Accelerating **data centre** performance

Dr Tolga Tekin and **Assistant Professor Nikos Pleros** explain how the PhoxTrot project is exploiting optical communications to streamline data centres and high performance computing systems

What constraints in computing and networking led to the inception of the PhoxTroT project?

NP: Power consumption and size are the main barriers to next-generation data centre (DC) and high-performance computing (HPC) environments. The limited capacity and physical quantity of electrically wired interconnects contrasts with the increased clock speeds and wiring densities inside machines. This creates a bottleneck to information exchange across all hierarchical communication levels: rack-to-rack, backplane, chip-to-chip and even on-chip.

TT: At the same time, parallelisation in processing brings completely new requirements with respect to the amounts of traffic that need to be exchanged within DCs and HPCs, essentially transforming computing into a shrunken networking environment. This reality has led to the accelerated penetration of optical technology into short distance, inter- and intra-rack transmission links.

Photonics holds great potential in the energy-efficient delivery of large amounts of data. What is the state-of-the-art in optical data transmission?

TT: Telecom-related products are able to provide the best solution for delivering large amounts of data. The main need is to learn from the experience of the telecom's system and implement the advantages from this environment to data communication.

In datacom, the main benefit is that value is added by incorporating the optical interconnect solutions, which reduces opex. Higher capex could at least be considered for the new generation of components and systems, which also allows and enables revolutionary solutions.

The project is developing optical interconnects across three hierarchies in data centre and high-performance computing systems, namely on-board, board-to-board and rack-to-rack. Could you outline what each entails?

NP: On-board interconnects are addressed within PhoxTrot through the development of a range of optochips for transmitted traffic (Tx), received traffic (Rx) and routing functionalities, properly adapted to optical printed circuit boards (PCBs).

Moreover, multi-and single-mode low-loss optical PCBs with appropriate sockets for interfacing in a 'plug-and-play' way with optochips will serve as the on-board platform, as well as for board-to-board interconnects through solutions that facilitate pluggable optical coupling between different PCBs. This will lead to optical motherboards capable of hosting optical daughterboards to offer the entire chain for intra-rack interconnectivity directly in the optical domain.

Concerning rack-to-rack interconnects, PhoxTroT is creating a 1.28 terabit per second active optical cable (AOC) module that will comprise silicon-based circuitry and is expected to rely on 16 Quadrature amplitude modulation (QAM) advanced modulation formats.

TT: State-of-the-art technological approaches provide isolated solutions to existing challenges, without considering the entire picture and overall compatibility. PhoxTroT is the first project to address the entire datacentre as an ecosystem.

Could you elaborate on the project's multidisciplinary approach?

NP: The idea of 'one-size-fits all' certainly does not apply to short-range datacoms, necessitating the use of a wide technology portfolio for realistic PCB and optochip developments. The recent revolution of silicon photonics brings an additional candidate technology into the game, since they have revealed a tremendous potential for complementary metal-oxide-semiconductor (CMOS)-compatible optical chip-scale circuitry.

TT: Reducing cost when such a variety of technologies is incorporated into the fabrication line is not a trivial task. Cost efficiency can only be achieved when all different technology platform requirements are aligned to a well-established process-level framework, and only if the proper balance in the use of the different on-board technologies is achieved.

PhoxTroT is the missing piece in advancing and synergising photonic interconnect fabrication platforms and technologies towards a clearly defined application field over the entire level of hierarchy. We collect the broadest possible technology and material portfolio including polymers, glass, III/Vs, silicon-on-insulator (SOI) photonics, plasmonics, Bi/CMOS electronics and high-density 3D integration and packaging. At the same time, the project involves the necessary system- and application-level expertise, bringing together the different material platforms confronted with interconnect testbed infrastructures and specifications set by HPC and DC end-users.

Are there any challenges with bringing such a diverse group of partners together?

TT: Coordination of such a large and diverse consortium is indeed very challenging. The most important thing is to establish a common language between people from different areas in order to achieve the right technological solution adapted to the specific needs of every hierarchical sector.

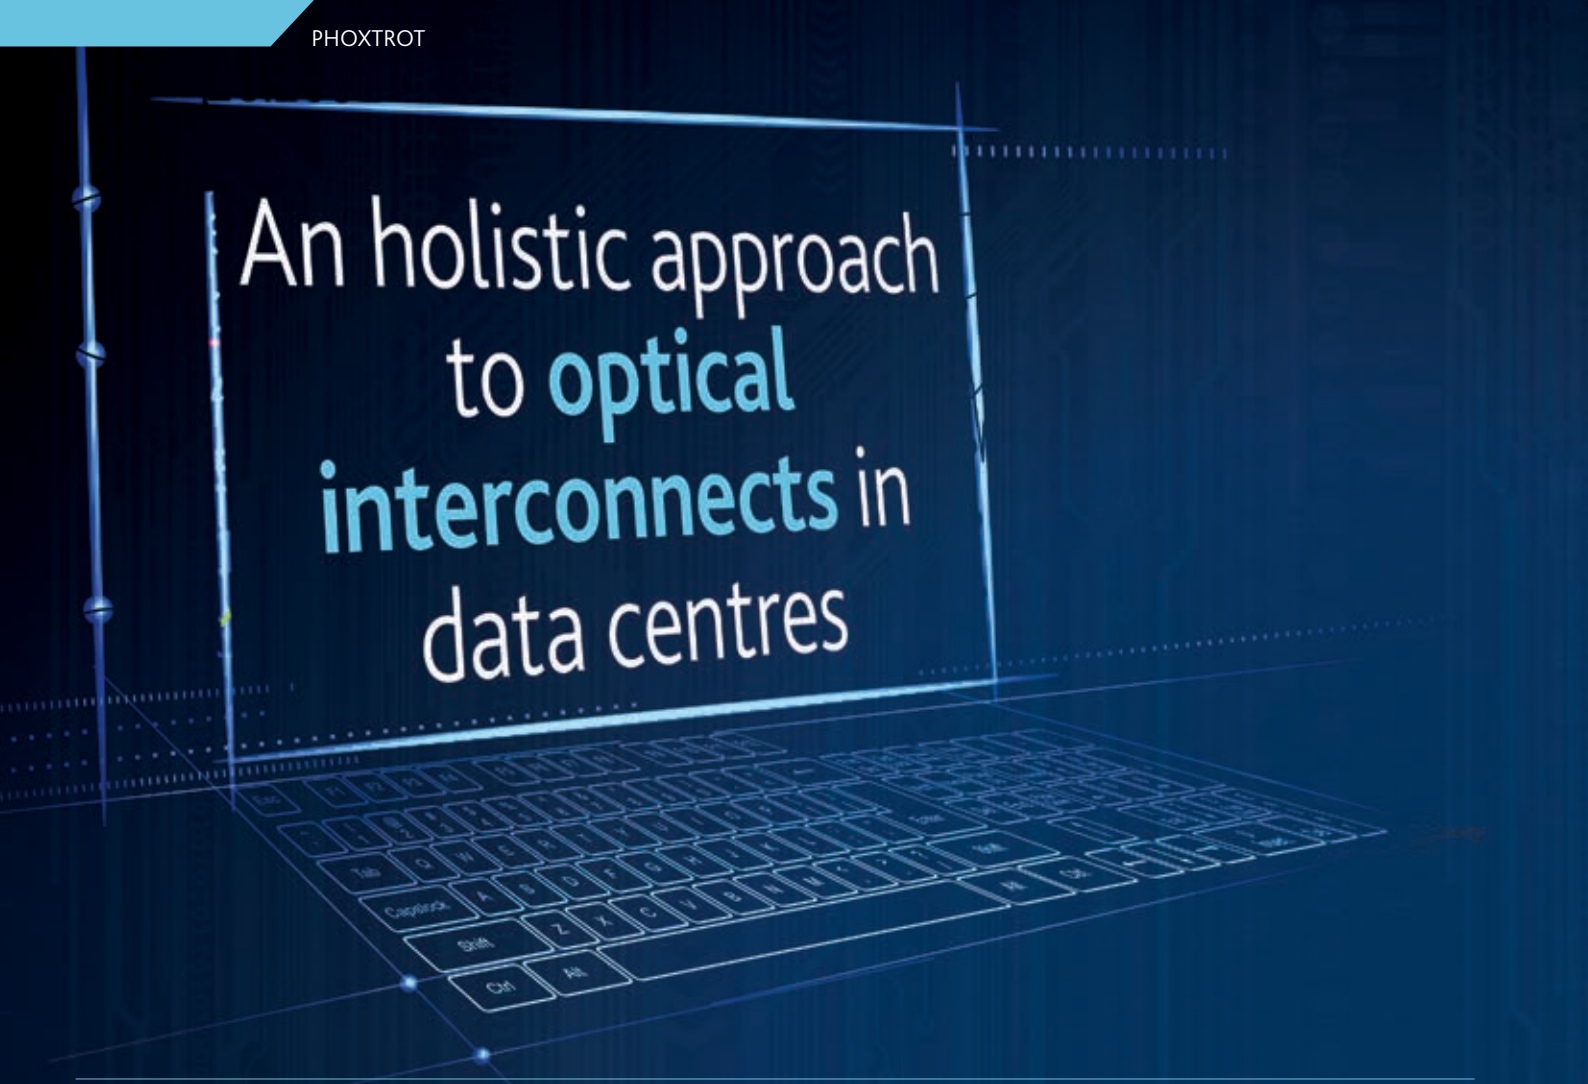

Led by the **Fraunhofer Institute for Reliability and Microintegration (IZM)**, the PhoxTroT project is aiming to lower operational costs whilst also increasing the performance of data centres

THE AMOUNT OF digital information used around the world has exploded in recent years. Posts on social media, movie downloads, email files and more all need a place to be stored so that they can be retrieved by users at any time. This is done through data centres. Although these digital warehouses keep consumers happy by allowing them to access their files on demand, data centres require a vast quantity of electricity to function. Globally, data centres use about 30 billion Watts of electricity, which is needed not only to keep the equipment running, but also to cool it in order to prevent overheating. This high energy use is very expensive, and with data use expected to continue to increase in the coming years, there is a need to find less costly, more efficient solutions for data centres.

The PhoxTroT project was founded with the aim of addressing the issues of power consumption and data usage in data centres and high performance computing environments. Coordinator Dr Tolga Tekin at the Fraunhofer Institute for Reliability and Microintegration (IZM) in Berlin describes the motivation behind the project: "Power consumption and high operational costs are the main drivers behind PhoxTroT. Since the main operational cost factor is directly related to the power consumption, new approaches to reduce those are the main interest of current developments and research".

PhoxTroT is being funded by the EU Seventh Framework Programme (FP7) for four years and brings together a highly interdisciplinary group of 19 partners from across Europe. Nikos Pleros, project member and Assistant Professor at the Department of Informatics, Aristotle University of Thessaloniki, Greece, explains how the consortium came to fruition: "PhoxTroT started with about five partners, working in different but complementary areas. This was very useful in identifying the key technological domains across all levels of interconnect hierarchy and subsequently approaching the right partners from all different technological sectors to provide the missing expertise in dedicated domains".

SPANNING HIERARCHIES

The research efforts within the initiative focus on high-performance, low-energy, low-cost and small-size optical interconnects across the different hierarchy levels in data centres and high-performance computing systems. These hierarchies include on-board, board-to-board and rack-to-rack. Such an holistic approach is expected to lead to a more optimised high performance computing and data centre environment by creating a synergy between the different optical interconnect technologies. "The use of optics at inter-rack communication level is not enough for sustaining performance enhancements," Pleros adds. "Low-energy photonic solutions have to penetrate at board-to-board and chip-to-chip data links, eventually also to intra-chip, in order to avoid an explosion in energy consumption."

Optical interconnect is a means of communication that uses optical fibres and waveguides. By using light, a much higher

bandwidth can be achieved compared to traditional copper cables. The technology is being explored as a link between supercomputers, as well as a means of facilitating communication within data centres. Another major advantage of using light to transmit data is that it requires much less energy than current methods. Although this technology, known as optical communication, has been widely researched in the past, the PhoxTroT project is the first to investigate the relationship between all of the components within data centres and how to best bring them all together to streamline how they work.

INNOVATIVE TECHNOLOGIES

In order to optimise the operation of data centres and high-performance computing systems, the consortium partners are working together to develop three prototype systems which are expected to be fully functional by the end of the project. The new technologies cover the entire chain of cost- and energy-efficiency, and include an optical printed circuit board with a throughput greater than one terabyte per second, a high-end optical backplane for board-to-board interconnection with a throughput greater than two terabits per second, and an active optical cable (AOC).

All of the new technologies are designed to reduce energy consumption by at least 50 per cent, which will in turn halve running costs. The savings in energy use do not come at the cost of a quality, high speed connection, however. In fact, data connection capacity is anticipated to double to two terabits per second. Tekin provides an example of how one of PhoxTroT's innovative technologies will provide cost savings: "The volume manufacturing processes expected to be established for the silicon-on-insulator-based Tx/Rx chips for PhoxTroT's AOC should bring its cost down to less than US $3 per gigabyte per second, even for the Quad Small Form-Factor Pluggable Plus (QSFP+), according to provisional estimates". To put this into perspective, that is 75 per cent less expensive than current modules.

Of course, achieving such revolutionary measures in computing is not without its challenges, as Pleros outlines: "This forward-looking pathway relies on the smooth transition from multi-mode to single-mode optical printed circuit board platforms and the enhanced degree of photonic chip-scale circuitry included in optochip deployments, taking into account the growing level of maturity of silicon photonics".

BUILDING BLOCKS FOR THE FUTURE

A year into the project, PhoxTroT group members are already taking steps to ensure that their work will have a long-lasting impact after its conclusion in 2016. "One of the main targets of PhoxTroT is to reinforce European industrial leadership and competiveness beyond 2020. That also means going beyond the state-of-the-art in order to explore new paths and solutions, by keeping the balance between the technologies," Tekin notes. While there are already a number of European initiatives that have focused on developing optical interconnects, the PhoxTroT project – with its large range of expertise – is expected to generate new ideas that will have a significant impact on achieving the full potential of optical interconnects.

One way in which PhoxTroT is helping to further Europe's leadership in the fields of optical interconnects and photonics is through the establishment of the European Cluster for Optical Interconnects (ECO). The initiative is designed to promote the visibility, sustainability and exploitation of European research and development efforts within optical interconnects. "This initiative seems to be on the right path – the first ECO workshop took place in London and was a big success with almost 100 people from all over the globe attending to discuss how to accelerate the commercial adoption of optical interconnects," Tekin highlights. Indeed, the PhoxTroT project is well on the way to revolutionising the computing environment and establishing the building blocks for the future of optical interconnects.

The PhoxTroT project was founded with the aim of addressing the issues of power consumption and data usage in data centres and high performance computing environments

INTELLIGENCE

PHOXTROT

OBJECTIVES

PhoxTroT is a large-scale research effort focusing on high-performance, low-energy and cost and small-size optical interconnects across different hierarchy levels in data centre and high-performance computing systems

PARTNERS

Fraunhofer Institute for Reliability and Microintegration (Fraunhofer IZM) • **Fraunhofer Insitute for Telecommunications Heinrich-Hertz-Institut** (Fraunhofer HHI) • **Vertilas GmbH** • **Xyratex Technology Ltd** • **ams AG** • **Meadville Aspocomp International Limited** • **Gesellschaft für Angewandte Mikro und Optoelektronik mit beschränkter Haftung mbH** (AMO GmbH) • **Institute of Communications & Computer Systems, National Technical University of Athens** • **DAS Photonics SL** • **Phoenix BV** • **Centre for Research and Technology Hellas** • **Compass Electro Optical Systems Ltd** • **Bright Photonics BV** • **Computer Technology Institute and Press** (Diophantus) • **Centre National de la Recherche Scientifique, Laboratoire Interdisciplinaire Carnot de Bourgogne**; **Laboratoire de photonique et de nanostructures** • **Karlsruhe Institute of Technology** • **University of Southern Denmark** • **Universitat Politècnica de València** • **Interuniversitair Micro-Elektronica Centrum vzw** • **TE Connectivity**

FUNDING

EU Seventh Framework Programme (FP7)

CONTACT

Dr Tolga Tekin

Fraunhofer IZM
Gustav-Meyer-Allee 25
13355 Berlin
Germany

T +49 30 46403 639
E project.phoxtrot@izm.fraunhofer.de
E tolga.tekin@izm.fraunhofer.de

www.phoxtrot.eu

DR TOLGA TEKIN is Group Manager of Photonic and Plasmonic Systems at Fraunhofer IZM, and is with the Research Center of Microperipheric Technologies, Technical University of Berlin. He is the current Chair of Technical Committee on 'Photonics - Communication, Sensing, Lighting' in the IEEE Components, Packaging and Manufacturing Technology Society.

DR NIKOS PLEROS is Assistant Professor and has been lecturing at the Department of Informatics, Aristotle University of Thessaloniki, Greece, since 2007. He is also a Collaborative Faculty Member at the Center for Research and Technology Hellas (CERTH).