index multimode waveguide technology," *Electron. Components Technol. Conf.*, pp. 260–267, 2013.

[84] L. Brusberg, H. Schroder, C. Ranzinger, M. Queisser, C. Herbst, S. Marx, J. Hofmann, M. Neitz, D. Pernthaler, and K.-D. Lang, "Thin glass based electro-optical circuit board (EOCB) with through glass vias, gradient-index multimode optical waveguides and collimated beam mid-board coupling interfaces," in *2015 IEEE 65th Electronic Components and Technology Conference (ECTC)*, 2015, pp. 789–798.

[85] K. Soma and T. Ishigure, "Fabrication of a Graded-Index Circular-Core Polymer Parallel Optical Waveguide Using a Microdispenser for a High-Density Optical Printed Circuit Board," *IEEE J. Sel. Top. Quantum Electron.*, vol. 19, no. 2, pp. 3600310–3600310, Mar. 2013.

[86] L. Schares, J. A. Kash, F. E. Doany, C. L. Schow, C. Schuster, S. Member, D. M. Kuchta, P. K. Pepeljugoski, J. M. Trewhella, C. W. Baks, R. A. John, L. Shan, Y. H. Kwark, R. A. Budd, P. Chiniwalla, F. R. Libsch, J. Rosner, C. K. Tsang, C. S. Patel, J. D. Schaub, R. Dangel, F. Horst, B. J. Offrein, D. Kucharski, D. Guckenberger, S. Hegde, H. Nyikal, C. Lin, A. Tandon, G. R. Trott, M. Nystrom, D. P. Bour, M. R. T. Tan, and D. W. Dolfi, "Terabus: Terabit / Second-Class Card-Level Optical Interconnect Technologies," *J. Quantum Electron.*, vol. 12, no. 5, pp. 1032–1044, 2006.

[87] H. Schroder, "Glass panel processing for electrical and optical packaging," *... (Ectc), 2011 Ieee ...*, pp. 625–633, 2011.

[88] R. Pitwon, A. Worrall, P. Stevens, A. Miller, K. Wang, and K. Schmidtke, "Demonstration of fully

enabled data center subsystem with embedded optical interconnect," in *Optical Interconnects Xiv*, 2014, vol. 8991, no. 0, p. 899110.

[89] R. Pitwon, K. Wang, J. Graham-Jones, I. Papakonstantinou, H. Baghsiahi, B. J. Offrein, R. Dangel, D. Milward, and D. R. Selviah, "FirstLight: Pluggable optical interconnect technologies for polymeric electro-optical printed circuit boards in data centers," *J. Light. Technol.*, vol. 30, no. 21, pp. 3316–3329, 2012.

[90] Y. Ito, S. Terada, M. K. Singh, S. Arai, and K. Choki, "Demonstration of High-Bandwidth Data Transmission above 240 Gbps for Optoelectronic Module with Low-Loss and Low-Crosstalk Polynorbornene Waveguides," *Proc. IEEE 62nd Electron. Components Technol. Conf.*, pp. 1526–1531, 2012.

[91] D. R. Selviah, A. C. Walker, D. A. Hutt, K. Wang, A. McCarthy, E. A. Fernandez, I. Papakonstantinou, H. Baghsiahi, H. Suyal, M. Taghizadeh, P. Conway, J. Chappell, S. S. Zakariyah, D. Milward, R. Pitwon, K. Hopkins, M. Muggeridge, J. Rygate, J. Calver, W. Kandulski, D. J. Deshazer, K. Hueston, D. J. Ives, R. Ferguson, S. Harris, G. Hinde, M. Cole, H. White, N. Suyal, H. U. Rehman, and C. Bryson, "Integrated optical and electronic interconnect PCB manufacturing research," *Circuit World*, vol. 36, no. 2, pp. 5–19, 2010.

[92] S. S. Zakariyah, P. P. Conway, D. A. Hutt, D. R. Selviah, K. Wang, H. Baghsiahi, J. Rygate, J. Calver, and W. Kandulski, "Polymer Optical Waveguide Fabrication Using Laser Ablation," *Electron. Packag. Technol. Conf.*, pp. 936–941, 2009.

[93] B. Swatowski, C. Amb, S. K. Breed, D. J. Deshazer, and K. W. W., "Flexible, stable, and easily processable optical silicones for low loss polymer waveguides," in *Spie …*, 2013, p. 11.

[94] Y. Takeyoshi, S. Member, and T. Ishigure, "High-Density 2 4 Channel Polymer Optical Waveguide With Graded-Index Circular Cores," *J. Light. Technol.*, vol. 27, no. 14, pp. 2852–2861, 2009.

[95] N. Bamiedakis, J. Chen, R. V Penty, and I. H. White, "Bandwidth Studies on Multimode Polymer Waveguides for 25 Gb/s Optical Interconnects," *IEEE Photonics Technol. Lett.*, vol. 26, no. 20, pp. 2004–2007, 2014.

[96] J. Chen, N. Bamiedakis, T. J. Edwards, C. T. A. Brown, R. V Penty, and I. H. White, "Dispersion studies on multimode polymer spiral waveguides for board-level optical interconnects," in *Optical Interconnects Conference (OI), 2015 IEEE*, 2015, pp. 26–27.

[97] N. Bamiedakis, J. Chen, P. Westbergh, J. S. Gustavsson, A. Larsson, R. V Penty, and I. H. White, “40 Gb/s data transmission over a 1-m-long multimode polymer spiral waveguide for board-level optical interconnects,” *J. Light. Technol.*, vol. 33, no. 4, pp. 882–888, 2015.

[98] S. Takenobu and T. Okazoe, “Heat Resistant and Low-Loss Fluorinated Polymer Optical Waveguides at 1310/1550 nm for Optical Interconnects,” *37th Eur. Conf. Expo. Opt. Commun.*, vol. 1, p. We.10.P1.31, 2011.

[99] H. Ma, A. K.-Y. Jen, and L. R. Dalton, “Polymer-Based Optical Waveguides: Materials, Processing, and Devices,” *Adv. Mater.*, vol. 14, no. 19, pp. 1339–1365, Oct. 2002.

[100] R. C. A. Pitwon, K. Wang, J. Graham-Jones, I. Papakonstantinou, H. Baghsiahi, B. Offrein, R. Dangel, D. Milward, and D. R. Selviah, “FirstLight: Pluggable Optical Interconnect Technologies for Polymeric Electro-Optical Printed Circuit Boards in Data Centers,” *J. Light. Technol.*, vol. 30, no. 21, pp. 3316–3329, 2012.

[101] K. Schmidtke, F. Flens, A. Worrall, R. Pitwon, F. Betschon, T. Lamprecht, and R. Krraehenbuhl, “960 Gb/s Optical Backplane Ecosystem Using Embedded Polymer Waveguides and Demonstration in a 12G SAS Storage Array,” *Journal of Lightwave Technology*, vol. 31, no. 24. pp. 3970–3975, 2013.

[102] R. Pitwon, H. Schröder, L. Brusberg, J. Graham-Jones, and K. Wang, “Embedded planar glass waveguide optical interconnect for data centre applications,” *Optoelectron. INTERCONNECTS XIII, Proc. SPIE*, vol. 8630, no. 0, p. 86300Z, 2013.

[103] L. Brusberg, D. Manessis, M. Neitz, B. Schild, H. Schröder, and T. Tekin, “Development of an electro-optical circuit board technology with embedded single-mode glass waveguide layer,” in *Proc. of ESTC,* 2014.

[104] L. Brusberg, H. Schroeder, R. Erxleben, I. Ndip, M. Toepper, N. F. Nissen, and H. Reichl, “Glass carrier based packaging approach demonstrated on a parallel optoelectronic transceiver module for PCB assembling,” *Electronic Components and Technology Conference (ECTC), 2010 Proceedings 60th*. pp. 269–274, 2010.

[105] H. P. Kuo, P. Rosenberg, R. Walmsley, S. Mathai, L. Kiyama, J. Straznicky, M. Mclaren, M. Tan, and S.-Y. Wang, “Free-space optical links for board-to-board interconnects,” *Appl. Phys. A*, vol. 95, no. 4, pp. 955–965, 2009.

[106] R. Dangel, C. Berger, R. Beyeler, L. Dellmann, M. Gmur, R. Hamelin, F. Horst, T. Lamprecht, T. Morf,

S. Oggioni, M. Spreafico, and B. J. Offrein, "Polymer-Waveguide-Based Board-Level Optical Interconnect Technology for Datacom Applications," *IEEE Transactions on Advanced Packaging*, vol. 31, no. 4. pp. 759–767, 2008.

[107] T. Lamprecht, F. Horst, R. Dangel, R. Beyeler, N. Meier, L. Dellmann, M. Gmur, C. Berger, and B. J. Offrein, "Passive alignment of optical elements in a printed circuit board," *56th Electronic Components and Technology Conference 2006*. p. 7 pp., 2006.

[108] B. J. Offrein, C. Berger, R. Beyeler, R. Dangel, L. Dellmann, F. Horst, T. Lamprecht, N. Meier, R. Budd, F. Libsch, and J. Kash, "Parallel optical interconnects in printed circuit boards," 2005, vol. 5990, p. 59900E–59900E–9.

[109] I. M. Soganci, A. La Porta, and B. J. Offrein, "Multichannel optical coupling to a silicon photonics chip using a single-step bonding process," *Photonics Conference (IPC), 2013 IEEE*. pp. 115–116, 2013.

[110] E. Zgraggen, I. M. Soganci, F. Horst, A. La Porta, R. Dangel, B. J. Offrein, S. A. Snow, J. K. Young, B. W. Swatowski, C. M. Amb, O. Scholder, R. Broennimann, U. Sennhauser, and G. L. Bona, "Laser Direct Writing of Single-Mode Polysiloxane Optical Waveguides and Devices," *Journal of Lightwave Technology*, vol. 32, no. 17. pp. 3036–3042, 2014.

[111] R. Kraehenbuhl, T. Lamprecht, E. Zgraggen, F. Betschon, and A. Peterhans, "High-Precision, Self-Aligned, Optical Fiber Connectivity Solution for Single-Mode Waveguides Embedded in Optical PCBs," *Journal of Lightwave Technology*, vol. 33, no. 4. pp. 865–871, 2015.

[112] J. Beals, N. Bamiedakis, A. Wonfor, R. V Penty, I. H. White, J. V DeGroot, T. V Clapp, and M. Glick, "Terabit capacity passive polymer optical backplane," *Lasers and Electro-Optics, 2008 and 2008 Conference on Quantum Electronics and Laser Science. CLEO/QELS 2008. Conference on*. pp. 1–2, 2008.

[113] T. Ishigure and Y. Nitta, "On-board fabrication of multi-channel polymer optical waveguide with graded-index cores by soft-lithography," *Electronic Components and Technology Conference (ECTC), 2010 Proceedings 60th*. pp. 275–281, 2010.

[114] Y. Matsuoka, D. Kawamura, K. Adachi, Y. Lee, S. Hamamura, T. Takai, T. Shibata, H. Masuda, N. Chujo, and T. Sugawara, "20-Gb/s/ch High-Speed Low-Power 1-Tb/s Multilayer Optical Printed Circuit Board With Lens-Integrated Optical Devices and CMOS IC," *IEEE Photonics Technology Letters*, vol. 23, no. 18. pp. 1352–1354, 2011.

[115] T. Yoshimura, M. Iida, and H. Nawata, “Self-aligned optical couplings by self-organized waveguides toward luminescent targets in organic/inorganic hybrid materials.,” *Opt. Lett.*, vol. 39, no. 12, pp. 3496–3499, 2014.

[116] E.-P. Li, W.-L. Yuan, J. Yi, L.-W. Li, and M. I. Montrose, “Analysis on the effectiveness of the 20-H rule using numerical simulation technique,” *Electromagnetic Compatibility, 2002. EMC 2002. IEEE International Symposium on*, vol. 1. pp. 328–333 vol.1, 2002.

[117] Polar Instruments, “Polar Instruments Company Web Page.” [Online]. Available: http://www.polarinstruments.com/index.html. [Accessed: 26-Feb-2016].

[118] H. Chen and J. Fang, “Effects of 20-H rule and shielding vias on electromagnetic radiation from printed circuit boards,” *Electrical Performance of Electronic Packaging, 2000, IEEE Conference on.* pp. 193–196, 2000.

[119] D. M. Kuchta, Y. Taira, C. Baks, G. McVicker, L. Schares, and H. Numata, “Optical interconnects for servers,” *Jpn. J. Appl. Phys.*, vol. 47, no. 8S1, p. 6642, 2008.

[120] K. Wang, D. R. Selviah, I. Papakonstantinou, G. Yu, H. Baghsiahi, and F. A. Fernandez, “Photolithographically Manufactured Acrylate Polymer Multimode Optical Waveguide Loss Design Rules,” in *2nd Electronics System-Integration Technology Conference Greenwich, UK*, 2008, no. September, pp. 1–27.

[121] N. Hendrickx, J. Van Erps, G. Van Steenberge, H. Thienpont, and P. Van Daele, “Laser ablated micromirrors for printed circuit board integrated optical interconnections,” *IEEE Photonics Technol. Lett.*, vol. 19, no. 11, pp. 822–824, 2007.

[122] N. Bamiedakis, R. V. Penty, and I. H. White, “Compact multimode polymer waveguide bends for board-level optical interconnects,” *J. Light. Technol.*, vol. 31, no. 14, pp. 2370–2375, 2013.

[123] M. Immonen, M. Karppinen, and J. K. Kivilahti, “Fabrication and Characterization of Polymer Optical Waveguides With Integrated Micromirrors for Optical Interconnects,” *IEEE Trans. Electron. Packag. Manuf.*, vol. 28, no. 4, pp. 304–311, 2005.

[124] W. Lee, S. H. Hwang, J. W. Lim, and B. S. Rho, “Polymeric Waveguide Film With Embedded Mirror for Multilayer Optical Circuits,” *IEEE Photonics Technol. Lett.*, vol. 21, no. 1, pp. 12–14, 2009.

[125] Y. Morimoto and T. Ishigure, “Low-loss light coupling design for a multimode polymer optical waveguide via 45-degree mirrors,” in *CPMT Symposium Japan (ICSJ), 2015 IEEE*, 2015, pp. 118–121.

[126] H. Baghsiahi, D. Selviah, and G. Yu, “Photolithographically manufactured acrylate multimode optical waveguide translation and rotation misalignment tolerances,” *, 2008. Estc 2008.*, pp. 617–622, 2008.

[127] R. C. A. Pitwon and K. M. Hopkins, “Optical connector, a communication system and a method of connecting a user circuit to an optical transceiver,” US Patent 7,625,134, 2006.

[128] R. Pitwon, K. Hopkins, and D. Milward, “An optical backplane connection system with pluggable active board interfaces,” *Circuit World*, vol. 33, no. 4, pp. 20–25, 2007.

[129] R. Pitwon, K. Hopkins, K. Wang, D. R. Selviah, H. Baghsiahi, B. J. Offrein, R. Dangel, F. Horst, M. Halter, and M. Gmür, “Design and implementation of an electro-optical backplane with pluggable in-plane connectors,” in *Proceedings of SPIE*, 2010, vol. 44, no. 0, p. 76070J–76070J–12.

[130] C. DeCusatis, “Optical Transmitters and Receivers,” in *The Optical Communications Reference, 1st Edition*, 2009.

[131] A. Corporation, “Board Design Guidelines for LVDS Systems,” 2010.

[132] T. Ishigure and Y. Nitta, “Polymer optical waveguide with multiple graded-index cores for on-board interconnects fabricated using soft-lithography.,” *Opt. Express*, vol. 18, no. 13, pp. 14191–201, Jun. 2010.

[133] D.-W. Kim, T.-W. Lee, M. H. Cho, and H.-H. Park, “High-efficiency and stable optical transmitter using VCSEL-direct-bonded connector for optical interconnection.,” *Opt. Express*, vol. 15, no. 24, pp. 15767–75, Nov. 2007.

[134] K. B. Yoon, I. K. Cho, S. H. Ahn, M. Y. Jeong, D. J. Lee, Y. U. Heo, B. S. Rho, H. H. Park, and B. H. Rhee, “Optical backplane system using waveguide-embedded PCBs and optical slots,” *J. Light. Technol.*, vol. 22, no. 9, pp. 2119–2127, 2004.

[135] H.-H. Hsu and T. Ishigure, “High-density channel alignment of graded index core polymer optical waveguide and its crosstalk analysis with ray tracing method.,” *Opt. Express*, vol. 18, no. 13, pp. 13368–78, Jun. 2010.

[136] T. Ishigure and Y. Takeyoshi, “Polymer waveguide with 4-channel graded-index circular cores for parallel optical interconnects.,” *Opt. Express*, vol. 15, no. 9, pp. 5843–50, 2007.

[137] C. L. Schow, F. E. Doany, A. V Rylyakov, B. G. Lee, C. V Jahnes, Y. H. Kwark, C. W. Baks, D. M. Kuchta, and J. A. Kash, “A 24-Channel, 300 Gb/s, 8.2 pJ/bit, Full-Duplex Fiber-Coupled Optical

Transceiver Module Based on a Single 'Holey' CMOS IC," *Journal of Lightwave Technology*, vol. 29, no. 4. pp. 542–553, 2011.

[138] W. H. Hofmann, P. Moser, and D. Bimberg, "Energy-Efficient VCSELs for Interconnects," *IEEE Photonics Journal*, vol. 4, no. 2. pp. 652–656, 2012.

[139] M. Karppinen, T. Alajoki, A. Tanskanen, K. Kataja, J.-T. T. Makinen, K. Kautio, P. Karioja, M. Immonen, and J. Kivilahti, "Parallel optical interconnect between ceramic BGA packages on FR4 board using embedded waveguides and passive optical alignments," in *56th Electronic Components and Technology Conference 2006*, 2006.

[140] S. Janz, P. Cheben, D. Dalacu, A. Delâge, A. Densmore, B. Lamontagne, M. J. Picard, E. Post, J. H. Schmid, P. Waldron, D. X. Xu, K. P. Yap, and W. N. Ye, "Microphotonic elements for integration on the silicon-on-insulator waveguide platform," *IEEE J. Sel. Top. Quantum Electron.*, vol. 12, no. 6, pp. 1402–1414, 2006.

[141] M. Diemeer, L. Hilderink, R. Dekker, and A. Driessen, "Low-cost and low-loss multimode waveguides of photodefinable epoxy," *IEEE Photonics Technol. Lett.*, vol. 18, no. 15, pp. 1624–1626, Aug. 2006.

[142] J. Van Erps and M. Wissmann, "Hot embossing of microoptical components prototyped by deep proton writing," *IEEE Photonics Technol. Lett.*, vol. 20, no. 18, pp. 1539–1541, 2008.

[143] P. Hendrickx, N.; Van Erps, J.; Van Steenberge, G.; Thienpont, H.; Van Daele, "Tolerance analysis for multilayer optical interconnections integrated on a printed circuit board," *Light. Technol. J. IEEE*, vol. 25, no. 9, pp. 2395–2401, 2007.

[144] I. D. Johnson, R. C. A. Pitwon, D. R. Selviah, and I. Papakonstantinou, "Optical printed circuit board and manufacturing method," US7936953 B2, 2011.

[145] R. Krahenbuhl, T. Lamprecht, E. Zgraggen, F. Betschon, and A. Peterhans, "High-precision, self-aligned, optical fiber connectivity solution for single-mode waveguides embedded in optical PCBs," *J. Light. Technol.*, vol. 33, no. 4, pp. 865–871, 2015.

[146] K. L. Deng, T. Gorczyca, B. K. Lee, H. Xia, R. Guida, and T. Karras, "Self-Aligned Single-Mode Polymer Waveguide Interconnections for Efficient Chip-to-Chip Optical Coupling," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 12, no. 5. pp. 923–930, 2006.

[147] Y. Ishii, T. Hayashi, and H. Takahara, "Demonstration of On-PCB Optical Interconnection Using Surface-Mount Package and Polymer Waveguide," *Integr. Vlsi J.*, pp. 1147–1152, 2003.

[148] A. Avital and E. Zussman, "Fluidic assembly of optical components," *IEEE Trans. Adv. Packag.*, vol. 29, no. 4, pp. 719–724, 2006.

[149] J. Redd, "Calculating statistical confidence levels for error-probability estimates," *Lightwave*, 2000.

[150] A. Ghiasi and R. Snively, "( 10 Gigabit Small Form Factor Pluggable Module )," *Electronics*, vol. 2003, no. 408, 2003.

[151] K. Wang, I. Papakonstantinou, D. R. Selviah, H. Baghsiahi, and G. Yu, "Design rules for Polymer Multimode Waveguide Interconnects and Design of waveguide layout for a highly constrained Electrical-Optical Backplanes," *Opt. Express*.

[152] I. Papakonstantinou, K. Wang, D. R. Selviah, and F. A. Fernández, "Transition, radiation and propagation loss in polymer multimode waveguide bends.," *Opt. Express*, vol. 15, no. 2, pp. 669–79, 2007.

[153] "Datasheet: 10 Gbps 850 nm Multi-mode SFP + Transceiver," *MRV Communications, Inc.*, 2010. [Online]. Available: www.mrv.com/.

[154] R. Pitwon, C. Smith, K. Wang, J. Graham-Jones, D. R. Selviah, M. Halter, and A. Worrall, "Polymer optical waveguides with reduced in-plane bend loss for electro-optical PCBs," *Proc. SPIE*, vol. 8264, no. 0, p. 82640Z–82640Z–10, 2012.

[155] H. Baghsiahi, K. Wang, W. Kandulski, R. C. a Pitwon, and D. R. Selviah, "Optical waveguide end facet roughness and optical coupling loss," *J. Light. Technol.*, vol. 31, no. 16, pp. 2959–2968, Aug. 2013.

[156] D. R. Selviah, F. A. Fernández, I. Papakonstantinou, K. Wang, H. Bagshiahi, A. C. Walker, A. McCarthy, H. Suyal, D. a. Hutt, P. P. Conway, J. Chappell, S. S. Zakariyah, and D. Milward, "Integrated optical and electronic interconnect printed circuit board manufacturing," *Circuit World*, vol. 34, no. 2, pp. 21–26, 2008.

[157] D. J. Thomson, F. Y. Gardes, J. M. Fedeli, S. Zlatanovic, Y. Hu, B. P. P. Kuo, E. Myslivets, N. Alic, S. Radic, G. Z. Mashanovich, and G. T. Reed, "50-Gb/s Silicon Optical Modulator," *IEEE Photonics Technology Letters*, vol. 24, no. 4. pp. 234–236, 2012.

[158] Q. Li, D. Nikolova, D. M. Calhoun, Y. Liu, R. Ding, T. Baehr-Jones, M. Hochberg, and K. Bergman, "Single Microring-Based 2x2 Silicon Photonic Crossbar Switches," *IEEE Photonics Technology Letters*, vol. 27, no. 18. pp. 1981–1984, 2015.

[159] X. Zheng, F. Liu, J. Lexau, D. Patil, G. Li, Y. Luo, H. D. Thacker, I. Shubin, J. Yao, K. Raj, R. Ho, J. E. Cunningham, and A. V. Krishnamoorthy, "Ultralow power 80 Gb/s arrayed CMOS silicon photonic transceivers for WDM optical links," *J. Light. Technol.*, vol. 30, no. 4, pp. 641–650, Feb. 2012.

[160] P. De Dobbelaere, S. Abdalla, S. Gloeckner, M. Mack, G. Masini, A. Mekis, T. Pinguet, S. Sahni, D. Guckenberger, M. Harrison, and A. Narasimha, "Si photonics based high-speed optical transceivers," *2012 38th European Conference and Exhibition on Optical Communications*. pp. 1–3, 2012.

[161] L. Brusberg, H. Schröder, C. Herbst, C. Frey, C. Fiebig, A. Zakharian, S. Kuchinsky, X. Liu, D. Fortusini, and A. Evans, "High Performance Ion-Exchanged Integrated Waveguides in Thin Glass for Board-Level Multimode Optical Interconnects," in *Proc. of European Conference on Optical Communication (ECOC),* 2015.

[162] L. Brusberg, N. Schlepple, and H. Schroder, "Chip-to-chip Communication by Optical Routing Inside a Thin Glass Substrate," *2011 Ieee 61st Electron. Components Technol. Conf.*, pp. 805–812, 2011.

[163] L. Brusberg, H. Schroeder, M. Queisser, and K. D. Lang, "Single-mode glass waveguide platform for DWDM chip-to-chip interconnects," *2012 IEEE 62nd Electronic Components and Technology Conference*. pp. 1532–1539, 2012.

[164] R. M. Dorward, K. Symington, L. Brusberg, J. R. Kropp, A. Miller, R. Pitwon, and S. Whalley, "Market drivers and architectural requirements for backplane inter-connect capacities in next generation PON head-end equipment in the access network," *Int. Conf. Transparent Opt. Networks*, vol. 7, pp. 1–4, 2013.

[165] H. Schröder, L. Brusberg, R. Pitwon, S. Whalley, K. Wang, A. Miller, C. Herbst, D. Weber, and K.-D. Lang, "Electro-optical backplane demonstrator with integrated multimode gradient-index thin glass waveguide panel," in *Proc. SPIE*, 2015, vol. 9368, p. 93680U–93680U–13.

[166] "Piano plus home page." [Online]. Available: http://www.pianoplus.eu/. [Accessed: 11-Mar-2014].

[167] R. Pitwon, H. Schröder, L. Brusberg, J. Graham-Jones, and K. Wang, "Embedded planar glass waveguide optical interconnect for data centre applications," in *OPTOELECTRONIC INTERCONNECTS XIII, Proceedings of SPIE*, 2013, vol. 8630, p. 86300Z.

[168] C. Hermanns and J. Middleton, "Laser separation of flat glass in electronic-, optic-, display- and bio-industry," *Phot. Process. Microelectron. Photonics IV*, vol. 5713, pp. 387–396, 2005.

[169] H. Schröder, M. Neitz, S. Whalley, C. Herbst, and C. Frey, "Multi-layer Electro-Optical Circuit Board Fabrication on Large Panel," in *Electronic Components and Technology Conference (ECTC), 2016 IEEE*

*66th*, 2016.

[170] R. Pitwon and K. Hopkins, "Optical connector, a communication system and a method of connecting a user circuit to an optical transceiver," US Patent 7,625,134, 2006.

[171] R. Pitwon, I. Johnson, D. Selviah, and I. Papakonstantinou, "Optical printed circuit board and manufacturing method," US Patent 7,936,953, 2005.

[172] R. Pitwon, "Optical circuit board, an optical backplane and an optical communication system," US Patent 7,454,097, 2006.

[173] R. Pitwon, "Optical printed circuit board blank, a kit and a method of making an optical printed circuit board," US Patent 7,422,374, 2006.

[174] R. Pitwon, "Optical printed circuit board, a method of making an optical printed circuit board and an optical waveguide," US Patent 8,731,343.

[175] R. Pitwon, "Optical wavelength division multiplexed multiplexer/demultiplexer for an optical printed circuit board and a method of manufacturing the same," US Patent 7,805,033, 2006.

[176] R. Pitwon, D. Selviah, and I. Papakonstantinou, "Optical wavelength division multiplexed multiplexer/demultiplexer for an optical printed circuit board and a method of manufacturing the same," US Patent 8,007,965.

[177] R. Pitwon, "Adapter for an optical printed circuit board, an optical printed circuit board and a method of connecting an adapter to an optical printed circuit board," US Patent 7,490,993, 2007.

[178] R. Pitwon, "Electro-optical printed circuit board, a blank and a method of making an electro-optical printed circuit board," US Patent 7,899,278, 2007.

[179] R. Pitwon, "Method of making a waveguide and a waveguide made thereby," US Patent 9,333,717, 2008.

[180] R. Pitwon, "An optical connector and a method of connecting a user circuit to an optical printed circuit board," US Patent Application 12/722908, 2010.

[181] R. Pitwon, "An optical connector and a method of connecting an optical connector to an optical printed circuit board," US Patent 8,306,374, 2010.

[182] R. Pitwon, "An interconnect for a data storage system," US Patent 8,861,975, 2010.

[183] K. Hopkins and R. Pitwon, "A data storage system, a modular printed circuit board, a backplane and a

backplane component," US Patent 8,417,071, 2010.

[184] R. Pitwon, "An amplification module for an optical printed circuit board and an optical printed circuit board," US Patent 9,325,146, 2010.

[185] R. Pitwon, "An optical PCB and a method of making an optical PCB," US Patent 8,488,920, 2010.

[186] K. Schmidtke, F. Flens, A. Worrall, R. Pitwon, F. Betschon, T. Lamprecht, and R. Kraehenbuehl, "960 Gb/s optical backplane using embedded polymer waveguides and demonstration in a 12G SAS storage array," *2013 Optical Interconnects Conference*. pp. 29–30, 2013.

[187] Y. Yang, J. Lee, K. Reichard, P. Ruffin, F. Liang, D. Ditto, and S. Yin, "Fabrication and implementation of a multi-to-single mode converter based on a tapered multimode fiber," *Opt. Commun.*, vol. 249, no. 1–3, pp. 129–137, 2005.

[188] S. A. Ibrahim, N. A. M. A. Hambali, K. D. Dambul, and S. Syahali, "Mode power evolution in tapered single-mode fiber," *Telecommun. Technol. (ISTT), 2012 Int. Symp.*, pp. 38–41, 2012.

[189] "International Electrotechnical Commission - Who we are." [Online]. Available: http://www.iec.ch/about/profile/. [Accessed: 07-Jul-2015].

[190] "International Electrotechnical Commission TC 86 Fibre optics JWG 9 Convenor & Members." [Online]. Available: http://www.iec.ch/dyn/www/f?p=103:14:0::::FSP_ORG_ID,FSP_LANG_ID:2775,25. [Accessed: 07-Jul-2015].

[191] I. Papakonstantinou, D. R. Selviah, K. Wang, R. Pitwon, K. Hopkins, and D. Milward, "Optical 8-channel 10 Gb/s MT pluggable connector alignment technology for precision coupling of laser and photodiode arrays to polymer waveguide arrays for optical board-to-board interconnects," in *Electronic Components and Technology Conference*, 2008, no. June 2015, pp. 1769–1775.

# 9 APPENDIX - FIRSTLIGHT PLATFORM FIRMWARE CODING

## *9.1 PCI command structure*

The author wrote the firmware for the StorConnTest3 FPGA including the complete PCI command structure for the FirstLight demonstration platform allowing the user to configure or read-back the values on the FPGA registers of a given StorConnTest3 daughtercard connected to the StorConnOpt3 backplane for general static control lines, StorConn3 transceiver registers, XFP control registers, LED control lines and the crosspoint switch.

### 9.1.1 General static control lines

**Table 9-1: PCI command map - PCI write and read command set for general static control status lines**

| **Write Format** | | | | |
|---|---|---|---|---|
| PCI Address - complete 32 bit address | | | hex “80000800” | |
| PCI Address BAR 2 (space 0) offset | | | 800 *(e.g. el s0+800 Data)* | |
| | | | | |
| **Data Bit** | **Function** | **Section** | **Description** | **User Input Method** |
| 32 - 23 | *not used* | *not used* | *not used* | *not used* |
| 22 | SC3_RXEN | SC3 Optical Receiver window | Global enable bit for Quad Optical Receiver on SC3 transceiver card<br>‘0’ = Disable<br>‘1’ = Enable | Button |
| 21 | SC3_SQEN | SC3 Optical Receiver window | Squelch enable bit for Quad Optical Receiver<br>‘0’ = Disable<br>‘1’ = Enable | Button |

| | | | | |
|---|---|---|---|---|
| 20 | SC3_VDEN | SC3 VCSEL Driver Window | Global enable bit for Quad VCSEL Driver | Button |
| 19 | SC3_Ctrl1 | *Reserved* | Spare control line to SC3 reserved for future functionality | *not used* |
| 18 | SC3_Ctrl2 | *Reserved* | Spare control line to SC3 reserved for future functionality | *not used* |
| 17 - 12 | *not used* | *not used* | *not used* | *not used* |
| 11 | XFP_TX_DIS3 | XFP interface Window | Disable bit for Tx laser of XFP 3 (right most XFP)<br>‘0’ = Tx Laser Enabled<br>‘1’ = Tx Laser Disabled | Button or Schematic Display |
| 10 | XFP_TX_DIS2 | XFP interface Window | Disable bit for Tx laser of XFP 2<br>‘0’ = Tx Laser Enabled<br>‘1’ = Tx Laser Disabled | Button or Schematic Display |
| 9 | XFP_TX_DIS1 | XFP interface Window | Disable bit for Tx laser of XFP 1<br>‘0’ = Tx Laser Enabled<br>‘1’ = Tx Laser Disabled | Button or Schematic Display |
| 8 | XFP_TX_DIS0 | XFP interface Window | Disable bit for Tx laser of XFP 0 (left most XFP)<br>‘0’ = Tx Laser Enabled<br>‘1’ = Tx Laser Disabled | Button or Schematic Display |
| | | | | |
| 7 | XFP_PDn_Rst3 | XFP interface Window | Power Down bit for XFP 3 (right most XFP)<br>‘0’ = Power up<br>‘1’ = Power down<br>Reset on transition from ‘1’ to ‘0’ | Button or Schematic Display |
| 6 | XFP_PDn_Rst2 | XFP interface Window | Power Down bit for XFP 2<br>‘0’ = Power up<br>‘1’ = Power down | Button or Schematic Display |

| | | | | |
|---|---|---|---|---|
| | | | Reset on transition from ‘1’ to ‘0’ | |
| 5 | XFP_PDn_Rst1 | XFP interface Window | Power Down bit for XFP 1<br>‘0’ = Power up<br>‘1’ = Power down<br>Reset on transition from ‘1’ to ‘0’ | Button or Schematic Display |

| | | | | |
|---|---|---|---|---|
| 4 | XFP_PDn_Rst0 | XFP interface Window | Power Down bit for XFP 0 (left most XFP)<br>‘0’ = Power up<br>‘1’ = Power down<br>Reset on transition from ‘1’ to ‘0’ | Button or Schematic Display |
| 3 - 1 | *not used* | *not used* | *not used* | *not used* |
| 0 | XFP_SCLK | XFP interface Window | Clocking rate select for all XFPs between 161.1328 MHz (10 GbE LAN) and 155.52 MHz (10 GbE WAN)<br>‘0’ = 161.1328 MHz<br>‘1’ = 155.52 MHz | Button |
| | | | | |
| **Read Format** | | | | |
| PCI Address - complete 32 bit address | | | hex “80000800” | |
| PCI Address BAR 2 (space 0) offset | | | 800 *(e.g. dl s0+800 4)* | |

**Table 9-2: PCI command sequence example - general IO PCI commands**

| **Write** | | | | | |
|---|---|---|---|---|---|
| | | | | | |
| **Command** | **PCI offset** | | | | |
| | el s0 + 800 | | | | |
| ***XFP*** | | | | | |
| Power down all XFPs | | XX | X | XX | FF0 |

| Power up all XFPs (reset) | | XX | X | XX | 000 |
|---|---|---|---|---|---|
| | | XX | X | | |
| ***SC3 Transceiver*** | | | | | |
| VCSEL Driver, Optical Receiver and Squelch enabled | | XX | 7 | XX | XX<br>X |
| VCSEL Driver, Optical Receiver and Squelch disabled | | XX | 0 | XX | XX<br>X |
| VCSEL Driver Enabled, Optical receiver and squelch disabled | | XX | 1 | XX | XX<br>X |
| VCSEL Driver disabled, Optical receiver and squelch enabled | | XX | 6 | XX | XX<br>X |
| | | | | | |
| **Read from general control register** | dl s0 + 800 4 | | | | |

## 9.1.2 XFP I$^2$C interface (XFPIO)

### *9.1.2.1. XFP I$^2$C read*

In order for the user to read data from an XFP register, two PLX API calls are required:

- PCI Write to execute the XFP read instruction and provide the parameters (internal XFP address)
- PCI read to read back from the FPGA the result of the previously instructed read operation

**Table 9-3: PCI command sequence - execute read of XFP I²C memory device**

| **PCI Command Seq:** | ***PCI Write*** | ***Offset*** | ***8 hex (32 bit) Data Word*** | | | |
|---|---|---|---|---|---|---|
| | el | s0+1900 | XFP sel | I²C Address | Data (R/W) | Proc Register (R/O) |
| XFP 3 (right most) | | | 7F | 00 - FF | XX | XX |
| XFP 2 | | | BF | 00 - FF | XX | XX |
| XFP 1 | | | DF | 00 - FF | XX | XX |
| XFP 0 (left most) | | | EF | 00 - FF | XX | XX |

This will command the FPGA to execute a read sequence on the selected XFP. The contents will be read back to the FPGA and stored at address: 80001c00 (or offset s0+1C00)

**Table 9-4: PCI command sequence - read back contents of previous read instruction**

| **PCI Command Seq:** | ***PCI Read*** | ***Offset*** | ***Specify number of Bytes to read back*** | | |
|---|---|---|---|---|---|
| | dl | s0+1C00 | 4 | | |
| | | | | | |
| **PCI Reply:** | *XFP Select* | *XFP Diagnostic 1* | *Read Address* | *Read back Data* | *XFP Diagnostic 2* |
| **Example:** | 7 (XFP3) | X | 00 | 06 | XX |

The contents of FPGA address 80001c00 (or offset s0+1C00) will be transferred to the user interface.

#### 9.1.2.1. *XFP $I^2C$ write*

In order for the user to write data to an XFP register, one PLX API call is required:

- PCI Write to execute the XFP write instruction and provide the parameters (internal XFP address and data to be written to that address)

**Table 9-5: PCI command sequence - execute write to XFP I²C memory device**

| PCI Command Seq: | *PCI Write* | *Offset* | *8 hex (32 bit) Data Word* | | | |
|---|---|---|---|---|---|---|
| | el | s0+1800 | XFP sel | I²C Address | Data (R/W) | Proc Register (R/O) |
| XFP 3 | | | 7F | 00 - FF | XX | XX |
| XFP 2 | | | BF | 00 - FF | XX | XX |
| XFP 1 | | | DF | 00 - FF | XX | XX |
| XFP 0 | | | EF | 00 - FF | XX | XX |

### 9.1.3 Crosspoint switch interface (CPSIO)

#### 9.1.3.1. *CPS read and write to configure switch*

The CPS is configured by the FPGA across a proprietary Vitesse parallel interface, by writing data to CPS registers.

In order for the user to write data to a CPS register, in order to, for instance, configure a given switch output to a given switch input, one PLX API call is required:

- PCI write to execute the CPS write instruction and provide the parameters (CPS output to input)

Optionally, a PCI read can be made to the same address the write was made, in order to confirm that the write instruction was correctly received by the FPGA

**Table 9-6: PCI command sequence - CPS write sequence to configure switch**

| | PCI offset write | CPS Address | CPS Data | Execute |
|---|---|---|---|---|

| Write to CPS | el s0 + 100 | 0 | Output 0 to 7 | Input 0 to 7 | XXXX1 |
|---|---|---|---|---|---|
| | | | | | |
| Example | | | | | |
| Map CPS output 3 to input 5 | el s0 + 100 | 0 | 3 | 5 | 00001 |
| | | | | | |
| Read from Write register | dl s0 4 | | | | |
| | | | | | |
| Example | | | | | |
| Read last write instruction | dl s0 4 | | | | |
| Expected | 0350XXXX | | | | |

In order for the user to read data from a CPS register, two PLX API calls are required:

- PCI write to execute the CPS read instruction and provide the parameters (CPS Address)
- PCI read to read back from the FPGA the result of the previously instructed read operation

**Table 9-7: PCI command sequence - CPS read sequence to read switch configuration**

| Execute read from CPS | el s0 + 200 | 0 | CPS Register (output) to read | X | XXXX1 |
|---|---|---|---|---|---|
| Example | | | | | |
| Read back input to which output 3 is mapped | el s0+ 200 | 0 | 3 | 0 | 00001 |
| Expected | | | | | |
| | | | | | |
| Read from Read register | dl s0 + 400 4 | | | | |

### 9.1.3.2. *Crosspoint Switch (CPS) output signal conditioning configuration*

The CPS also allows pre-emphasis of all its output ports. The range of output functionality is as follows:

- The following functionalities can be applied either globally or to individual ports
  - Force all outputs to 1, 0 or normal operation
  - Output power **nominal** or **high**

- ○ **Pre-emphasis** enabled or disabled
- ○ **Pre-emphasis adjustment** for varying line lengths between range 0 to 15 (450 ps to 700 ps respectively)
- **Boost mode** enabled / disabled for each channel

**Table 9-8: PCI command sequence - CPS Output pre-emphasis configuration**

| | PCI offset write | CPS Address | CPS Data | Execute |
|---|---|---|---|---|
| **Global Pre-emphasis Setting** | el s0 + 100 | 08 | Range 0 to F (450 ps to 700 ps) | XXXX1 |
| | | | | |
| **Example** | | | | |
| Set Global pre-emphasis to 700 ps | el s0 + 100 | 08 | F | 00001 |
| Set Global pre-emphasis to 450 ps | el s0 + 100 | 08 | 0 | 00001 |
| | | | | |
| **Individual Pre-emphasis Setting** | el s0 + 100 | 10 – 17 (ports 0 to 8) | Range 0 to F (450 ps to 700 ps) | XXXX1 |
| Set Global pre-emphasis on port 3 to 700 ps | el s0 + 100 | 13 | F | 00001 |
| Set Global pre-emphasis on port 0 to 450 ps | el s0 + 100 | 10 | 0 | 00001 |
| Set Global pre-emphasis on port 8 to 700 ps | el s0 + 100 | 18 | F | 00001 |
| Set Global pre-emphasis on port 8 to 573 ps (mid point) | el s0 + 100 | 18 | 7 | 00001 |

**Table 9-9: PCI command sequence - CPS global output level configuration**

| | offset | Address | Data | | | | ex. |
|---|---|---|---|---|---|---|---|
| | | | bit 3 | bit 2 | bit 1 | bit 0 | |
| **Global Output Configuration** | el s0 + 100 | 09 | 00 = normal<br>01 = outputs to 0<br>10 = outputs to 1<br>11 = outputs to 0 | | 0 = nominal output level<br>1 = high output level | 0 = pre-emph. disabled<br>1 = pre-emph. enabled | XXXX 1 |
| **Example** | | | | | | | |
| Set global outputs to normal, high level with pre-emph. enabled | el s0 + 100 | 09 | 3 | | | | 00001 |
| | | | *0* | *0* | *1* | *1* | |
| Force global outputs to 1, nominal level, no pre-emph. | el s0 + 100 | 09 | 8 | | | | 00001 |
| | | | *1* | *0* | *0* | *0* | |

**Table 9-10: PCI command sequence - CPS individual output level configuration**

| | offset | Address | Data | | | | ex. |
|---|---|---|---|---|---|---|---|
| | | | bit 3 | bit 2 | bit 1 | bit 0 | |
| **Individual Output Configuration** | el s0 + 100 | 18 – 1F | 00 = normal<br>01 = outputs to 0<br>10 = outputs to 1<br>11 = outputs to 0 | | 0 = nominal output level<br>1 = high output level | 0 = pre-emph. disabled<br>1 = pre-emph. enabled | XXXX 1 |
| **Example** | | | | | | | |
| Set output 0 to | el s0 + | 18 | 3 | | | | 00001 |

| normal, high level with pre-emph. enabled | 100 | | *0* | *0* | *1* | *1* | |
|---|---|---|---|---|---|---|---|
| Force output 7 to 1, nominal level, no pre-emph. | el s0 + 100 | 1F | 8 | | | | 00001 |
| | | | *1* | *0* | *0* | *0* | |

**Table 9-11: PCI command sequence - CPS output boost mode configuration**

| | offset | Add. | Data | | | | ex. |
|---|---|---|---|---|---|---|---|
| | | | bit 3 | bit 2 | bit 1 | bit 0 | |
| **Individual Output Configuration <u>Outputs 0 - 3</u>** | el s0 + 100 | 0D | output 3<br>0 = boost off<br>1 = boost on | output 2<br>0 = boost off<br>1 = boost on | output 1<br>0 = boost off<br>1 = boost on | output 0<br>0 = boost off<br>1 = boost on | XX<br>XX<br>1 |
| **Individual Output Configuration <u>Outputs 0 - 3</u>** | el s0 + 100 | 0E | output 7<br>0 = boost off<br>1 = boost on | output 6<br>0 = boost off<br>1 = boost on | output 5<br>0 = boost off<br>1 = boost on | output 4<br>0 = boost off<br>1 = boost on | XX<br>XX<br>1 |

### *9.1.3.3. CPS Input Signal Equalisation Configuration*

The CPS also allows equalisation of all its input ports. The range of functionality is as follows:

Global and individual equalisation level on the inputs between four settings:

- No equalisation
- Minimum equalisation
- Medium equalisation
- Maximum equalisation

**Table 9-12: PCI command sequence - CPS input equalisation configuration**

| | PCI offset write | CPS Address | CPS Data | Execute |
|---|---|---|---|---|
| **Global Equalisation Setting** | el s0 + 100 | 0A | X 000 = No eq<br>001 = Min eq<br>011 = Med eq<br>111 = Max eq | XXXX1 |
| | | | | |
| **Example** | | | | |
| Set Global equalisation to no equalisation | el s0 + 100 | 0A | 0 (0000) | 00001 |
| Set Global equalisation to maximum equalisation | el s0 + 100 | 0A | 7 (0111) | 00001 |
| | | | | |
| **Individual Equalisation Setting** | el s0 + 100 | 20 - 27 (ports 0 to 8) | X 000 = No eq<br>001 = Min eq<br>011 = Med eq<br>111 = Max eq | XXXX1 |
| Set equalisation on port 3 to medium | el s0 + 100 | 23 | 3 (0011) | 00001 |
| Set equalisation on port 0 to no equalisation | el s0 + 100 | 10 | 0 (0000) | 00001 |
| Set equalisation on port 8 to maximum | el s0 + 100 | 18 | 7 (0111) | 00001 |
| Set equalisation on port 3 to minimum | el s0 + 100 | 18 | 1 (0001) | 00001 |

### 9.1.4 StorConn3 transceiver interface (SC3IO)

#### *9.1.4.1. Configurable StorConn3 parameters:*

- VCSEL global functionality

    - Waveform control
    - Temperature compensation for bias current (Imax)
    - Temperature compensation for modulation current (Imod)
- VCSEL per channel functionality
    - Enable
    - Bias current
    - Modulation current
    - *Force output high (not displayed in example GUI)*
- Optical receiver global functionality
    - Global signal detect sensitivity
    - Signal detect threshold
    - Hysteresis on/off
    - Optical output swing (nominal/high)
- Optical receiver per channel functionality
    - Enable
- VCSEL temperature read-back
- VCSEL drive current read-back

#### *9.1.4.2. SC3 IO test sequence examples*

**Table 9-13: PCI command sequence - execute I²C write of VCSEL Driver to turn channel 4 on**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Write VCSEL block 1 | el s0+1000 | A0 | 0C | 29 | 79 |
| Write VCSEL block 2 | el s0+1020 | 02 | 97 | 90 | 29 |
| Write VCSEL block 3 | el s0+1040 | 79 | 02 | 97 | A0 |
| Write VCSEL block 4 | el s0+1060 | 00 | 00 | 00 | 00 |
| | | | | | |
| Execute VCSEL Write | el s0+10E0 | FF | FF | FF | FF |
| Check that Write Done flag (bit 0) = ‘1’ | dl s0+ 1780 4 | XX | XX | XX | X1 |

**Table 9-14: PCI command sequence - execute I²C write of VCSEL Driver to turn all VCSELs off**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Write VCSEL block 1 | el s0+1000 | A0 | 0C | 29 | 79 |
| Write VCSEL block 2 | el s0+1020 | 02 | 97 | 90 | 29 |
| Write VCSEL block 3 | el s0+1040 | 79 | 02 | 97 | 90 |
| Write VCSEL block 4 | el s0+1060 | 00 | 00 | 00 | 00 |
| | | | | | |
| Execute VCSEL Write | el s0+10E0 | FF | FF | FF | FF |
| Check that Write Done flag (bit 0) = ‘1’ | dl s0+ 1780 4 | XX | XX | XX | X1 |

**Table 9-15: PCI command sequence - execute I²C write of VCSEL driver to turn all VCSELs on**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Write VCSEL block 1 | el s0+1000 | A0 | 0C | 29 | 7A |
| Write VCSEL block 2 | el s0+1020 | 02 | 97 | A0 | 29 |
| Write VCSEL block 3 | el s0+1040 | 7A | 02 | 97 | A0 |
| Write VCSEL block 4 | el s0+1060 | 00 | 00 | 00 | 00 |
| | | | | | |
| Execute VCSEL Write | el s0+10E0 | FF | FF | FF | FF |
| Check that Write Done flag (bit 0) = ‘1’ | dl s0+ 1780 4 | XX | XX | XX | X1 |

Read diagnostic status of VD write cycle: dl s0+1600 4

**Table 9-16: PCI command sequence - execute I²C read of VCSEL Driver**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Execute VCSEL Read | el s0+1080 | FF | FF | FF | FF |
| Check that Read Done flag (bit 0) = ‘1’ | dl s0+ 1780 4 | XX | XX | XX | X1 |

| | | | | | |
|---|---|---|---|---|---|
| | | | | | |
| Read VCSEL block 1 | dl s0+1720 4 | Block 1 | Block 2 | Block 3 | Block 4 |
| **Expected Reply** | | A0 | 0C | 29 | 7A |
| Read VCSEL block 2 | dl s0+1740 4 | Block 5 | Block 6 | Block 7 | Block 8 |
| **Expected Reply** | | 02 | 97 | A0 | 29 |
| Read VCSEL block 3 | dl s0+1760 4 | Block 9 | Block 10 | Block 11 | Block 12 |
| **Expected Reply** | | 7A | 02 | 97 | A0 |
| Read VCSEL block 4 | dl s0+1780 4 | Block 13 | XX | XX | XX |
| **Expected Reply** | | 00 | 00 | 00 | 00 |
| | | | | | |

Read diagnostic status of VD read cycle: dl s0+1620 4

**Table 9-17: PCI command sequence - execute I²C write of optical driver (TIA) to turn all photodiode channels on**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Write RA block 1 | el s0+1100 | A5 | 9B | 00 | 00 |
| Write RA block 2 | el s0+1120 | 00 | XX | XX | XX |
| | | | | | |
| Execute RA Write | el s0+11E0 | FF | FF | FF | FF |
| Check that Write Done flag (bit 0) = ‘1’ | dl s0+ 17C0 4 | XX | XX | XX | X**1** |

**Table 9-18: PCI command sequence - execute I²C write of optical driver (TIA) to turn all photodiode channels off**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Write RA block 1 | el s0+1100 | A5 | 9B | 22 | 22 |
| Write RA block 2 | el s0+1120 | 00 | XX | XX | XX |
| | | | | | |
| Execute RA Write | el s0+11E0 | FF | FF | FF | FF |
| Check that Write Done flag (bit 0) = '1' | dl s0+ 17C0 4 | XX | XX | XX | X**1** |

Read diagnostic status of RA write cycle: dl s0+1680 4

**Table 9-19: PCI command sequence - execute I²C read of optical receiver**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Execute VCSEL Read | el s0+1200 | FF | FF | FF | FF |
| Check that Read Done flag (bit 0) = '1' | dl s0+ 17C0 4 | XX | XX | XX | X**1** |
| | | | | | |
| Read RA block 1 | dl s0+17A0 4 | Block 1 | Block 2 | Block 3 | Block 4 |
| **Expected Reply (to case six above)** | | A5 | 9B | 22 | 22 |
| Read RA block 2 | dl s0+17C0 4 | Block 5 | XX | XX | XX |
| **Expected Reply (to case six above)** | | 00 | 00 | 00 | 00 |

Read diagnostic status of RA read cycle: dl s0+1660 4

**Table 9-20: PCI command sequence - Execute I²C read of VCSEL driver temperature**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Execute VCSEL Temperature Read | el s0+1260 | FF | FF | FF | FF |
| Check that Read Done flag | dl s0+ 17C0 4 | XX | XX | XX | X**1** |

| (bit 0) = '1' | | | | | |
|---|---|---|---|---|---|
| | | | | | |
| Read temp block | dl s0+1400 4 | Block 1 | Block 2 | XX | XX |

Read diagnostic status of VD Temp read cycle: dl s0+1660 4

**Table 9-21: PCI command sequence - execute I²C read of VCSEL driver monitor current**

| **PLX Mon Command:** | offset | 1st byte | 2nd byte | 3rd byte | 4th byte |
|---|---|---|---|---|---|
| Execute VCSEL Current Read | el s0+1280 | FF | FF | FF | FF |
| Check that Read Done flag (bit 0) = '1' | dl s0+ 17C0 4 | XX | XX | XX | X**1** |
| | | | | | |
| Read current block | dl s0+1420 4 | Block 1 | Block 2 | XX | XX |

Read diagnostic status of VD Temp read cycle: dl s0+1660 4

### 9.1.5 LED control interface (LEDIO)

**Table 9-22: PCI command map - LED control interface**

| **Write Format** | | | | |
|---|---|---|---|---|
| PCI Address - complete 32 bit address | | | hex “80000C00” | |
| PCI Address BAR 2 (space 0) offset | | | C00 *(e.g. el s0+C00 Data)* | |
| | | | | |
| **Data Bit** | **Function** | **Section** | **Description** | **User Input Method** |
| 32 - 23 | *not used* | *not used* | *not used* | *not used* |
| 22 - 20 | LED Mode | Diagnostic window | “000” Direct mode in which bits 18 – 1 map directly to LEDs 18 – 1<br>“001” LED automatic flashing mode 1 (default)<br>“010” LED automatic flashing mode 2<br>“011” LED automatic flashing mode 3<br>“100” LED automatic flashing mode 4 | Button |
| 19 | *not used* | *not used* | *not used* | |
| 18 - 1 | LED 18 - 1 | Diagnostic interface window | Enable bit for diagnostic LEDs 18 -1<br>‘0’ = Disable<br>‘1’ = Enable | Button |
| 0 | *not used* | *not used* | *not used* | *not used* |
| | | | | |
| **Read Format** | | | | |
| PCI Address - complete 32 bit address | | | hex “80000C00” | |
| PCI Address BAR 2 (space 0) offset | | | C00 *(e.g. dl s0+C00 4)* | |