Grant Agreement No. 619572

COSIGN

Combining Optics and SDN In next Generation data centre Networks

Programme: Information and Communication Technologies

Funding scheme: Collaborative Project – Large-Scale Integrating Project

Deliverable D5.1 – Demonstrator Results of Data Plane Integration of Different Switch and Fibre Solutions

Due date of deliverable: March 31 2017
Actual submission date: May 23 2017

Start date of project: January 1, 2014 Duration: 39 months

Lead contractor for this deliverable: DTU

Project co-funded by the European Commission within the Seventh Framework Programme		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Executive Summary

The COSIGN project has produced a number of unique innovations in optical technologies for datacentre networks. Below we will outline the main achievements within the tree major topics of fibre technology development carried out in the COSIGN project.

Multi-core fibre

Through a coordinated and persistent collaboration between several project partners key elements for a multicore fibre platform for data networks has been developed. A novel 4-core fibre concept has been developed and two different fibre types have been produced to satisfy the different requirements of switch-integration and transmission performance. A prototype of a multicore fibre switch has also been developed allowing the use of 4-core fibre in an optically switched data plane. This constitutes the first demonstration of integration of multi-lane fibre transmission and switching and constitutes a major milestone on the path to practical applications of multi-lane optical data planes.

To complete the multi-core platform, the project has initiated development glass interposers to interface the arrays of VCSEL light sources and photodiodes used in transceivers to the cores of a multi-core fibre. Once this technology is fully developed, a full multi-core fibre data plane can be realised.

Few-mode fibre

As an alternative or potentially a supplement to multi-core technology, the COSIGN project has also investigated mode multiplexing as an optical technology for multi-lane transmission. Appropriate optical fibres (few-mode fibres) have been identified and characterised for multi-lane transmission confirming the potential of the technology. Finally, integrated silicon devices for multiplexing and de-multiplexing of modes have been developed and combined with few-mode fibre to create an entire transmission link based on the mode-multiplexing technology. This constitutes the first demonstration of a mode-multiplexed link based on integrated devices compatible with cheap mass production. As such, it is a key step in demonstrating the practical potential of the technology.

Hollow-core fibre

Low latency optical transmission has been investigated using hollow-core fibre technology. Specifically, the combination of optically time-shared connections and low latency fibre has been tested as a way of reducing the impact of the increased latency associated with temporal partitioning of optical connections. This is the first investigation of its kind and has mapped out the trade-offs between latency contributions from different sources for different link configurations. It is clearly indicated that as the latency associated with sharing of a link is brought down by e.g. faster transceiver circuits, the benefit from using hollow-core fibre for transmission will increase.

The details of these results are reported in this document.

Document Information

Status and Version:	Final version	
Date of Issue:	23/05/2015	
Dissemination level:	Public	
Author(s):	**Name**	**Partner**
	Michael Galili	DTU
	Hans Christian Mulvad	UNISOUTH
	Martynas Beresna	UNISOUTH
	Tim Lee	UNISOUTH
	Valerija Kamchevska	DTU
	Jan Markus Baumann	DTU
	Chris Jackson	UNIVBRIS
Edited by:	Michael Galili	DTU
Reviewed by:	Salvatore Spadaro	UPC
	Giada Landi	NXW
Checked by :	Sarah Ruepp	DTU

Table of Contents

Executive Summary **2**

Table of Contents **5**

1 Introduction **6**

1.1 Reference Material 6
1.1.1 Reference Documents 6
1.1.2 Acronyms and Abbreviations 6

1.2 Document History 7

2 Multilane switch with integrated MCFs **8**

2.1 Performance of the switch 8

2.2 Prospects of scaling the switch to more ports / more cores 10

2.3 Tolerance towards lateral (x, y) misalignments 13

2.4 Summary of multi-core optical switch efforts in COSIGN 14

3 Multi-lane few mode fibre transmission **15**

3.1 Integrated devices for mode multiplexing and de-multiplexing 15

3.2 Transmission performance 16

3.3 Prospects of integrating multi-lane few mode transmission in DCNs 18

4 Laser-written silica interposer for VCSEL-MCF interface **19**

4.1 FPC design for 4 channel 25Gbps 1530nm single-mode(SM) transmitter 19

4.2 Fabrication of the interposer 19

5 Hollow-core fibre (HCF) for low latency DCN links **22**

5.1 Motivation for HCF-based links in DCNs 22

5.2 Description of the demonstration 22

5.3 TDM network behaviour differences with HCFs 22
5.3.1 Throughput Measurements 23
5.3.2 Latency Measurements 24
5.3.3 Improved Timeslot Allocation 25
5.3.4 Summary and future experiments 26

6 Summary and conclusion **27**

1 Introduction

This deliverable reports on the results obtained on the integration of fibre and switch technologies developed in the COSIGN project.

In chapter 2 we describe the development and performance of the multi-core fibre switch. The chapter includes description of the fibre which is specially designed for this purpose as well as details regarding alignment and scaling of the switch.

In chapter 3 we describe the realization of a mode-multiplexed multi-lane transmission link based on integrated silicon multiplexer/de-multiplexer devices. The integrated devices are discussed and the performance of the entire link is reported.

In chapter 4 we describe the efforts on developing glass interposers to interface conventional transmitter/receiver technology to multicore fibre. In the chapter we discuss how such a component could interface linear VCSEL arrays to 2D core geometries as the one used in the multi-core fibre switch.

In chapter 5 we present the results of detailed investigation of the latency improvements achievable by using hollow-core fibre in conjunction with time-division multiplexed optical transmission.

Finally, in chapter 6 we summarise and conclude on the work presented in the document.

1.1 Reference Material

1.1.1 Reference Documents

[1]	COSIGN FP7 Collaborative Project Grant Agreement Annex I - "Description of Work"
[2]	H. C. H. Mulvad, A. Parker, B. King, D. Smith, M. Kovacs, S. Jain, J. R. Hayes, M. Petrovich, D. J. Richardson, and N. Parsons, "Beam-Steering All-Optical Switch for Multi-Core Fibers," in *Optical Fiber Communication Conference*, OSA Technical Digest (online) (Optical Society of America, 2017), paper Tu2C.4.
[3]	D2.3 Report on Development of First Generation MCF and PBGF Fibres
[4]	D2.4 Report on New Parallel Optics Transceiver Modules Based on 2.5D and 3D Integrated Components
[5]	D2.5 Report on Development of Optimised Second Generation Fibres
[6]	D2.7 Final Report on Fibre development
[7]	D2.9 Report on Performance of MCF Designed for Multi-Lane Switching Experiments
[8]	D5.0 Definition, Design and Test Plan for Use Cases
[9]	Y. Ding et al., CLEO'2015, paper STh1F.1
[10]	F. Poletti, N.V. Wheeler, M.N. Petrovich, N. Baddela, E. Numkam Fokoua, J.R. Hayes, D.R. Gray, Z. Li, R. Slavik, D.J. Richardson, "Towards high-capacity fibre-optic communications at the speed of light in vacuum," *Nature Photon.,* vol. 7, no. 4, pp. 279–284, (2013)

1.1.2 Acronyms and Abbreviations

Most frequently used acronyms in the Deliverable are listed below. Additional acronyms may be defined and used throughout the text.

DoW	Description of Work
SDN	Software Defined Networks

MCF	Multi core fibre
MFD	Mode field diameter
FMF	Few mode fibre
MMI	Multi mode interferometer
SOI	Silicon on Insulator
BOX	Buried Oxide
SSMF	Standard single mode fibre
SMF	Single mode fibre
Al	Aluminium
MDE	Mode dependent coupling efficiency
MDM	Mode division multiplexing
MIMO	Multiple input Multiple output
DSP	Digital signal processing
PIC	Photonic integrated circuit
VCSEL	Vertical cavity surface emitting laser
FPC	Flexible printed circuit boards
SM	Single mode
HeNe	Helium Neon
LTD	Long term demonstrator
TDM	Time division multiplexing
HCF	Hollow core fibre
FPGA	Field programmable gated arrays
NIC	Network interface card
OS	Operating system
TCP	Transmission Control Protocol

1.2 Document History

Version	Date	Authors	Comment
01	04/03/2017	See the list of authors	TOC
02	12/05/2017		Integration of final version
03	22/05/2017		Quality check
04	23/05/2017		Final version

2 Multilane switch with integrated MCFs

During the course of the COSIGN project, it was decided to investigate, within WP2, the integration of multicore fibre (MCF) into the Polatis switch. A key aim of the COSIGN project is to enable software-defined dynamic reconfiguration at the physical (fibre) layer. As previously explained in [7], the integration of MCF fibres with an optical circuit switch technology such as the POLATIS switch offers the prospect of routing multilane traffic on a single fibre pair, thereby substantially reducing the cost and control complexity of such dynamically reconfigurable links. Hence, POLATIS and UNISOUTH undertook the development of a multilane switch with integrated MCF fibres. This work included several stages which have been described in detail in previous deliverables: 1) the development of a 4-core MCF suitable for integration within the switch (cf. [3] and [5]); 2) a theoretical and numerical analysis of a telecentric alignment of the switch optics to enable low-loss MCF-to-MCF switching performance followed by initial experimental tests of the configuration (cf. [7]); and, finally 3) the connectorisation of the MCFs with collimators followed by their integration within the switch assembly (cf. [6]). In deliverable D2.7, we reported in detail on the performance of the packaged multi-lane switch prototype with integrated 4-core MCFs. In the following, we summarise the main results from D2.7, including the low-loss performance of a 1x2 switch (presented at OFC 2017 [2]), and a 2x2 switch setup with a 1 km transmission fibre and fan-in/out devices on all MCFs (built for the long-term demonstrator in WP5). Furthermore, we perform some calculations on the scalability of the number of ports when integrating higher core-count MCFs into the switch. Finally, some basic calculations on the lateral misalignment tolerance of the collimator and microlens arrays are included, since the switch performance was found to be very sensitive to this alignment.

2.1 Performance of the switch

An illustration of the switch setup, which has the potential for holding up to 96 ports with 4-core MCFs, is shown in Fig. 2.1 (a). The fabricated multilane switch prototype was packaged in a compact module, shown in Fig. 2.1 (b). The individual cores were assessed by using fan-in/out devices from Optoscribe. The switching loss results are shown in Fig. 2.1 (c). Including the fan-in/outs, the total core-to-core losses are in the range 5-8 dB, which is largely attributed to the fan-in/out device losses (‘full system’). If the fan-in/out losses are subtracted, we obtain the loss values for the switch including the integrated 4-core MCFs. The obtained switch losses are in the range 1.2-2.2 dB for all cross-connections between the 3 ports (labelled P7, P8, P9). These results demonstrate that the telecentric setup can successfully perform 1:1 imaging between input and output MCFs over a large working distance (here 174mm). This working distance potentially allows for integrating up to 96 ports in the switch, as illustrated in Fig. 2.1 (a). The results were reported at the OFC 2017 conference [2], and the detailed characterizations can be found in [7].

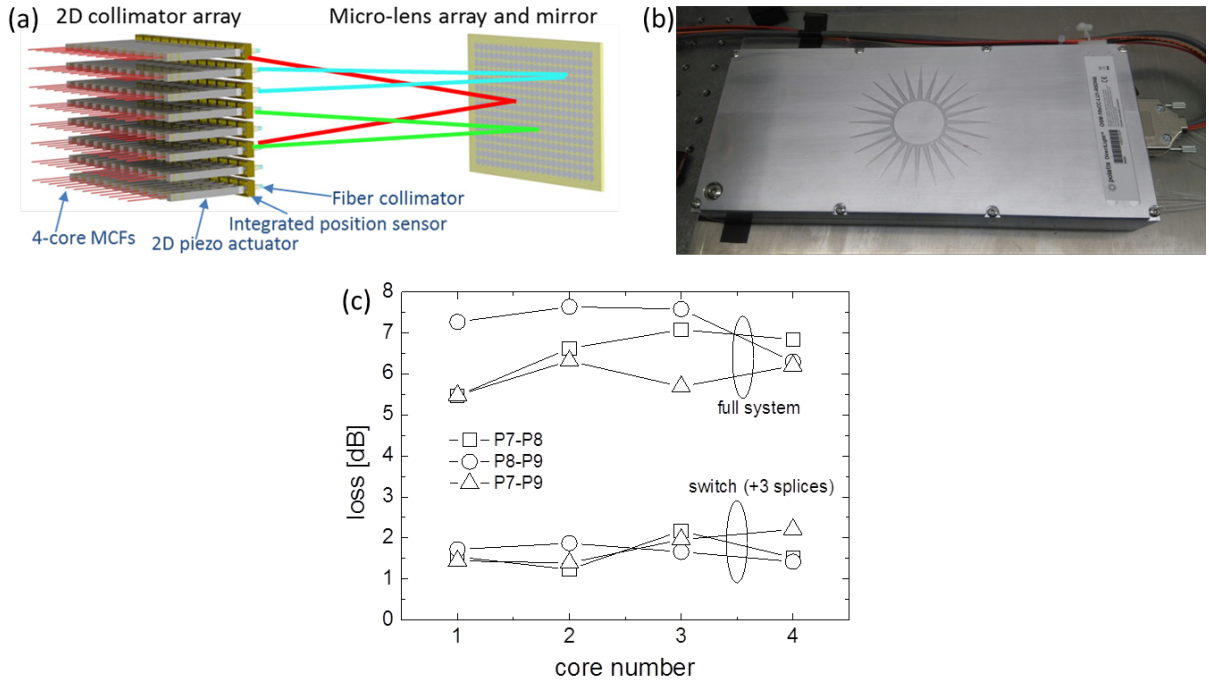

Fig. 2.1. (a) Illustration of the multilane switch configuration, (b) Photograph of the fabricated switch module (c) Switching loss performance between 3 ports with integrated 4-core MCFs (full-system: including fan-in/out devices to assess individual cores, switch: fan-in/out losses subtracted).

We subsequently prepared a 2x2 switch setup including a 1-km long transmission span of the 2nd generation MCF, as shown in Fig. 2.2 (a). Cross-section images of the fabricated 1st and 2nd generation MCFs are shown in Fig. 2.2 (c). The 1st generation MCF was fabricated specifically for integration into the Polatis switch, and included a central 'shunt' core to facilitate collimator alignment. The 2nd generation MCF was fabricated for transmission purposes, and it is based on a trench-assisted index profile preform from OFS that enables a low cross-talk of -54 dB/km. This cross-talk level is sufficiently low for intra data-centre transmission distances. The pitch and MFD value matches the 1st generation MCF, enabling low-loss splicing between the two MCF types. The total losses of the 2x2 switch setup, including fan-in/out devices, are shown in Fig. 2.2 (b). The losses are in the range 5-9 dB. The higher losses obtained for some of the cores in cross-connections involving the port labeled #6, are attributed to a lateral misalignment of the corresponding collimator. The tolerance towards such misalignment is addressed in more detail in section 2.3.

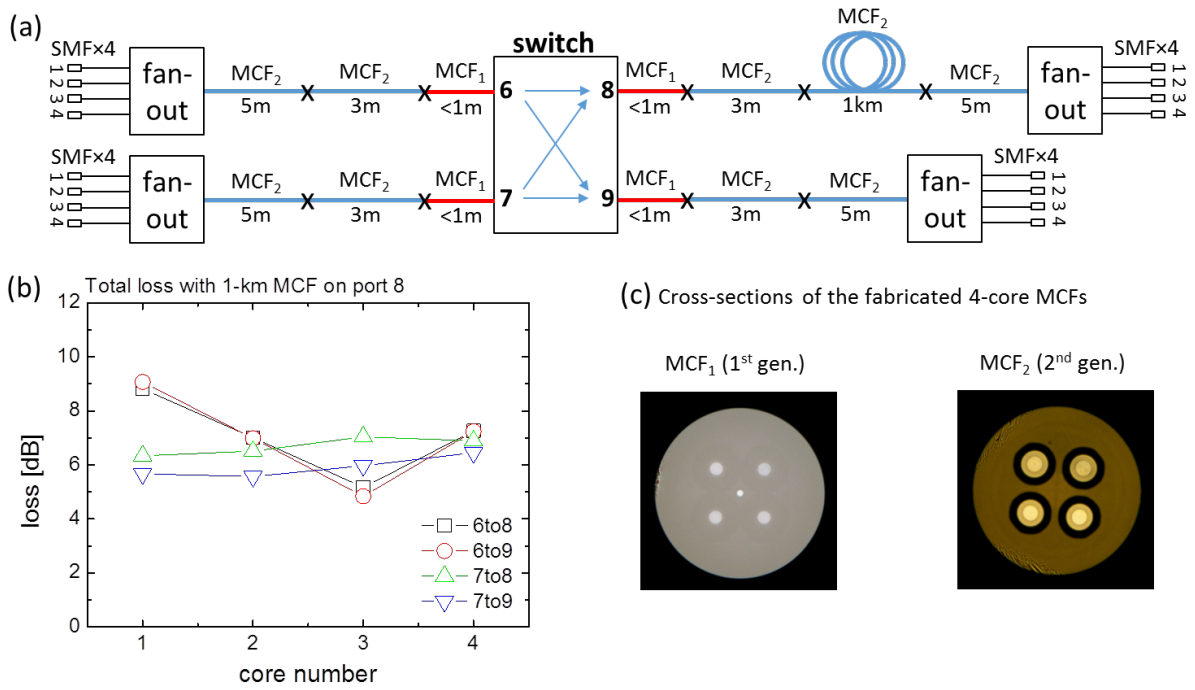

Fig. 2.2. (a) 2x2 multilane switch set-up for the long-term demonstrator, (b) Total core-to-core loss of the full setup, incl. fan-in/out devices, (c) Images of the fabricated 4-core MCFs (1st generation MCF for integration into the switch, 2nd generation MCF for transmission).

2.2 Prospects of scaling the switch to more ports / more cores

The number of ports is determined by the scan range of an ingress collimator across the egress collimator array when varying the pointing direction of the collimator. In the paraxial approximation, the scan range in the x- and y-directions is given by:

$$\Delta_{x,y} = WD \cdot \theta_{x,y},$$

where WD is the working distance and $\theta_{x,y}$ is the angular range of the collimator in the x- and y-directions, respectively. In the telecentric configuration of the multilane switch, the working distance is given by $WD = 2z' = 2(f_1 + 2f_2)$, where f_1 and f_2 are the focal lengths of the collimator and the central micro-lens array lenslets, respectively (cf. [7]).

The number of egress collimators $n_{x,y}$ that can be targeted in the x,y-directions are given by (before rounding to the lower nearest integer):

$N_{x,y} = 1 + \Delta_{x,y}/\delta_{x,y}$, where $\delta_{x,y}$ is the collimator spacing in the x,y-directions.

Since the angular range and spacing of the collimators are fixed, the WD should be maximised by adapting the focal lengths f_1 and f_2. However, adjustments of these focal lengths will also be limited by the apertures of the corresponding lenses, as described in the following.

Microlens aperture limitation. First of all, the possible WD is limited by the aperture of the central microlens. In [7], it was shown that the overall MCF field profile at the microlens (total field diameter D_{lens}) is a magnified image of the MCF output field profile (diameter D_{MCF}) . The relation is given by:

$D_{lens} = |M| \cdot D_{MCF}$, where $M = -\frac{2f_2}{f_1}$ is the magnification factor.

The total field diameter D_{MCF} is determined by the number of cores and the core lattice geometry. If the distance from the collimator to the microlens is too large, the diverging beams from the outmost cores of the MCF will exceed the microlens aperture, leading to distortion and ultimately beam

clipping. If the microlens aperture is given by A_{lens}, we therefore need $D_{lens} \leq A_{lens}$ to avoid beam-clipping, leading to an upper limit M_{max} for the magnification:

$$|M| = \frac{2f_2}{f_1} \leq \frac{A_{lens}}{D_{MCF}} = |M_{max}|.$$

Collimator aperture limitation. A further limitation to the working distance is the collimator lens aperture A_{coll}. This aperture must not be exceeded by the expanding (parallel) output beams of the MCF cores when propagating the distance $z = f_1(1 + f_1/2f_2)$ between the MCF output facet and the collimator lens. This condition is expressed by the following inequality:

$$4z\lambda/\pi MFD + D_{MCF} \leq A_{coll},$$

Here, the first term on the left-hand side is the beam diameter for a single core at the collimator lens, after propagating the distance z (assuming Gaussian beam propagation, well beyond the Rayleigh range), and $MFD{\sim}10\mu m$ is the mode-field diameter of a single core mode in the MCF (at $\lambda = 1550nm$).

Maximum number of ports

The number of ports in each dimension can be rewritten as:

$$N_{x,y} = 1 + 2(f_1 + 2f_2)\,\theta_{x,y}/\delta_{x,y} = 1 + 2f_1\theta_{x,y}(1 + |M|)/\delta_{x,y}$$

An estimate for the maximum number of ports can be derived if the propagating signal is allowed to extend to the full aperture of both the collimator lens and microlens. Using the upper limits for M and z given above leads to solutions for f_1 and f_2. These can then be inserted in the expression for $N_{x,y}$ to obtain:

$$N_{x,y}^{max} = 1 + \frac{\theta_{x,y}\pi MFD}{2\lambda\delta_{x,y}}(A_{coll} - D_{MCF})|M_{max}| = 1 + \frac{\theta_{x,y}\pi MFD \cdot A_{lens}}{2\lambda\delta_{x,y}}\left(\frac{A_{coll}}{D_{MCF}} - 1\right)$$

As expected, the maximum port number in a single dimension scales inversely with the MCF output field diameter D_{MCF}. The total port number is given by $N = N_x^{max} \cdot N_y^{max}$, resulting in an inverse square scaling vs D_{MCF}:

$$N \cdot D_{MCF}^2 \sim constant$$

As function of the number of cores in the MCF (n_c), the value of D_{MCF} will scale approximately as $\sqrt{n_c}$ (at least for large core numbers). This implies as a rule of thumb that the total port count (in two dimensions) will scale inversely with the core count, which can be written as:

$$N \cdot n_c \sim constant$$

Here, the constant will depend on the parameters of the switch, including the geometry, apertures, and angular range. For fixed switch parameters, the maximum number of ports for a given MCF core count is obtained with an optimum core packing geometry that minimizes the diameter D_{MCF}. Here, it is worth noting that a few-mode/multimode fiber would enable a significantly larger number of switch ports, since a much a higher number of spatial channels (modes) can be packed within a given diameter compared to MCF. Furthermore, multimode fibers do not require rotational alignment as does MCF. In terms of optimizing switch parameters, there is limited room for improvement. The microlens aperture A_{lens} is upper-limited by the microlens array pitch, which is half the collimator array pitch ($\delta_{x,y}$). It can also be noted that the minimum collimator array pitch is limited by the angular range $\theta_{x,y}$ of the collimators, since these have a finite length and rotate about the fiber termination. Increasing the collimator aperture A_{coll}, if possible, might be the best approach towards improving the scalability constant.

In the following, we provide some numerical examples of the port scalability using the parameters from the implemented 4-core MCF switch as a starting point. Fig. 2.3 shows the maximum possible number of ports in a single dimension N_x^{max} (in one switch slice) as a function of D_{MCF}.

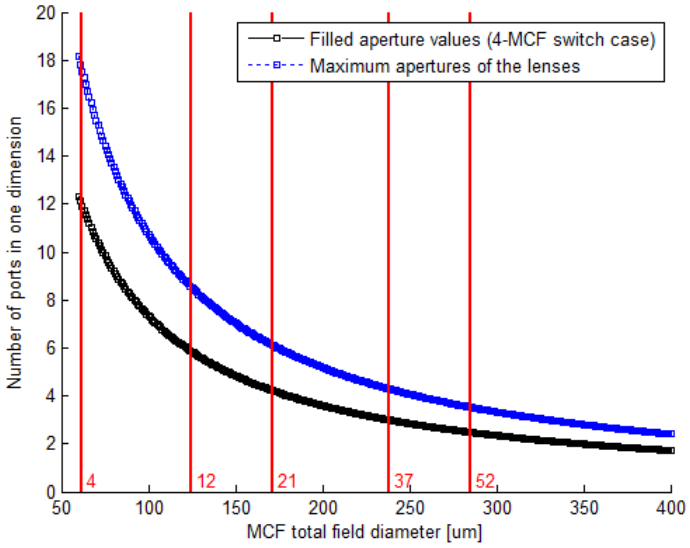

Fig. 2.3. Switch scalability, showing the possible number of ports in one dimension vs. the total MCF field diameter D_{MCF}. The vertical lines indicate the values of D_{MCF} for a number of square-lattice MCFs with different numbers of cores ($n_c = 4, 12, 21, 37, 52$), all illustrated in Fig. 2.4.

One graph (black) shows the port count scaling when using the filled aperture values in the 4-core MCF switch as reference. This means that the actual field diameters at the collimator lens and micro-lens in the 4-core MCF switch case are used as fixed values for the maximum apertures. The other graph (blue) shows the port count when the higher (specified) aperture values of the lenses are used as maxima. The red vertical lines indicate the D_{MCF} values for MCFs with various core-counts, and can be used to estimate to corresponding port count. The MCF core geometries are illustrated in Fig. 2.4. All cases are based on a square lattice geometry with the same pitch and single-mode core MFD as the fabricated 4-core MCFs (pitch = $36\mu m$, $MFD = 10.7\mu m$). As it can be seen in Fig. 2.3, the implemented 4-core MCF switch has capacity for 12 ports in a switch slice. If the maximum lens apertures are explored, the switch slice could in principle hold up to 17 ports. If D_{MCF} is scaled up to the value of a 52-core MCF with the same pitch, the maximum capacity would only be 3 ports.

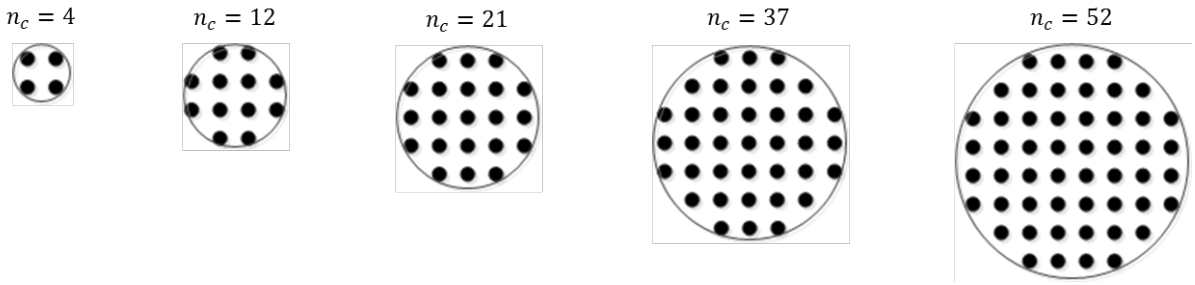

Fig. 2.4: Examples of square lattice MCFs with different core counts. In each case, the enclosing circle indicates the corresponding diameter D_{MCF} that determines the maximum port count.

Table 2.1 shows the port counts for the various MCF cases in both the x- and y-dimensions (note that the collimator spacings δ_x and δ_y are slightly different, resulting in different x/y port counts). Note also that the product $N \cdot n_c$ deviates from a constant due to the integer rounding when obtaining core- and port-numbers, but the general scaling rule $N \cdot D_{MCF}^2 \sim constant$ holds.

Table 2.1: Maximum port counts for various MCFs (core configurations are illustrated in Fig. 2.4)

n_c	D_{MCF} [μm]	N_x^{max}	N_y^{max}	$N = N_x^{max} \cdot N_y^{max}$	$N \cdot n_c$
4	61.6	17	11	187	748
12	124.5	8	5	40	480
21	171.7	6	4	24	504
37	238.4	4	3	12	444
52	284.9	3	2	6	312

2.3 Tolerance towards lateral (x, y) misalignments

During the development of the switch, it was found that the loss performance is very sensitive to the lateral (x,y) alignment of the collimator-array and microlens array. Ideally, the ingress and egress collimator and the microlens should be positioned on a single line for all possible cross-connections, such that the propagating signal always passes through the center of a microlens. Since such ideal alignment is difficult to achieve in practice, it is relevant to investigate the tolerance of the switch geometry to lateral misalignments. A certain tolerance is obtained if the microlens can be used for beam-steering, allowing to re-direct the signal towards a shifted position of the egress collimator relative to the ideal, in-line position. Such beam-steering is possible when the signal diameter D_{lens} at the microlens is smaller than the aperture A_{lens}. Using ray transfer matrix calculations in the thin-lens and paraxial rays approximation, it can be shown that the misalignment tolerance (relative to the collimator spacing) is given by:

$$\Delta\delta_{x,y} = \frac{f_1}{2f_2} - \frac{D_{MCF}}{A_{lens}} = \frac{1}{|M|} - \frac{1}{|M_{max}|}$$

Beam-steering can be used to correct x/y-misalignments of the egress collimator up to $\pm\Delta\delta_{x,y}$ relative to the optimum in-line position, without resulting in beam clipping by the microlens aperture. Targeting a specific value for $\Delta\delta_{x,y}$ leads to a reduced maximum value for the magnification factor M, given by:

$$\left|\widetilde{M}_{max}\right| = \frac{|M_{max}|}{|M_{max}| \cdot \Delta\delta_{x,y} + 1}$$

Assuming that the focal lengths are adjusted for $|M| = \left|\widetilde{M}_{max}\right|$ and that the maximum collimator aperture is used (similarly to above), we obtain a modified maximum number of ports, which can be expressed as:

$$N_{x,y}^{max} = 1 + \frac{\theta_{x,y}\pi MFD}{2\lambda\delta_{x,y}}(A_{coll} - D_{MCF})\left|\widetilde{M}_{max}\right|$$

It can be seen from the above that a given target $\Delta\delta_{x,y}$ leads to a larger reduction in port count for MCFs with a small number of cores. This is essentially because less beam-steering is possible in this case due to the larger distance between the collimator and microlens. Using the parameters of the implemented 4-core MCF switch, the value of $\Delta\delta_{x,y}$ is only 0.6%, which indicates the high precision required when aligning the collimator and microlens arrays. Fig. 2.5 shows the maximum port count vs D_{MCF} in one dimension (similarly to Fig. 2.3), now plotted for various values of $\Delta\delta_{x,y}$ ranging from 0% to 50%. For the 4-core MCF, a misalignment tolerance of just $\pm 1\%$ requires a reduction in maximum port capacity from 17 to 13 ports (~75%). For a more practically useful tolerance of $\pm 5\%$, the capacity must be reduced to 7 ports (~40%). For MCFs with higher core counts, the reductions in

port capacity become less significant, as can be inferred from the figure. For a tolerance as large as the collimator spacing ($\Delta\delta_{x,y} = 50\%$), it is barely possible to have 2 ports, irrespective of D_{MCF}. This is because the WD is too low for a collimator to scan across the distance between two egress collimators.

These results clearly indicate the high degree of precision required for large port counts when mounting the collimator array and microlens array in the switch. Furthermore, the lens in each collimator must be accurately centered relative to the MCF fibre. If this is not the case, the collimator angle must be slightly shifted relative to the ideal orientation, which also results in a lateral misalignment.

The above calculations were based on ray matrix calculations and on the requirement of no beam clipping/distortion. Note that a slightly higher misalignment tolerance is possible if some excess core-to-core loss can be accepted (estimations of such tolerance is the subject of future work).

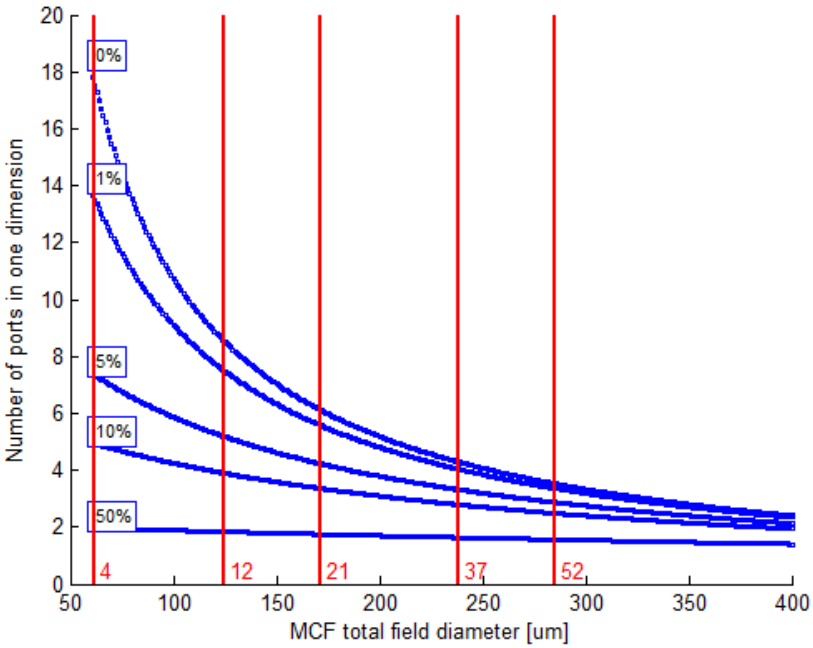

Fig. 2.5. Switch scalability for various lateral misalignment tolerances $\Delta\delta_x$ from 0% to 50%. The graphs shows the possible number of ports in one dimension vs. the total MCF field diameter D_{MCF}. The vertical lines indicate the values of D_{MCF} for a number of square-lattice MCFs with different numbers of cores ($n_c = 4, 12, 21, 37, 52$), all illustrated in Fig. 2.4.

2.4 Summary of multi-core optical switch efforts in COSIGN

In the COSIGN project, we have developed the first beam-steering optical multilane switch with integrated multicore fibres. We have fabricated 4-core MCFs specifically for this purpose and have integrated them directly in the Polatis beam-steering platform. The switch configuration is based on a telecentric setup including a microlens array to counteract beam divergence, enabling a large working distance for large port count. We have demonstrated low MCF-to-MCF switching losses in the range 1.2-2.2 dB in the fabricated multilane switch prototype. A 2x2 switch setup including an MCF transmission span has been fabricated for the long-term demonstrator experiment (characterisations of the switch setup were reported in D2.7 [6]). The switch has the potential for holding up to 96 ports with 4-core MCFs. In preparation for future work, we have provided estimates for the alignment tolerance when mounting ports in the switch. The required precision is high and improved alignment procedures are currently being planned. Finally, the scalability properties of the switch when moving to higher core-count MCFs have been estimated, showing an inverse scaling of the port count vs. the core count.

3 Multi-lane few mode fibre transmission

During the course of the COSIGN project, it was decided to investigate the use of few-mode fibres (FMF) as a means of increasing the capacity of single fibre. Mode multiplexing does have attractive features compared to MCF: simplified alignment of fibre components not requiring rotational alignment, more favourable scaling between the number of channels and the area to be imaged in a Polatis switch and more mature fibre technology. To make mode multiplexing relevant for transmission in a datacentre setting, it must be required that stable transmission can be achieved without requiring complex digital signal processing (i.e., MIMO processing is required for multimode transmission where the modes couple to each other). In order to achieve stable transmission we consider transmission using modes belonging to mode groups which have weak coupling in the fibre. Detailed characterisations of the coupling between these mode groups in fibres developed by OFS are reported in [3]. From these results it is clear that coupling inside the fibre can be low enough to allow practical transmission using several modes. As discussed in [6], a key challenge for practical applications of mode group multiplexing in FMF is the unavailability of compact practical multiplexers and de-multiplexers for the individual mode groups. In order to be able to send and receive data in different modes, there has to be a way to access the individual spatial channels as well as to multiplex or demultiplex them, respectively at the transmitter/receiver. The results presented in this chapter illustrate the joint effort of OFS and DTU in investigating mode multiplexed transmission. OFS provided and characterised the few-mode fibre used and DTU developed and fabricated the integrated silicon devices used as a mode multiplexer and demultiplexer.

3.1 Integrated devices for mode multiplexing and de-multiplexing

The developed mode multiplexer and demultiplexer [9] are equivalent, but used to either multiplex or demultiplex the spatial channels. Fig. 3.1 illustrates the topology of the mode multiplexer. Up to 3 LP modes, i.e. LP01, LP11 and LP21, can be multiplexed. Additionally, polarization multiplexing is possible, thus allowing for the number of distinct channels to double. At last, by using two different orthogonal orientations of the LP11 mode, the number of spatial channels can increase to a total of 8.

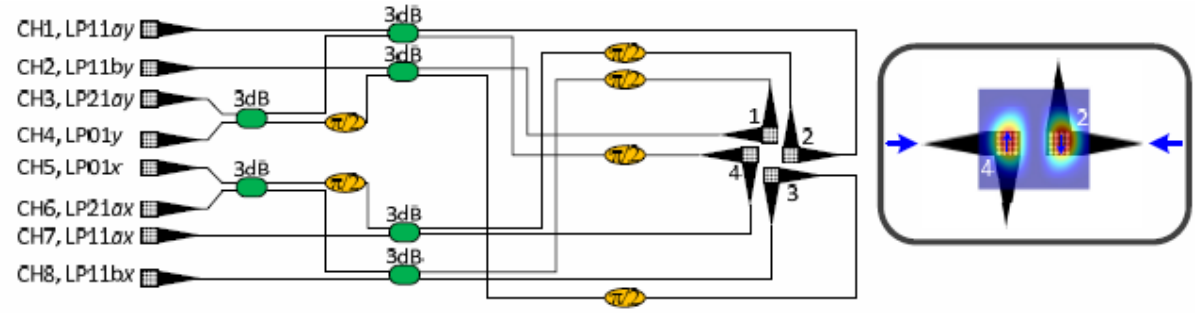

Fig. 3.1. Topology structure of the mode multiplexer device. The inset shows the principle to excite the LP11ay mode.

The mode multiplexer is based on a four-grating coupler based mode exciter. Light injected into CH1 attempting to excite LP11*a*y mode is first split into two beams by a 2×2 MMI. Note that there will be a π/2 phase difference between the upper and lower output arms of the 2×2 MMI. With another π/2 phase shift, the two beams are injected into the grating no.2 and no.4 respectively resulting in vertical polarization with a phase difference of π, thereby exciting the LP11*a*y mode for the FMF, as shown in the inset of Fig. 3.1. The same principle is applied to excite the other LP11 modes for CH2, CH7 and CH8. For CH3, light is first split into two beams by a 2×2 MMI. With extra π/2 phase shift, the two beams are injected into the following two 2×2 MMIs respectively with phase difference of π. Each beam is then split into two beams afterwards by the corresponding 2×2 MMI, and injected into the

corresponding grating couplers. The phases of these four beams in the grating coupler no.1~no.4 will be (0, π, 0, π) polarizing at vertical direction, and LP21*a*y can be excited. While light from the CH4 is split to four beams being injected into grating coupler no.1~no.4 with the same phase, and LP01y can be excited. The same methods can be applied to excite the other LP01x and LP21*a*x modes.

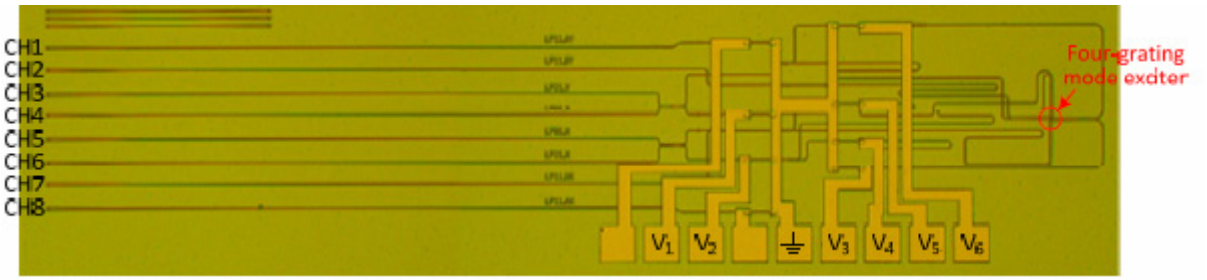

Fig. 3.2. Fabricated silicon chip for mode multiplexing/demultiplexing.

The devices were fabricated on a commercial SOI sample with top silicon thickness of 250 nm and BOX layer thickness of 3 μm, as shown in Fig. 3.2. Fully etched photonic crystal grating couplers are used for coupling from standard single mode fibers (SSMFs) to the corresponding channels. In order to increase the coupling efficiency, an Al mirror was introduced by a flip-bonding method.

3.2 Transmission performance

One chip was initially characterised using a 2-m long FMF supporting LP01 and LP11 modes with effective radius of 5.35 μm^2 and 5.49 μm^2, respectively. By applying proper voltages on the heaters V1 to V6, LP01 and LP11 modes polarized at both x and y polarizations were successfully excited, as shown in Fig. 3.3. From Fig. 3.3, one could also find that LP21ay and LP21ax modes show very weak field profile at the output of the fibre, the reason being that the FMF does not support the LP21 mode. We further measured the coupling efficiency as a function of wavelength for all the supported LP modes. One could find that a highest coupling efficiency (from the input SSMF to the output of the FMF) of –10.6 dB with 3.7 dB MDE is achieved. In addition, the 2 dB coupling bandwidth covers the whole C-band. Such coupling efficiency considers coupling loss of the input grating coupler, the chip loss, the coupling loss to the FMF, and propagation loss of the FMF, and is a 13-dB improvement compared to previous demonstrations.

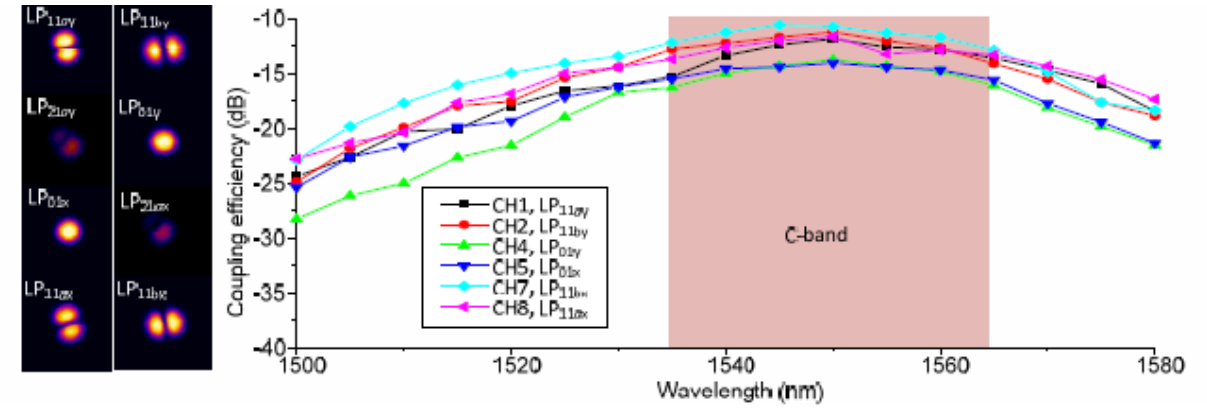

Fig. 3.3. Experimental results. Measured field profiles for different LP modes output from the FMF, and corresponding mode coupling loss.

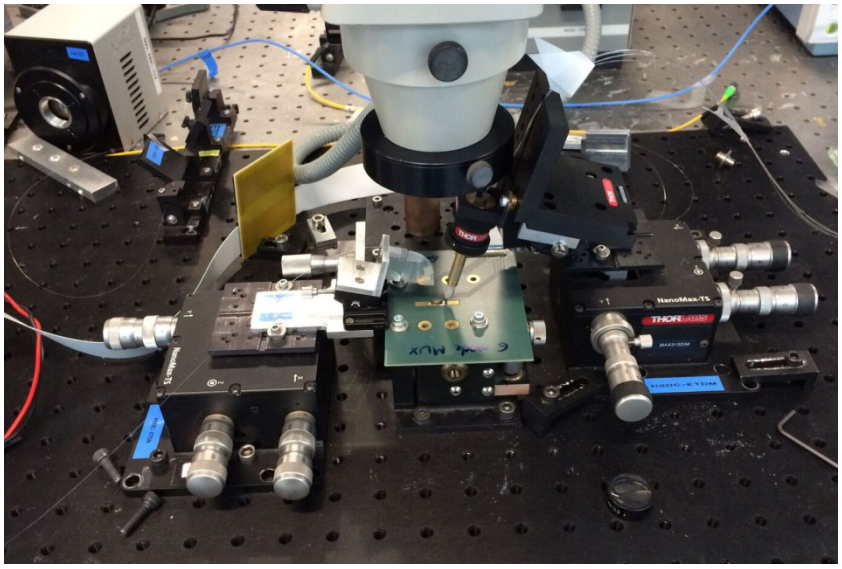

Fig. 3.4. Photo of the mounted device. On the right side, a fibre array is aligned to the input waveguides and on the left side the few-mode fibre is aligned to the grating couplers.

Due to the superior crosstalk performance of the FMF fabricated by OFS, we have performed all further studies using their two-mode graded index fibre. The mounted device with the aligned fibre array at the input and FMF at the output are shown in Fig. 3.4. The initial test were performed with a shorter length of fibre but the crosstalk performance for 15 km has been found to be -27 dB and -28 dB for LP01 and LP11, respectively.

For characterization purposes we captured the mode profile of the different LP modes from both chips, when both are used as multiplexers. The mode profiles are shown in Fig. 3.5. Similar behaviour can be observed for both chips, with fibre alignment and heater tuning contributing to fine optimization.

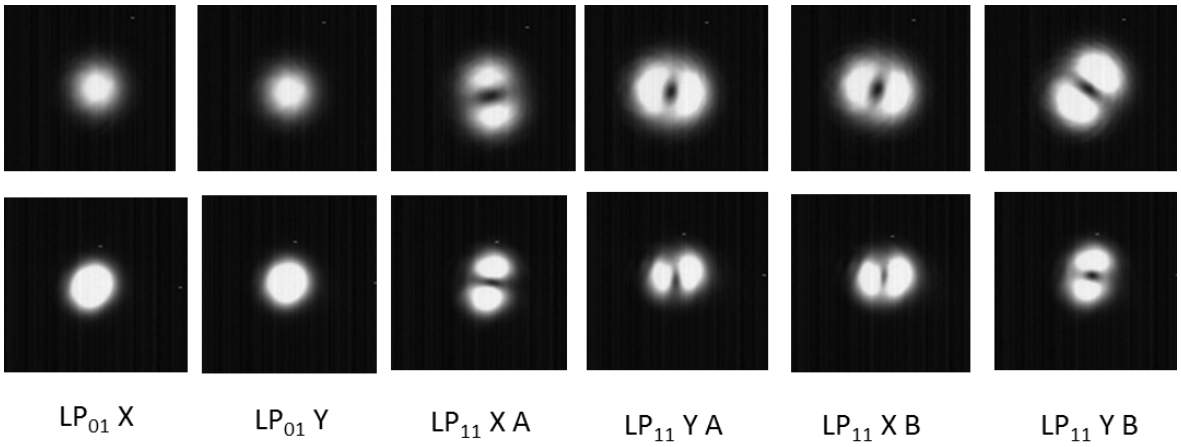

Fig. 3.5. Measured field profiles for the different LP modes at the output from the FMF from the multiplexer chip (up) and from the demultiplexer chip when used as a multiplexer (down).

In addition, we transmitted 10 Gbit/s data using single polarization of two modes only, namely the LP01 and LP11 modes that are supported by the specific few mode fibre used and belong to different mode groups ensuring low coupling between them. Error-free performance is achieved for both channels, demonstrating the first few mode fibre transmission with on-chip mode multiplexer and demultiplexer. This not only confirms the feasibility of the scheme, but also opens doors for future demonstrations scaling transmission reach and mode-count.

3.3 Prospects of integrating multi-lane few mode transmission in DCNs

Mode-division multiplexing (MDM) is emerging as an interesting alternative to multicore fibres for spatial multiplexing. FMFs are mainly attractive because unlike MCF, they have cylindrical symmetry and do not require precise rotational alignment during interconnection and splicing. However, the crosstalk between the modes due to distributed mode coupling in the fibre can greatly affect the overall system performance. Although the use of MIMO-DSP can help mitigate the crosstalk, its use in data centre networks is not desired because of its relatively high cost and added transmission latency. Thus, the crosstalk in the FMF has to be as low as possible. The fibres fabricated by OFS have demonstrated acceptable crosstalk values for lengths of few 10's of kilometres, which is sufficient for data centre applications.

Another important component for MDM is an efficient mode multiplexer, which should launch orthogonal LP modes into a FMF independently with low insertion loss and a large bandwidth. A number of non-integrated technologies using spatial light modulators, and phase plates have recently been demonstrated. However, in order to make MDM transceivers attractive, and especially to be able to integrate with other functionalities, it is preferable to develop solutions using photonic integrated circuits (PICs). By using photonic integration, mode multiplexers and demultiplexers can be made compact, which is extremely important for building future generation transceivers. Moreover, these devices can be used as a basic building block for fabricating more complex devices. An example is a few mode fibre switch that can reconfigure the different modes by first converting the mode channels to data carried in the different waveguides, subsequently switching them and multiplexing them again into a mode multiplexed data stream. The demonstrated performance of the on-chip mode multiplexer and demultiplexer indicates that few-mode fibre transmission is a feasible alternative for increased capacity for intra data centre networks.

4 Laser-written silica interposer for VCSEL-MCF interface

The work on the laser-written silica interposer for interfacing a VCSEL array and the 4-core MCF complements the work done in TU/e on novel transceiver assembly techniques (see [4]).

Since these transceivers make use of linear arrays of emitters/detectors it is necessary to create a unique optical interposer to take the light from a linear array (4 in this demonstration) and re-configure the waveguides to conform with the 4-core MCF which has a rectangular core arrangement. In the following sections we give more details on the different efforts carried out in the TU/e and UNISOUTH to create this unique transceiver module.

4.1 FPC design for 4 channel 25Gbps 1530nm single-mode(SM) transmitter

1530nm SM VCSELs are capable of transmitting high-data communications transfers over long distances. In the transmitter module, the Vertilas 25Gbps 1530nm SM VCSELs will be used. The carrier for this transmitter is based on a two-layer FPC with coplanar transmission line and the components will be flip-chipped on this FPC carrier. This work was carried out in collaboration with the University of Ghent for the long wavelength VCSEL drivers. The FPC for this transmitter has been designed and fabricated, as shown in Fig. 4.1.

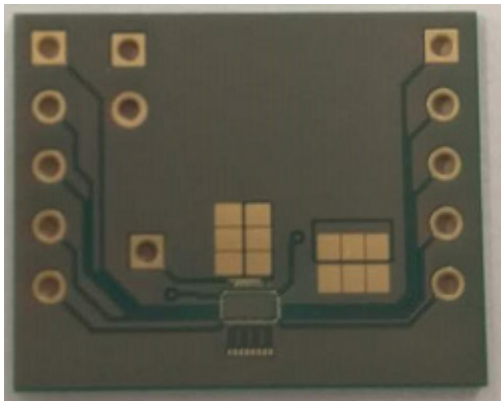

Fig. 4.1. FPC for 4 channel 25Gbps 1530nm single-mode(SM) transmitter.

4.2 Fabrication of the interposer

Femtosecond lasers have been exploited for inscribing waveguides in various types of glasses for more than a decade. The typical losses for femtosecond laser written single mode waveguides are close to 0.3 dB/cm. With refractive index change of 0.005, turning diameter smaller than 15 mm can be achieved without introducing significant losses (<0.2 dB). These parameters together with the ability of a femtosecond laser to implement complex geometries makes this technology ideal for 3D waveguide writing.

Waveguide writing is implemented using a high precision three axial translation system together with a high repetition rate amplified femtosecond laser (pulse duration 200 fs). Thanks to high stability and power budget of the laser system we were able to inscribe waveguides with speed >5 mm/s. For waveguide inscription we decided to use a multi-scan approach, where multiple lines are written parallel to each other with a spacing of less than 300 nm to form a single waveguide with a square cross-section. This approach allowed us to easily implement the complex 3D structures required for the interposer. A single waveguide could be inscribed in less than 5 min. The total fabrication time for the 4-core interposer was less than 20 min. After inscription, both sides of the samples were lapped and polished to expose the waveguides' cross-sections and to enable efficient light coupling.

Fig. 4.2 shows a 3D schematic diagram of the waveguide structure of the interposer. There are 4 waveguides, designed to couple light from 4 VCSELS in a 1-dimensional array with a pitch of

250 μm to the 4-core MCF with cores in a square lattice of pitch 36 μm. The outer dimensions of the glass sample are length 10mm, width 4mm, and height 1mm. The waveguides are written in the centre of the glass sample. The core size is about 10 μm to provide a good mode match to the MCF fibre cores.

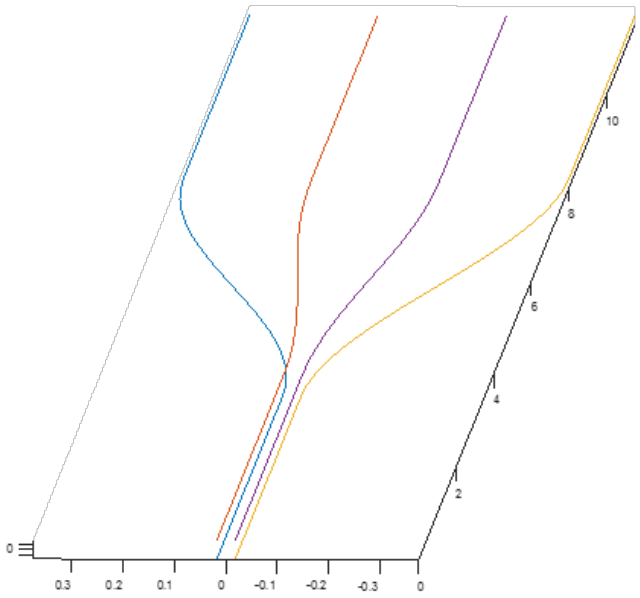

Fig. 4.2. 3D schematic of the waveguide geometry inside the laser-written interposer. The units on the three axes are mm.

Fig. 4.3 shows some microscope images of various sections of the interposer, with the VCSEL input end in (a), part of the 'tapered' section in (b), and the MCF output end in (c). The middle four images shows the top view of the interposer when red light from a HeNe laser is coupled into each of the 4 cores. Preliminary characterisations of the interposer show that single-mode propagation with a loss of about 1.5-2.5 dB (including coupling loss and propagation loss) is obtained for the two central cores (2 and 3). The output mode profile for cores 2 and 3 is shown at the bottom of Fig. 4.3. The mode profile of the outer cores (1 and 4) shows a deviation from the desired single-mode behaviour. At the time of submission of this deliverable, characterisations and further optimisations are underway to improve the output mode profile and if possible further reduce the propagation loss of all waveguides.

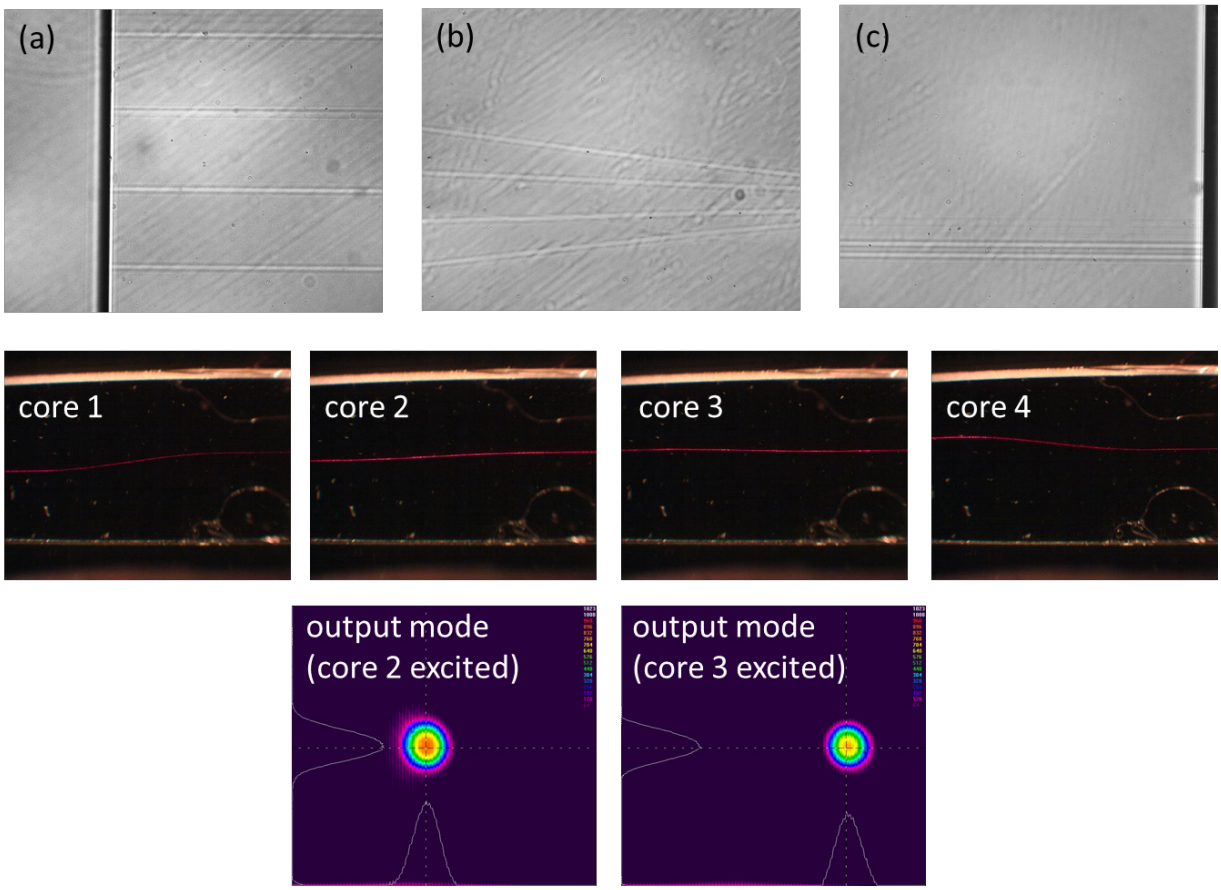

Fig. 4.3. Microscope images of the interposer (top view). (a) VCSEL input end (1-D pitch 250μm), (b) central section of the interposer where the 4 waveguides converge towards a square lattice, (c) MCF output end (2-D square lattice of pitch 250μm). Middle four images shows the interposer top view when HeNe laser light is coupled into each of the 4 cores. Lower two images shows the output mode profile at the MCF end when cores 2 and 3 are excited at the VCSEL end with 1550 nm light (using butt-coupling with an SSMF fibre).

5 Hollow-core fibre (HCF) for low latency DCN links

5.1 Motivation for HCF-based links in DCNs

Hollow Core Photonic Band-Gap Fibre offers a 30% reduction in propagation delay when compared to Single Mode Fibre [10]. The proposed Long Term Demonstrator (LTD) architecture includes an optical TDM data plane [9]. The nature of TDM imposes relatively high latency on data transmission, due to the time taken to serialize data into a TDM frame slot.

We aim to explore the extent to which the lower propagation delay of HCF affects the behaviour of the TDM network in terms of both throughput and latency. We also include measurements of packet loss.

5.2 Description of the demonstration

For the purposes of exploring the impact of HCF on the optical TDM plane, we connected two servers using FPGA TDM NICs and two standard optical NICs. The standard NICs can be recognised by the OS (the TDM NICs lack drivers for this) and enable application-layer traffic to be transmitted over the TDM network. This arrangement is illustrated in Fig. 5.1.

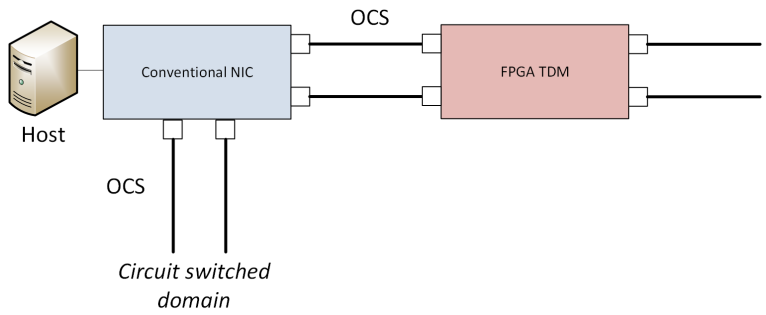

Fig. 5.1. Arrangement of ICs to enable application layer use of the TDM network. N.B. in the HCF/SMF experiments the OCS network is not used.

We measure throughput and latency of the TDM connection with a range of slot allocations and patterns with both SMF and HCF. We used 100m of SMF vs 100m of HCF on the TDM NIC. The bandwidth of the TDM network is configurable at the hardware layer by altering the ratio of real data transmitted to key characters and so we obtain results for different bandwidths. In addition, the TDM scheme exposed to the application layer enables provisioning of fractions of the hardware-configured bandwidth. The scheme has 96 slots per TDM frame, which can be enabled individually, and we present results with a range of slots enabled.

5.3 TDM network behaviour differences with HCFs

In this section, we present results from experimentation with HCF, vs SMF. The primary aim was to understand the impact of HCF on the performance of the TDM data plane and whether performance would be affected. If this was the case, we aimed to derive recommendations for provisioning of the TDM place for applications. A fibre length of 100m was used for experiments to allow performance differences to be more easily detected. The test configuration is illustrated in Fig. 5.2.

	SMF	**HC-PBGF**
Propagation delay	5 ns/m	3.5 ns/m

Latency over 10m fibre	50ns	35ns
Latency over 100m fibre[1]	500ns	350+5ns ([1] Includes SMF adapters from SC-APC to LC SFP+ connector)
Latency over NIC-NIC direct connection without TDM with 100m fibre	8140ns	8030ns
Average latency with TDM over 100m fibre (maximum bandwidth)	50036ns	41664ns
Minimum latency with TDM over 100m fibre (maximum bandwidth)	41427ns	34162ns
Maximum latency with TDM over 100m fibre (maximum bandwidth)	57153ns	48540ns

Fig. 5.2. Latencies for the various components in the TDM experiment

5.3.1 Throughput Measurements

The TDM scheme allows a variable number of slots to be utilised for transmission. With the optimal configuration, the peak throughput attainable is 8.6Gbps. To establish if the lower latency of HCF could affect the timeslot-to-throughput settings we measured the bandwidth over the 96 possible slots with IP traffic.

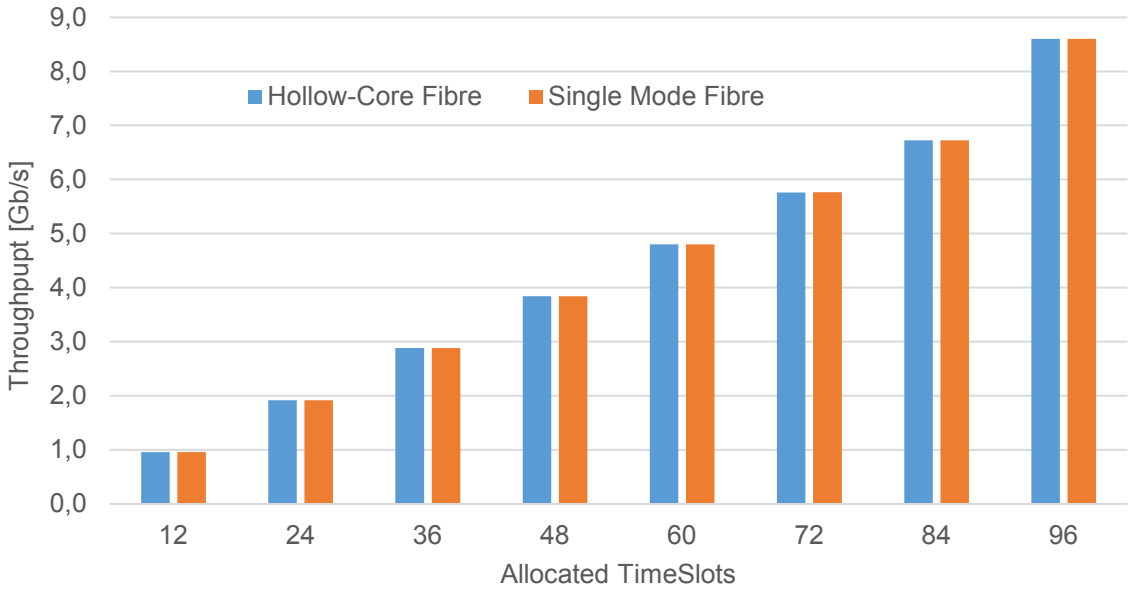

Fig. 5.3. Throughput measurements for HCF and SMF fibre with increasing time slot allocations

As can be observed in Fig. 5.3 there is no noticeable effect on throughput when using the Hollow Core Fibre. Each allows the TDM connection to achieve the maximum theoretical throughput when the full slot allocation is used. Fig. 5.4 shows the experimental setup evaluation to compare HCF and SMF connections with FGPA TDM NIC cards.

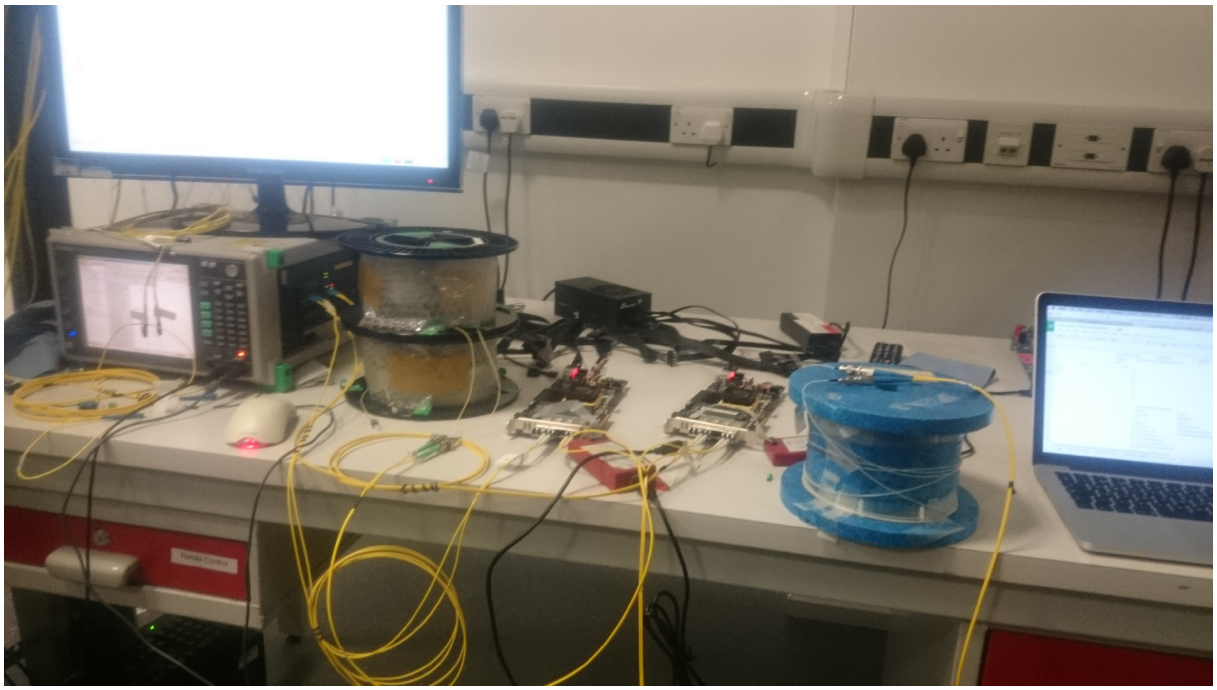

Fig. 5.4. Experimental setup for comparing HCF and SMF performance in the TDM data plane. Pictured PCBs are the FPGA TDM NIC cards.

5.3.2 Latency Measurements

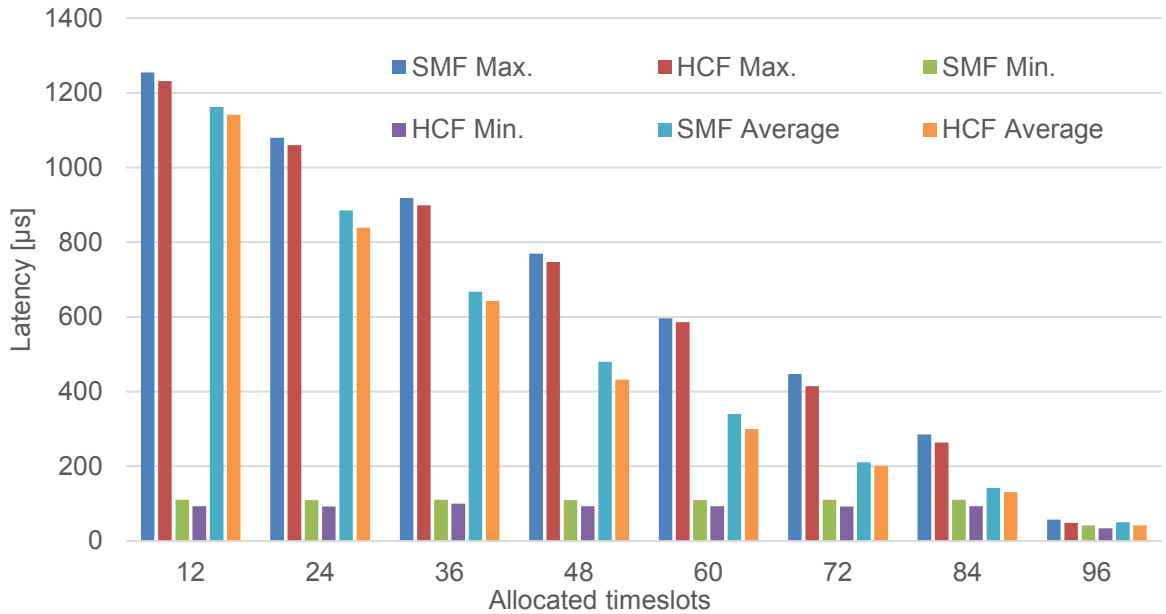

Fig. 5.5. Latency measurements for HCF and SMF with diffferent time slice allocations

Using HCF should offer improvements in latency characteristics. As TDM schemes incur additional latency when slicing data into timeslots, it is possible that HCF can help alleviate some of this latency penalty. Again, we took multiple measurements with different time slice allocations. The maximum, minimum and average latencies were recorded.

As it can be observed from Fig. 5.5 HCF exhibits lower maximum, minimum and average latencies. The nearly flat results at 96 timeslots are caused because all transmit slots are full. The other slot allocations have several slots during which no data is transmitted, incurring a latency penalty. However, the data in Fig. 5.2 shows us that even with 100m of fibre, by far the largest component of

latency is contributed by the buffering and logic of the FPGA TDM NIC. We observe the greatest percentage latency saving when all time slots are utilised, this is because we have removed another source of latency – a timeslot in which data is not transmitted.

5.3.3 Improved Timeslot Allocation

We attempted to improve on the previous performance measurements. It is clear from the results in Sections 5.3.1 and 5.3.2 that when all timeslots are utilised the latency behaviour changes significantly and that empty (no-transmit) timeslots contribute a significant fraction of latency. Therefore, we repeated our experiments by interleaving timeslots instead of allocating them in a contiguous block. In the experiment this means that only half the slots can be utilised, but two or more applications could run concurrently with interleaved timeslots provisioned by the orchestration layer.

For the same number of timeslots, we show the change in throughput when using interleaved instead of contiguous allocation in Fig. 5.6. It is clear that identical throughput increases can be obtained regardless of the fibre type. Therefore, the latency of the link does not affect the throughput of the TDM scheme. However, we will examine the latency changes also. From the data shown in Fig. 5.5 and Fig. 5.6, we would expect that change in latency will differ for the different fibre types.

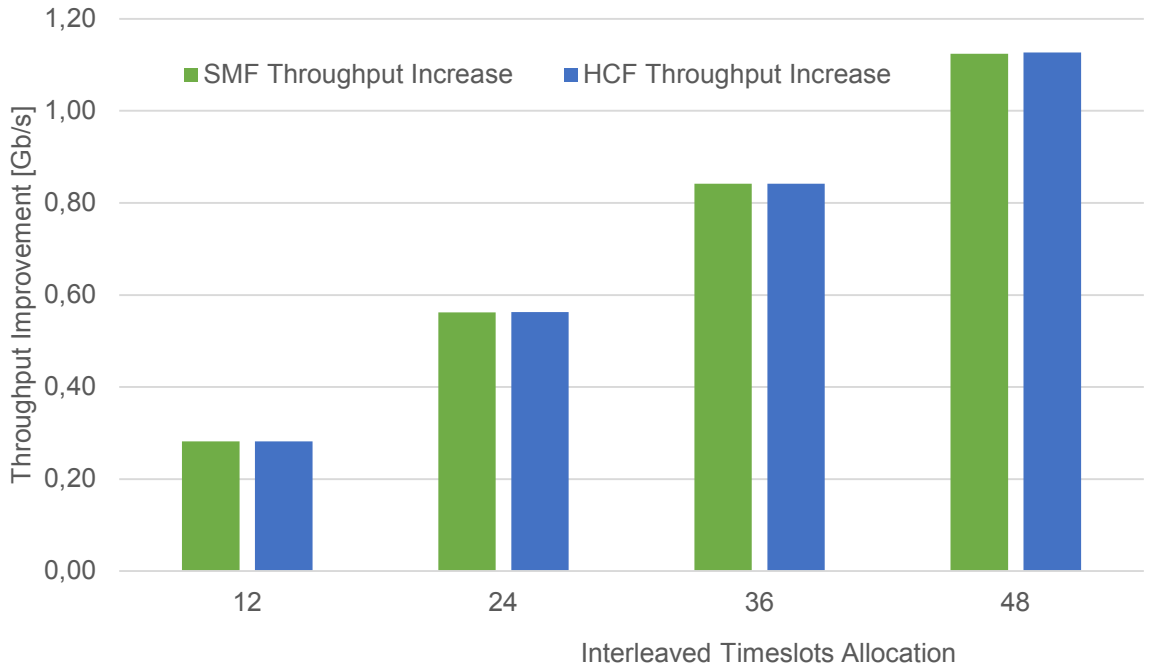

Fig. 5.6. Throughput improvement when using interleaved instead of contiguous timeslot allocation for both SMF and HCF

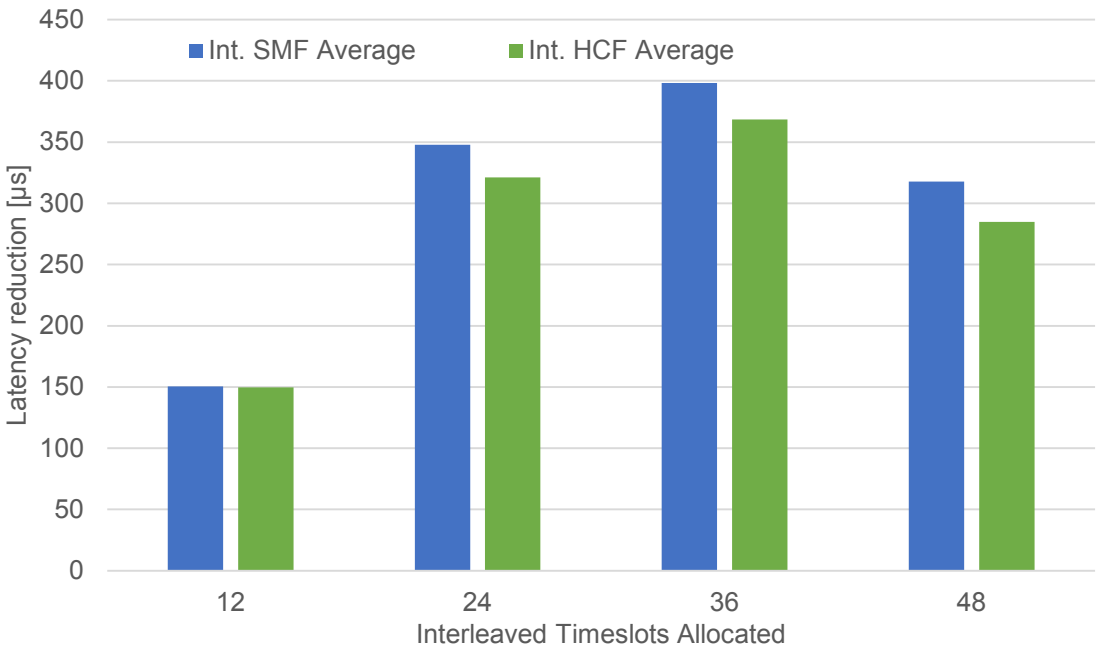

Fig. 5.7. Latency improvement gained through interleaved timeslot allocations instead of contiguous

In Fig. 5.7 we observe improvements in average latency readings across the spectrum of timeslots. The allocated timeslots occupy a greater area of the TDM frame, thus reducing average latency. However, overall the average latency improvements are greater for SMF than HCF.

5.3.4 Summary and future experiments

It is clear that HCF introduces latency savings when used in conjunction with the TDM scheme. However, the results show that the scheme is largely bound by the latency of the TDM NIC rather than the fibre interconnect. Both HCF and SMF TDM implementations can gain latency and throughput improvements from non-contiguous time slice allocation where this is possible.

The next steps are to experiment with TCP settings to establish if this protocol can work effectively over the TDM network. The results with HCF are encouraging. With the magnitude of latencies we encounter with this TDM scheme, a large TCP window size will be required. However, the improved latency performance of HCF reduces the necessary size of the window and thus impacts system resources less.

6 Summary and conclusion

The COSIGN project has produced a number of unique innovations in optical technologies for datacentre networks. Below we will outline the main achievements within the tree major topics of fibre technology development in the COSIGN project.

Multi-core fibre

Through a coordinated and persistent collaboration between several project partners key elements for a multicore fibre platform for data networks has been developed. A novel 4-core fibre concept has been developed and two different fibre types have been produced to satisfy the different requirements of switch-integration and transmission performance. A prototype of a multicore fibre switch has also been developed allowing the use of 4-core fibre in an optically switched data plane. This constitutes the first demonstration of integration of multi-lane fibre transmission and switching and constitutes a major milestone on the path to practical applications of multi-lane optical data planes.

To complete the multi-core platform, the project has initiated development glass interposers to interface the arrays of VCSEL light sources and photodiodes used in transceivers to the cores of a multi-core fibre. Once this technology is fully developed a full multi-core fibre data plane can be realised.

Few-mode fibre

As an alternative or potentially a supplement to multi-core technology the COSIGN project has also investigated mode multiplexing as an optical technology for multi-lane transmission. Appropriate optical fibres (few-mode fibres) have been identified and characterised for multi-lane transmission confirming the potential of the technology. Finally, integrated silicon devices for multiplexing and de-multiplexing of modes have been developed and combined with few-mode fibre to create an entire transmission link based on the mode-multiplexing technology. This constitutes the first demonstration of a mode-multiplexed link based on integrated devices compatible with cheap mass production. As such, it is a key step in demonstrating the practical potential of the technology.

Hollow-core fibre

Low latency optical transmission has been investigated using hollow-core fibre technology. Specifically, the combination of optically time-shared connections and low latency fibre has been tested as a way of reducing the impact of the increased latency associated with temporal partitioning of optical connections. This is the first investigation of its kind and has mapped out the trade-offs between latency contributions from different sources for different link configurations. It is clearly indicated that as the latency associated with sharing of a link is brought down by e.g. faster transceiver circuits, the benefit from using hollow-core fibre for transmission will increase.

The above results are remarkable achievements and are all indicating that significant benefits can be realised by continued innovation of optical data plane technologies. None of the results reported here would have been possible without the collaboration of several project partners, confirming the tight collaboration between various constellations of partners.

Multi core- and few mode fibre technology both provide a promising path to achieving high spatial efficiency of highly aggregated data links. Hollow core fibre technology shows great promise for latency sensitive links and could alleviate the added latency associated with TDM links.

The level of integration of data plane technologies has surpassed the expectations prior to the project. The successful integration of multi-core fibre in the Polatis switch engine and the realisation of a complete mode-multiplexed link are both results that lie beyond what was believed achievable within the duration and scope of COSIGN.