

As the length or density of electrical interconnection increases they suffer from increased wire resistance, residual wire capacitance, fringing fields and inter-wire cross-talk.

Optical interconnection does not suffer from any of these problems due to its non-interacting nature: however this makes optics inefficient when trying to implement any type of gate.

Optoelectronics attempts to make the best of both worlds by using electronics for switching and optics for communication.

Non-Local Interconnection

Free space optical channels can pass through each other to form any desired interconnection topology without cross-talk. Interconnects such as the perfect shuffle (seen below) and the hypercube thus become relatively simple to implement. This also reduces skew as very large variations in wire length can be avoided.



Optical implementation

Since electrons carry mass and charge they interact strongly (Coulomb Interaction). Ideally suited for switching.

Electrons

A single channel may be easily routed to multiple output channels as shown below.

Channel splitting

DOE or SLM

CGH.

Single input channel



Photons do not carry mass or charge and are noninteracting in free space. They are ideally suited to interconnection.

2D detector array



Multiple diffracted output channels

The optical elements used for routing can either be active (SLM) or passive (DOE and CGH). Passive elements rout channels as defined during fabrication, however active elements can be changed in system to give any desired configuration.

Optical Interconnect Technology Heriot-Watt University

This poster outlines some of the most promising optical interconnect technologies.

Vertical Cavity Surface Emitting Laser (VCSEL)

The VCSEL array is an array of laser diodes which are fabricated on the same chip. Such arrays are attractive as they have a high optical output powers.

It has also been shown that VCSELs can be driven at data rates in excess of 10GHz.

The array shown here is currently being used by Heriot-Watt University and has a worst case turn on delay of 1.6ns (625MHz).





Diffractive Optic Element (DOE)

Phase-only DOEs are created by selectively etching fused silica to form phase profiles which have been optimised to produce the desired intensity patterns in the far-field of a Fourier lens. Elements with efficiencies of >70% and non-uniformities of <3%are routinely produced by the Diffractive Optics Group at Heriot-Watt University.

These elements are used as array generators and interconnection elements as well as more complex beam shaping elements for laser material processing.

Sample DOE



Multiple Quantum Well (MQW) Modulator



Photodetector and Drivers

Photodetectors act as input devices and are currently available in a wide range off-the-shelf. They are already responsive enough to handle input from any emitter: however the faster they are driven the more power they require.

The important considerations are input sensitivity, power dissipation, area usage and speed.

The packet switch controller described here uses the device shown to the right.



DOE Output (Single beam input)



Bonding Techniques: Smart Pixels Heriot-Watt University

Inline System

Optical

input

Electrical IO

In general the optoelectronic interface is constructed of hybrid processing chip technologies, which employ GaAs optical chips hosting detectors and emitters, flip-chipped on top of an Si FPGA using solder bumps. The combination of input, processing and output elements is generally known as a smart pixel.

A three layer chip is shown here which makes visualisation of circuits simpler (actual components at present are two layer with input and output on the same surface).



pads



We will refer to the combination of [dynamically reconfigurable] FPGAs and either optical input or optical output as an Optical FPGA (OFPGA). Such systems are capable of communicating internally electronically, or with any other local electronics for that matter, and can be viewed as a standard DR-FPGA with additional optical I/O pins.

The real advantage of a system using flip-chip bonding is that there are essentially no transmission lines, therefore its potential bandwidth is consequently higher. If the FPGA is dynamically reconfigurable then the inline system could be used for applications such as a telecommunications router or as an inline DSP-like system which could reconfigure its hardware to perform signal processing on an incoming data stream.



Flip-chip bonding between layers

Mapping Interconnects to FPGAs

OFPGAs can be conceptualised as having of three basic stages. The first stage, the input stage, consists of a detector array that is capable of receiving digital optical input.

The second stage is the processing stage and consists of a dynamically reconfigurable FPGA system that could be considered as one or more configurable logic blocks (CLBs) corresponding to a single detector input. The final stage is the output stage and consists of an optical emitter or modulator, again corresponding to one or more CLBs.

Heriot-Watt University

The CLBs seen here are not considered to be of any specific architecture - they are only used for illustration purposes.

Connection Block to Optical Channels



Sample CLB with Optical Input Channel



Vertical

output Mapping a detector or VCSEL/MQW only indicates that a CLB could potentially have access to the component. This grouping is referred to as an OFPGA element.

Single CLB Element



Optical

input





High Bandwidth Dynamically Reconfigurable Architectures using Optical Interconnects

Heriot-Watt University Edinburgh

Keith J. Symington¹, John F. Snowdon¹ and Heiko Schroeder² ¹Heriot-Watt University, ²Loughborough University

Project funded by EPSRC "Analysis and Modelling of Optoelectronic Systems" grant.

The maximum bandwidth of electronic systems has been estimated by Burton Smith (Tera Corp.) and David Miller (Bell Labs) as:

Consider chip connections for a 10x10mm chip (100mm²)

- Edge connections: 400 with ~100 μ m diameter lines \rightarrow A=3mm²
- 2D solder-bump array: 2,000 with ~15 μ m diameter pads \rightarrow A=3mm² Thus for a 10cm electrical connection across a board: B_{max} ~150GHz

The Attraction of Optics

- Off Chip Data Rates: Currently we can drive 4,096 channels from 1 cm^2 and see no real obstacle to reaching >10,000 channels.
- Bandwidth in Busses: A 1cm² relay can carry >100,000 channels. Currently we are driving at 200MHz giving a bandwidth of approximately 20Tbs⁻¹. Devices may routinely be driven at 10Gbs⁻¹ so the relay can handle 1,000Tbs⁻¹ if we are not CMOS limited (the theoretical limit is actually much higher).
- Data Acquisition: The naturally parallel nature of the connections implies high parallelism around any machine: e.g. to and from memory and peripherals.
- Pin-Out Limitations: The physical pin-out limit can be overcome with additional optical pins.
- Distance: Optical signal transmission lengths of the order of meters are attainable without a significant increase in driving power.
- Power Consumption: Off-chip optical interconnection uses less power in comparison to electronics.

SIA Roadmap - Electronic Bandwidths Required

Year	1995	1998	2001	2004	2007	2010
Process size (µm)	0.35	0.25	0.18	0.13	0.10	0.07
Chip size (mm ²)	450	660	750	900	1000	1400
On-chip clock (Mhz)	400	600	800	1250	1500	1900
I/O Bus speed (Mhz)	150	200	250	300	375	475
I/O Bus width	256	256	512	512	512	1024
Off-chip data rate (Gbs ⁻¹)	38	51	128	153	192	486

MEL-ARI Roadmap - Optoelectronic VCSEL & Driver Arrays

Year	1997	2002	2007
VCSEL+driver pitch (µm)	125	80	60
Chip size (mm ²)	6.25	9	16
Optical channels/chip	256	1024	4096
VCSEL data rate (Gbs ⁻¹)	0.6	1	2
I/O Bus Width	256	1024	4096
Off-chip data rate (Gbs ⁻¹)	154	1024	8192

 $B_{max} = 500. \frac{A}{1^2}$ Thz

Aspect ratio



Sorting Demonstrator

The architecture of the demonstrator utilises optoelectronics exploiting non-local interconnection: in this case the perfect shuffle. Below is a schematic of the sorting demonstrator and right its implementation using OFPGAs. The data to be sorted are entered sequentially into the processing loop through electrical I/O.

Heriot-Watt University

Sixteen bit planes of 32x32 bits (the number of optical communication channels) may be entered in this version. At run time, a 2D perfect shuffle is performed by a lens operation during each cycle of the machine and all computations are performed in parallel by the FPGA.



Circulating Inline System



The total number of cycles scales as $(\log N)^2$ for Batcher's bitonic sort of N data points. When computation is complete, the sorted set of data can be sequentially downloaded. OFPGAs introduce a level of flexibility that allows network reconfiguration to optimise execution time for different data array sizes.

The iterative perfect shuffle forms (with suitable node switching) an omega class network capable of arbitrarily permuting data. This set up could be used to implement any algorithm or multistage switching function given the right FPGA logic configuration. Indeed it is known that many algorithms map exceedingly efficiently onto this topology, the FFT being the classic example.



Incomina

packets

Neural Network Crossbar Switch Controller The assignment problem is essentially task allocation optimisation amongst all available resources to maximise throughput. Solving the assignment problem is computationally intensive and complexity grows exponentially with problem size.

This project examined specifically the assignment problem in a crossbar switch for packet routing: maximising throughput and minimising delay.

Neural networks are capable of solving the assignment problem. Their inherent parallelism and mode of operation allows them easily to outperform any other known method at higher orders.

The system as it stands at present has a fixed set of weights and is designed to perform this single task. The inclusion of the FPGA stages will allow, in the first instance, programmed weights

to be considered and consequently configuration of the network for a wider variety of tasks. The most exciting possibility is being able

to reconfigure these weights in near or at real time so that fully adaptive, supervised and unsupervised learning Lens 2 schemes may be implemented.

array

(Lens 1)

DOE **Microlens**

Object (VCSEL)



Hopfield network

 $\bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc$

Request

connections





Rows: i

Neurons

Neuron

The key to utilising the parallelism of a neural network is matching the network as closely as possible to the problem by altering the updating rule. The updating rule determines the next value a neuron will take based on the previous outputs of other neurons:

$$_{j}(t-1)+\Delta t\left(-ax_{ij}-A\sum_{k\neq j}^{n}y_{ik}-B\sum_{k\neq i}^{n}y_{kj}+\frac{C}{2}\right)$$

x_{ii}: Summation of all the inputs to the neuron referenced by ij: including



Optical Highways Heriot-Watt University

The concept of optical highways is to provide a general purpose multiprocessor harness with several thousands of channels passed node to node via an OFPGA interface.

The optical interconnect is 'hard-wired' using polarising optics to define computational topology. In essence, the destination of any data channel output into the optical domain is determined by the spatial location of the emitter output upon the OFPGA. Thus by re-routing signals within the OFPGAs, a particular global topology may be established.

The combination of FPGAs with high bandwidth optoelectronics enables an intelligent communications interface to be constructed which allows maximum utilisation of the ultra-high throughput available. In turn this enables real time optimisation and load balancing of the whole machine over a range of computational models.





The optical design package Code V was used to simulate relays built using conventional bulk optics and components.





algorithms.



Engineering Issues and Conclusions

Engineering Issues

When using optics in any practical system, various factors must be considered.

- Active effects: <1Hz thermal changes and component creep.
- Static effects: Tolerances in fabricated components could lead to misalignment in final system.
- Adaptive effects: Vibrational effects >1Hz e.g. 10kHz.

One way around these problems is to use active optic alignment or adaptive optics (AO) which perform measurement and correction of focusing and positional error in real time.

The commercial viability of such techniques is easily seen by looking at a CD player, now generally regarded as a disposable piece of machinery, which maintains focus and position of a light spot in real time on a rapidly rotating optical disk.

Dynamic Reconfiguration and Optical Interconnects - Conclusions

There are essentially two reasons why optical interconnects are specifically of interest to dynamically reconfigurable FPGAs:

- Bandwidth: FPGAs are routed dynamically and any signal must therefore traverse both switching and routing blocks, each with an associated RC delay, to reach an I/O block. Optical I/O can relieve such bottlenecks.
- Reconfiguration: To make dynamically reconfigurable computing viable, new FPGA configurations must be downloaded at a rate which puts the component out of action for the shortest possible time period.

Optical Interconnects are not meant as a replacement but as an enhancement.

Without optical interconnection, be it through free-space or waveguide, dynamically reconfigurable computing will hit a premature performance ceiling due to bandwidth limitations.





Neural Network Project

