Heriot-Watt University
Edinburgh

Optoelectronics and a Neural Network Packet Switch Controller

http://www.optical-computing.co.uk

K. J. Symington and J. F. Snowdon.

Who are we?

Collaborating with other members of the optoelectronics group at Heriot-Watt University.

We examine novel applications and architectures that can be put under the umbrella of “Optically Interconnected Computing Systems”.

Introduction

- This presentation examines an optoelectronic neural network that can be used to solve a variety of problems.
- The design and motivation for the system will be discussed.
- Results will be presented from the first generation system.
- Performance and optical system scalability issues for current hardware will discussed.
- Design and engineering issues will be examined.

Will essentially examine two demonstrators.

First generation complete and working which was designed with a specific task in mind, in this case telecomms routing.

Second generation enhances routing capability but also opens up a range of other possible applications because of its programmability.

The Assignment Problem

This problem can be found in situations such as:

- Network service management.
- Distributed computer systems.
- Work management systems.
- General scheduling, control or resource allocation.

Solving the assignment problem is computationally intensive.

Neural networks are capable of solving the assignment problem.

Their inherent parallelism allows them to outperform any other known method at higher orders.

What is our system designed to do? Solve assignment problems.

The assignment problem is essentially task allocation optimisation amongst all available resources to maximise throughput.

The assignment problem is computationally expensive and its complexity grows exponentially with problem size.

The inherent parallelism of neural networks and modus operandi allows them easily to outperform any other known method at higher orders.

We will start with a look at the optoelectronics used to create this system.

Optoelectronic Interconnects

Electrons

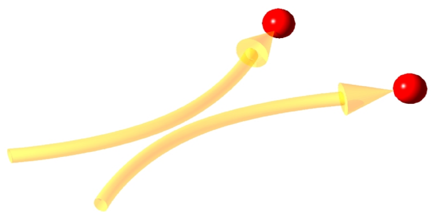

Since electrons carry mass and charge they interact strongly (Coulomb Interaction). Ideally suited for switching.

Photons

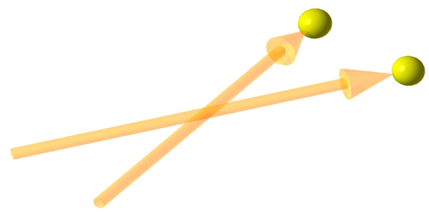

Photons do not carry mass or charge and are non-interacting in free space. They are ideally suited to interconnection.

We solve an intricate interconnection problem using the non interacting nature of photonics.

Interconnection is a problem which we have in the neural network system.

As far as switching is concerned, photonics will not be able to beat electronics for the foreseeable future.

Non-Local Interconnection

Two layer metallisation **Optical implementation**

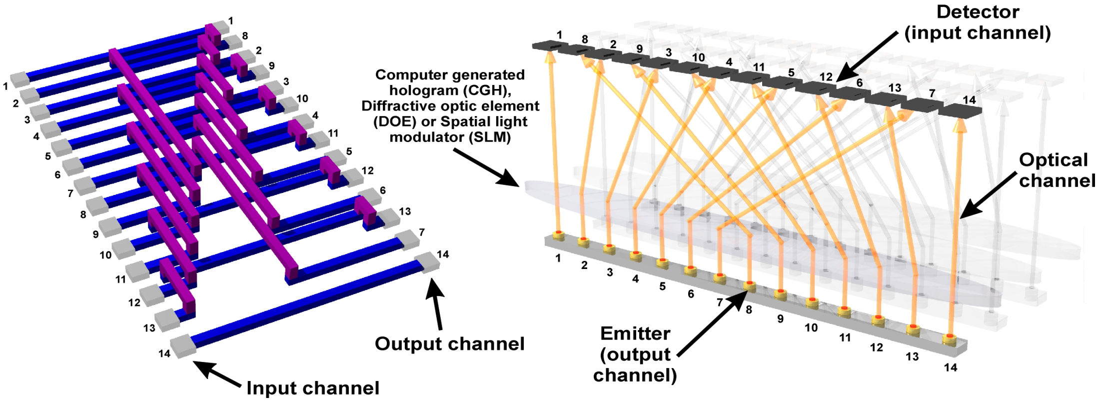

Free space optical channels can pass through each other to form any desired interconnection topology without cross-talk.

Interconnects such as the perfect shuffle and the hypercube thus become relatively simple to implement.

This also reduces skew as very large variations in wire length can be avoided.

The VCSEL Array

The VCSEL (Vertical Cavity Surface Emitting Laser) is a laser diode that emits from the surface of the substrate.

Typical speeds are >1GHz.

VCSEL array photograph
2.8mm

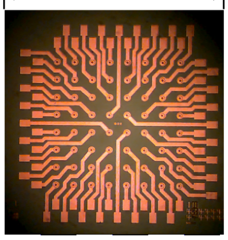

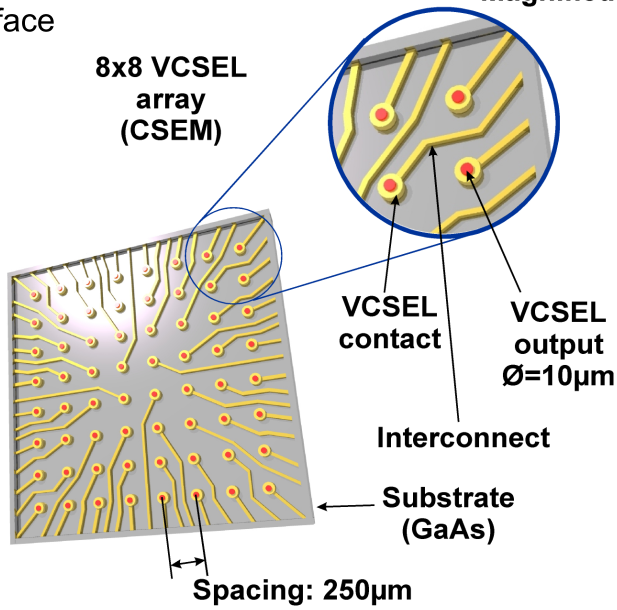

Firstly let us talk about some of the components we use in our system.

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.

Spacing is typically less than or equal to 250 microns.

Wavelengths are typically 850 to 980 nm.

Diffractive Optic Elements (DOEs)

Sample DOE

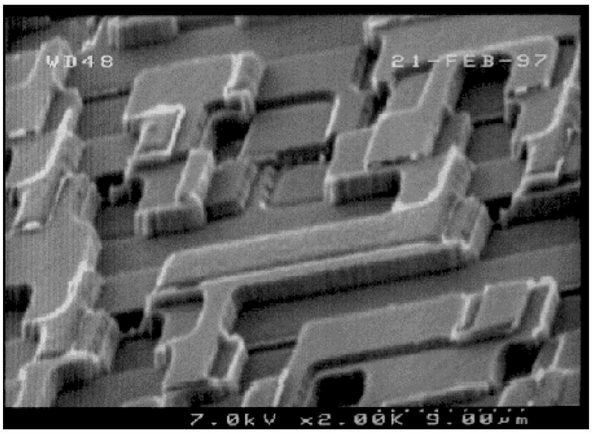

DOE Output (Single beam input)

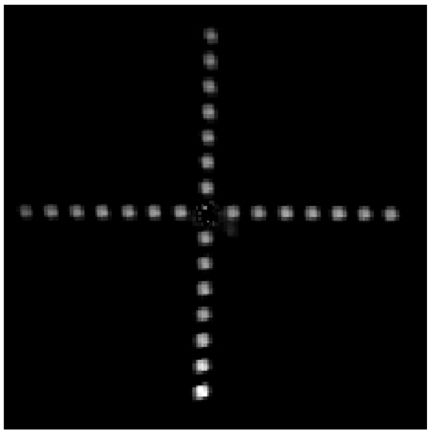

These elements are used as array generators and interconnection elements

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.

They can also be used for more complex beam shaping elements for laser material processing.

Cheap to manufacture.

Detector Arrays

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 (speed >1GHz): however the faster they are driven the more power they require.

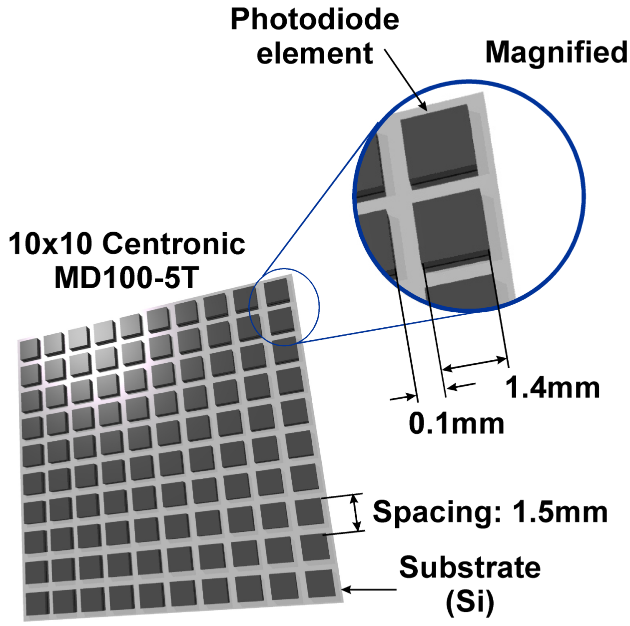

The important considerations are input sensitivity, power dissipation, area usage and speed.

The packet switch controller described here uses the device shown.

Spacing is typically less than or equal to 250 microns.

Can typically detect between 600 and 1700 nm (InGaAs).

Optoelectronic Neural Networks

- Neural networks are intractable to build to any extent on silicon.
- Free-space optics can be used to perform interconnection.
- Optoelectronics allows scaleable networks.
- Input summation is also done in an inherently analogue manner.
- Noise added naturally.

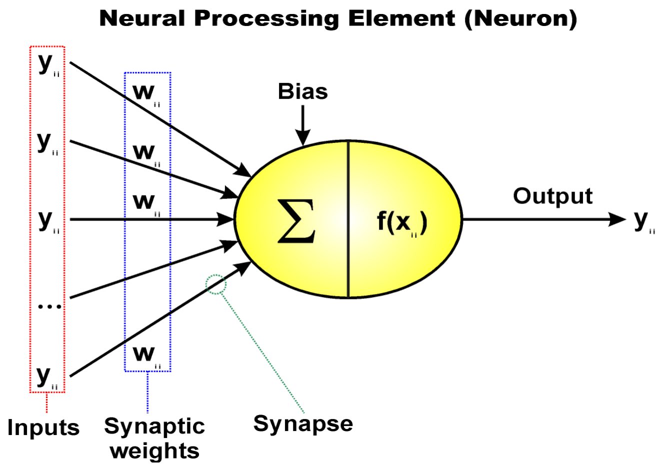

We have a good solution for the assignment problem that proves scalable in optoelectronics.

Interconnection of neural networks on silicon becomes prohibitive as the number of neurons increases. For example, in a Hopfield neural network the output of every neuron needs to be fed back to the input of every other. As network size increases, the number of interconnects increases exponentially. (Diagram)

Free space optics can be used to supply a very high connectivity by interconnecting output from an optical emitter to target detectors using a diffractive optic element.

A large amount of time can be consumed in neural networks performing weight summation. This is done in an inherently analogue manner here as the detector receives all neural activation levels optically and therefore gives an activation level proportional to the sum of all incident light.

Noise needs to be added to a neural network to stimulate convergence. This is added inherently by analogue summation.

Our optical scheme enables the deployment of scaleable neural networks. All that needs to be performed in silicon is the activation function.

In optoelectronics the summation is performed in optics.

Crossbar Switching

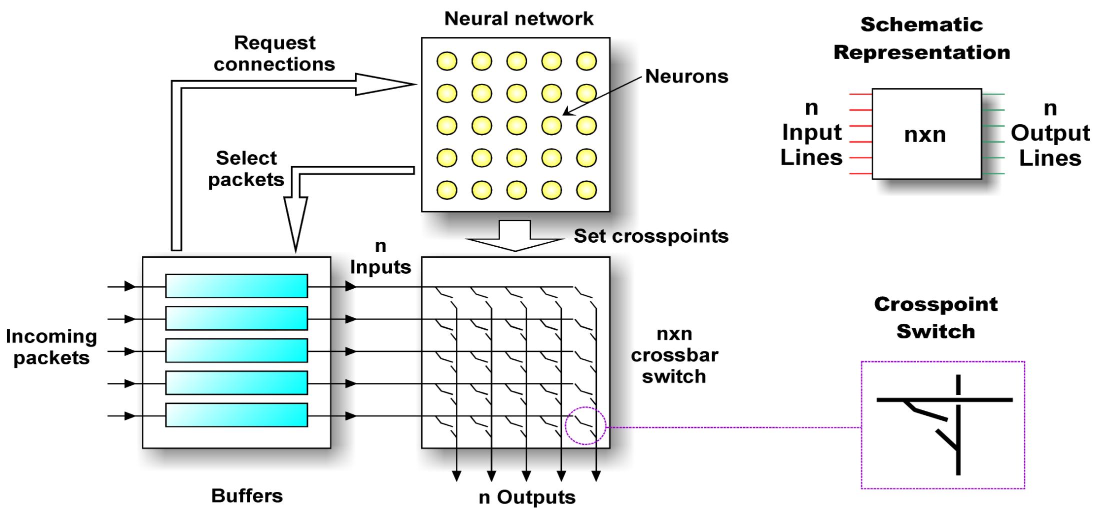

This project examines specifically the assignment problem in a crossbar switch for packet routing: maximising throughput and minimising delay.

Explain the diagram.

Crossbar switches are very common in telecommunications systems and computer networks, one good example being ATM (Asynchronous Transfer Mode) networks.

With the increase of traffic from sources such as the internet, efficient packet switching is becoming a very important issue.

Emphasise that this system is a packet scheduler and not a switch.

Algorithm

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

where:

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

y_{ij}: the output of neuron ij.

A, B and C: Optimisation parameters.

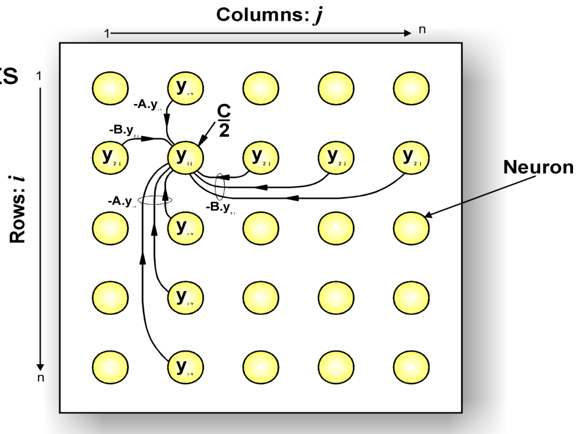

Modified Hopfield Neural Network Interconnection

This slide shows the updating rule for each neuron.

Each neuron has one detector and one VCSEL mapped to it.

Incident light on a neuron's detector acts in an inhibitory fashion as can be seen from the updating rule.

Similar to placing a rook on a chess board.

Optical Interconnect

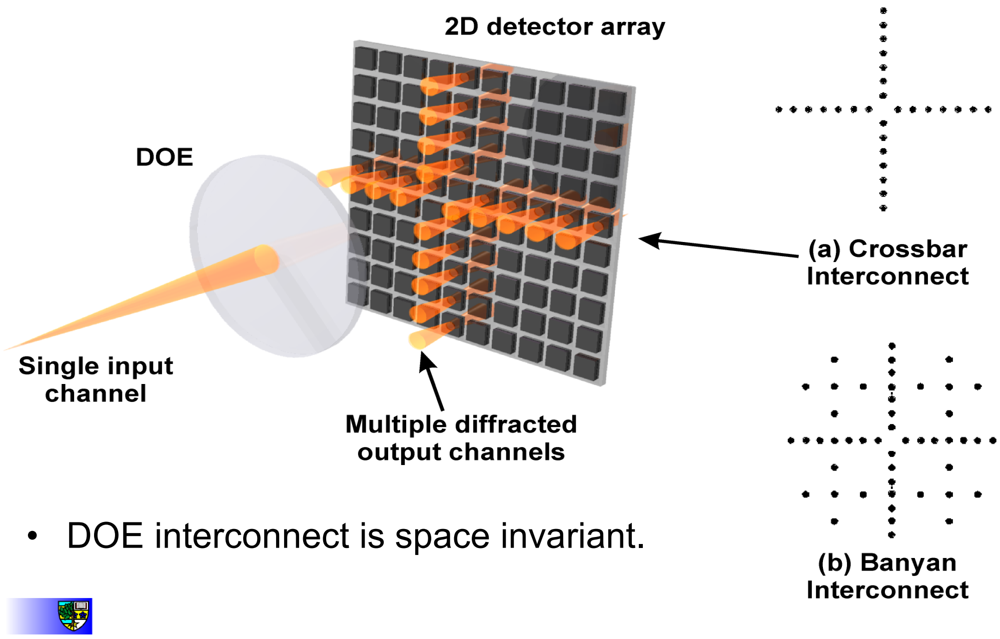

- DOE interconnect is space invariant.

Talk about interconnection method explaining how it is space invariant.

Banyan interconnect is another which is studied but for illustration purposes the crossbar switch is easier to comprehend and visualise.

Can create other interconnects to solve other problems such as chess problems.

Crossbar interconnect is similar to placing rooks on a chess board so none can take each other.

Can create DOEs for problems such as 8 queens, knights, bishops etc.

The systems strength is in allocation of resources optimally to solving a problem.

First Generation System

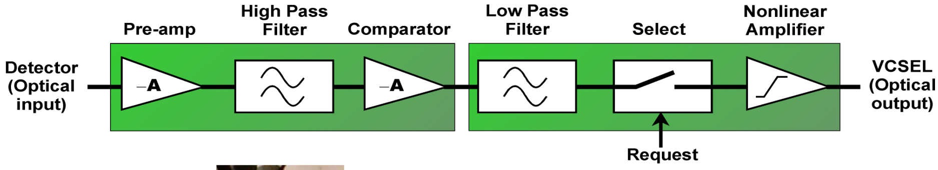

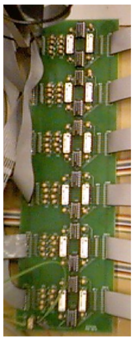

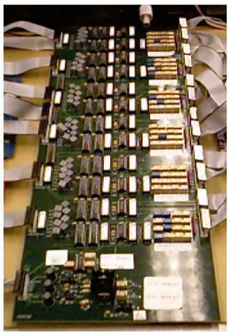

- Constructed using discrete components.
- Lacked the ability to prioritise input packets.

Explain progression in first generation system.

State that it lacked the ability to prioritise.

Optical System

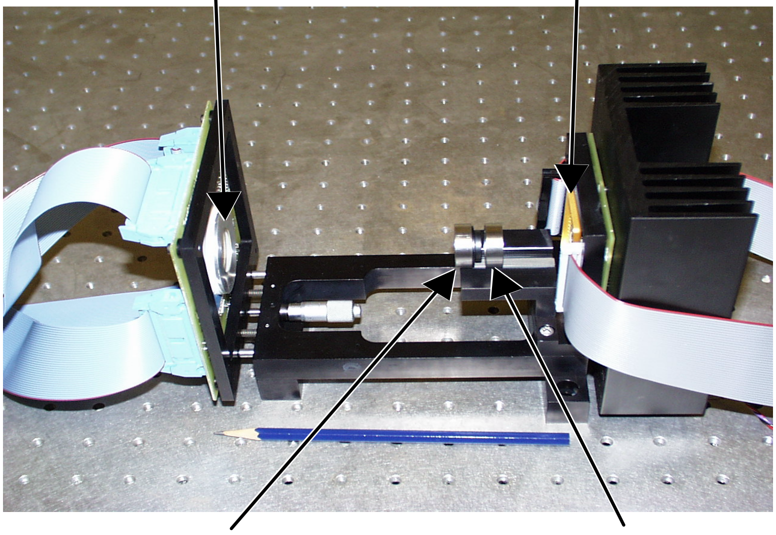

About 230mm.

The system is very simple and there is little to go out of alignment.

We don't create a lot of problems such as power loss (max 30%) or aberrations.

Note that the heat sink shown here is overkill.

This is not a 50cm pencil.

Could reduce system size to ~15cm or less by inserting a more powerful lens or by using a microlens array.

Crossbar Switch Results

- Histogram of packets routed successfully in a crossbar switch.

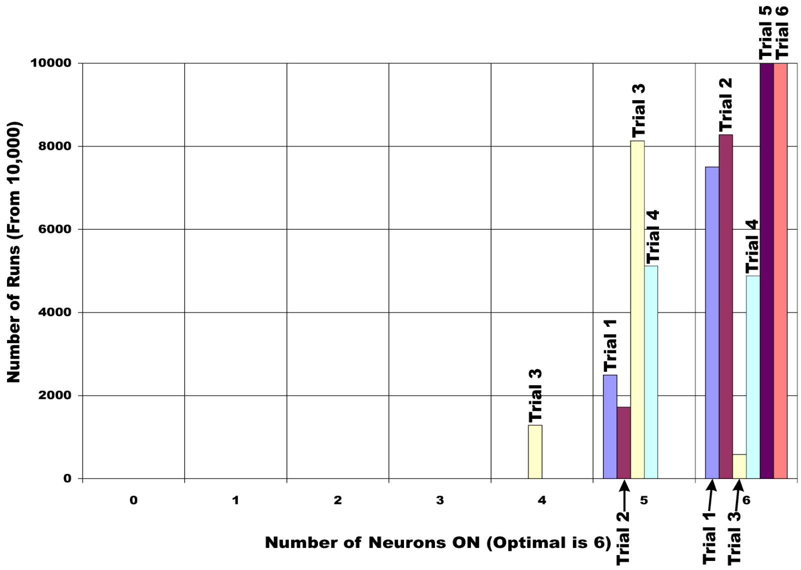

Explain histograms.

Emphasize that these are experimental results.

Banyan Switch Results

- Histogram of packets routed successfully in a banyan switch.

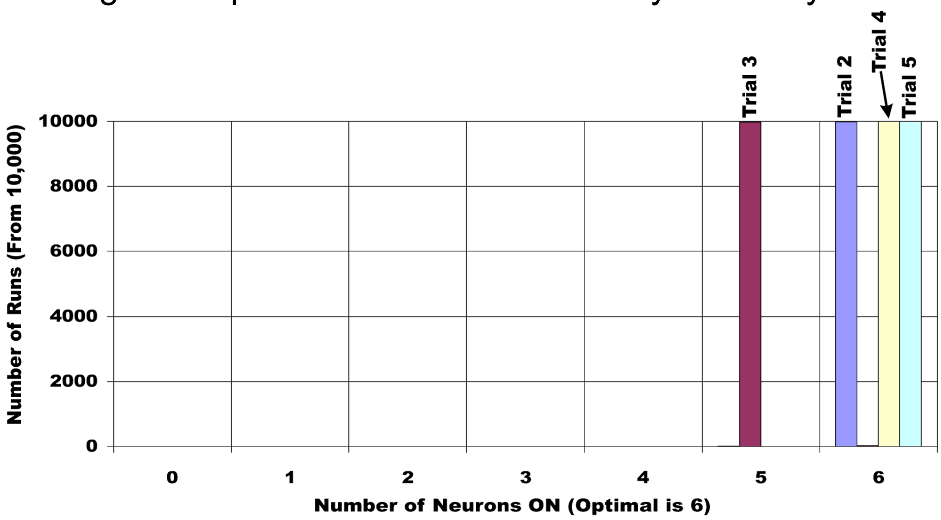

Explain histograms.

Emphasize that these are experimental results.

Performance

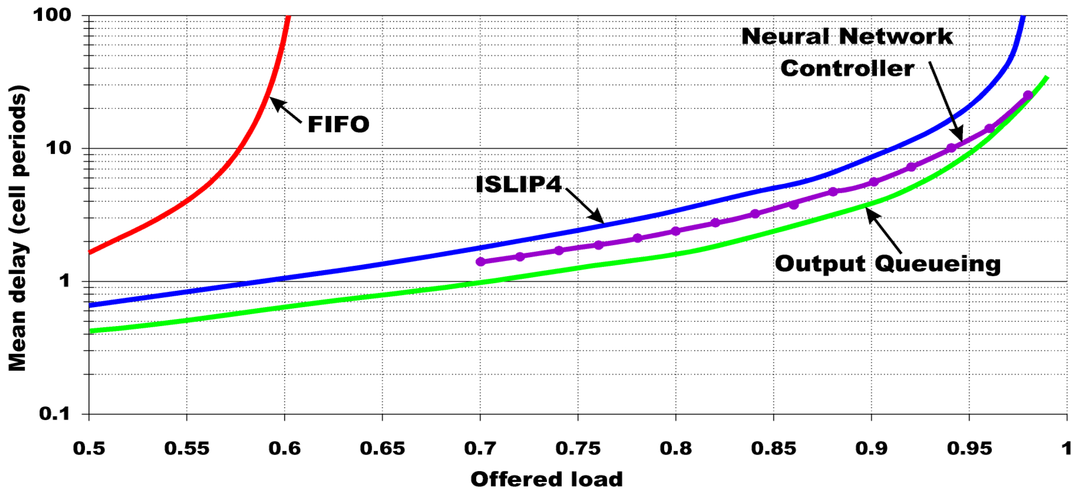

A comparison of the neural network controller against a state of the art scheduler, ISLIP4, clearly indicates its advantage at high levels of offered load. ISLIP4 cannot be implemented larger than about 16x16 because of fabrication constraints.

Not so many saturation problems.

The output queuing curve indicates a theoretical optimum value. This best result is described as output queuing and is calculated under the assumption of an ideal switching fabric in which packets have only to wait for a vacant slot on the output line.

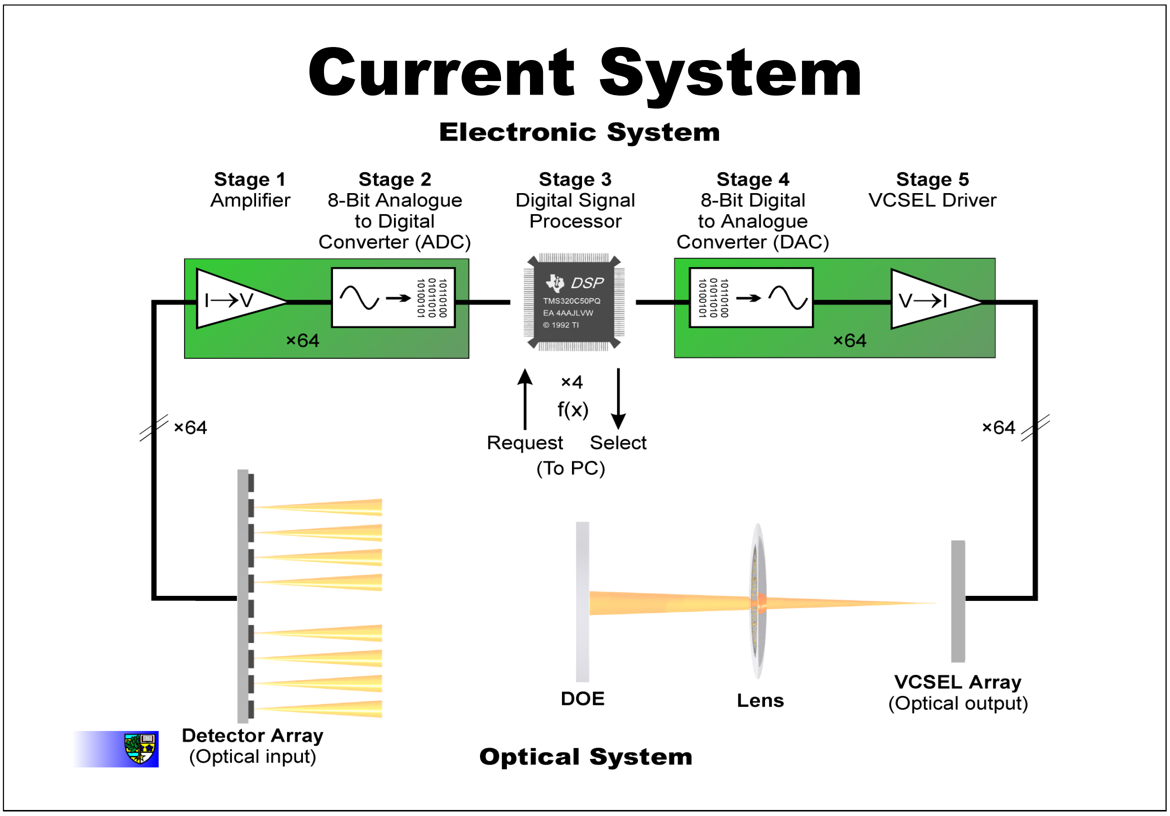

This is the second generation system. It is currently under construction. We can also drive VCSELs in a digital manner which we will discuss later. This system supports prioritisation.

Analogue Scalability

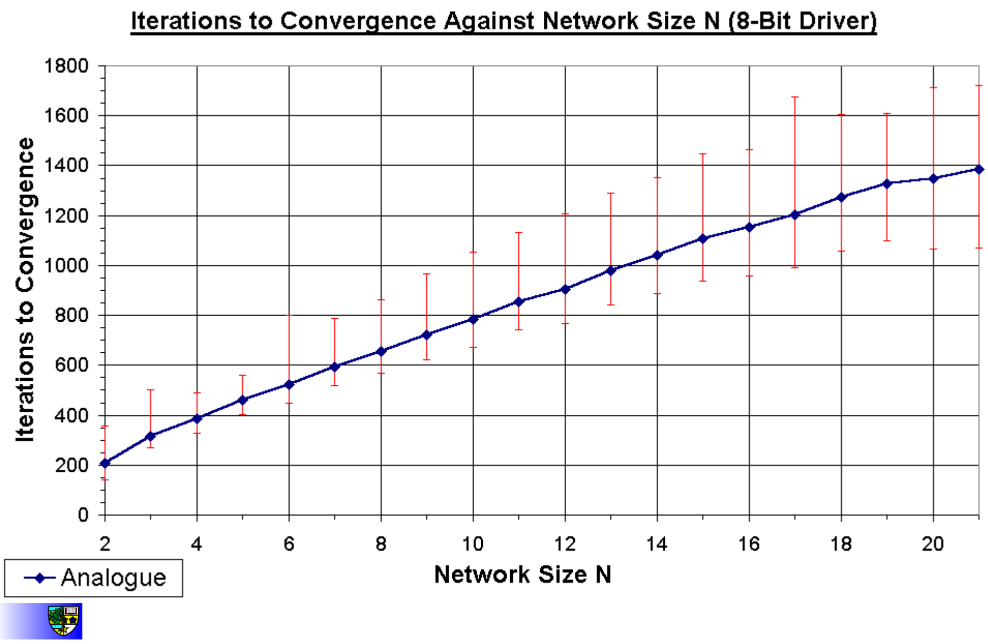

Simulation closely mapped to current component values – assessing limits of our system with current technology.

Explain the method of convergence.

Network size N is where we have an equal number of inputs and outputs (N) therefore NxN neurons.

One iteration represents a single pass through the system and is analogous to bit rate through the system.

For example, a system with components that have a maximum bit rate of 1GHz can perform up to 1Gig iterations per second.

Error bars represent minimum and maximum convergence times.

Distribution of solutions is of a guassian nature.

Maximum value on graph represents limit before the system cannot function correctly any more.

Invalid results become noticeable at N=22 or more.

Digital Scalability

Iterations to Convergence Against Network Size N (1-Bit Driver)

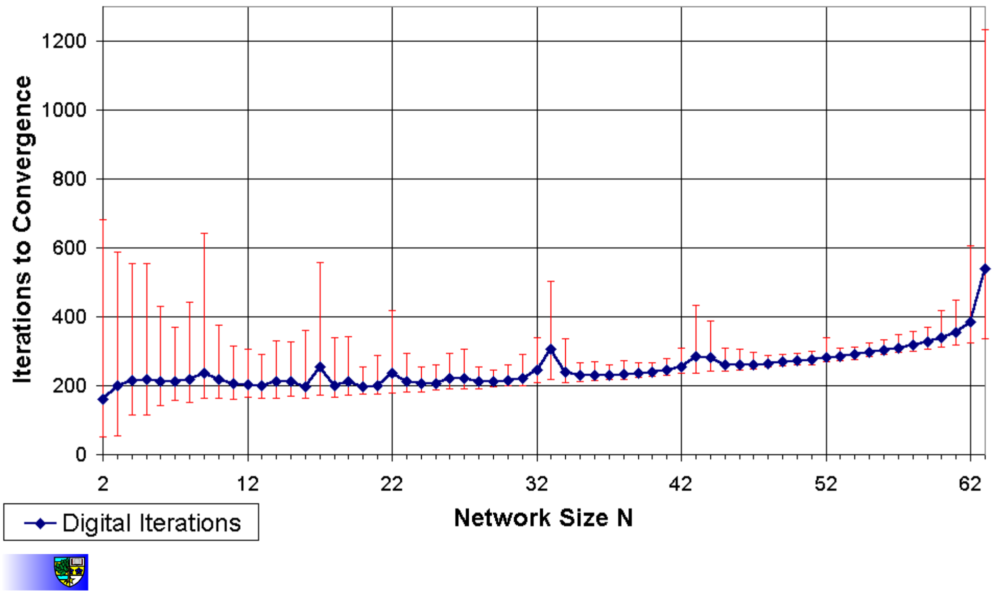

Surprisingly, a digitally driven system can be scaled better than an analogue one.

This seems to be because of impetus given to the neurons by digital driving.

Could create digitally driven system that takes the results after a generous 400 cycles.

If we had a custom ASIC that supported 1GHz then we could produce 2.5 million switch configurations per second.

Whether we have an 8 input/ 8 output crossbar or 60 input/ 60 output crossbar we can still reach a decision in 400 iterations.

Digital vs. Analogue

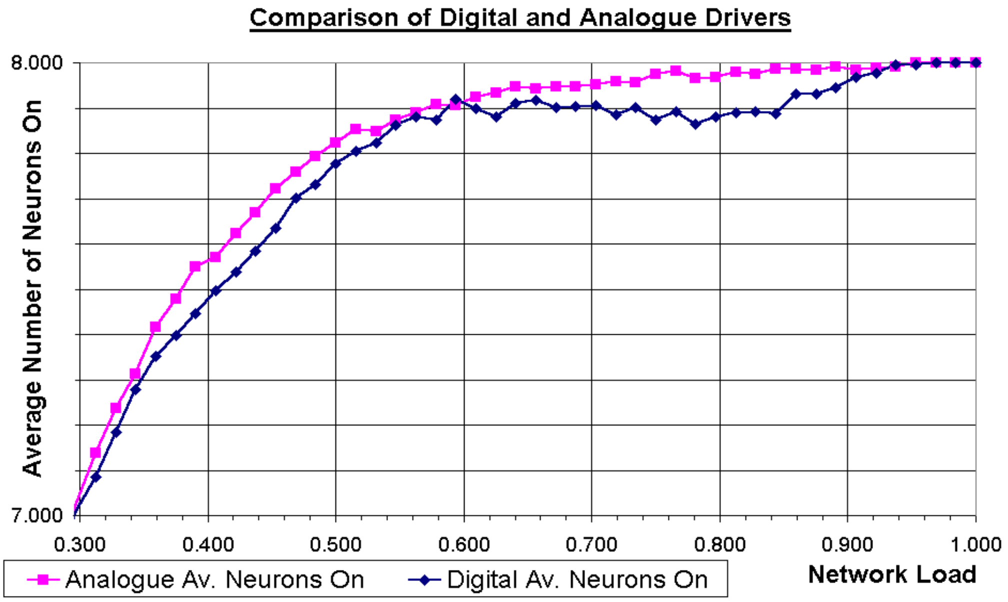

Analogue: Optimal ~97%. Digital: Optimal ~91%.

Always a price to pay – no such thing as a free lunch.

Slight reduction in overall solution optimality of N-1 neurons on rather than N neurons on.

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.

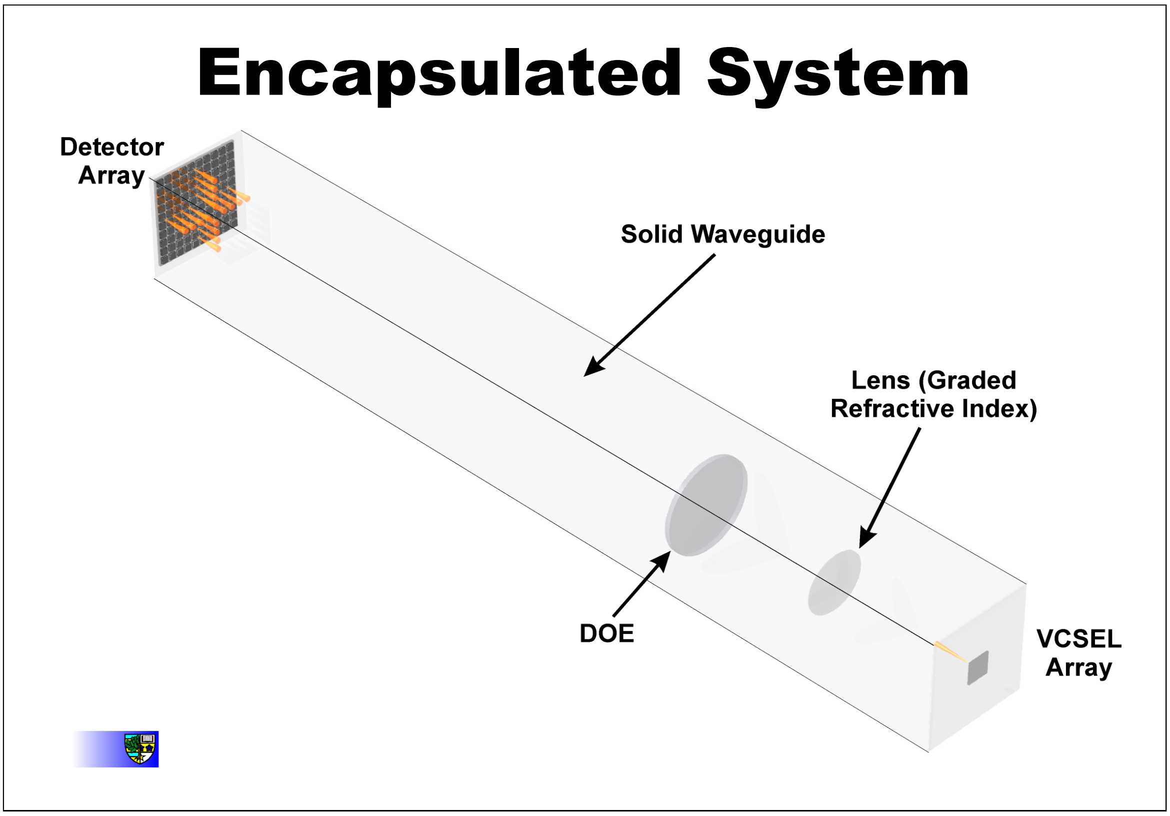

Encasing our system in an optical material would make it deployable. System would be 1.5cm square by 15cm or less in length.

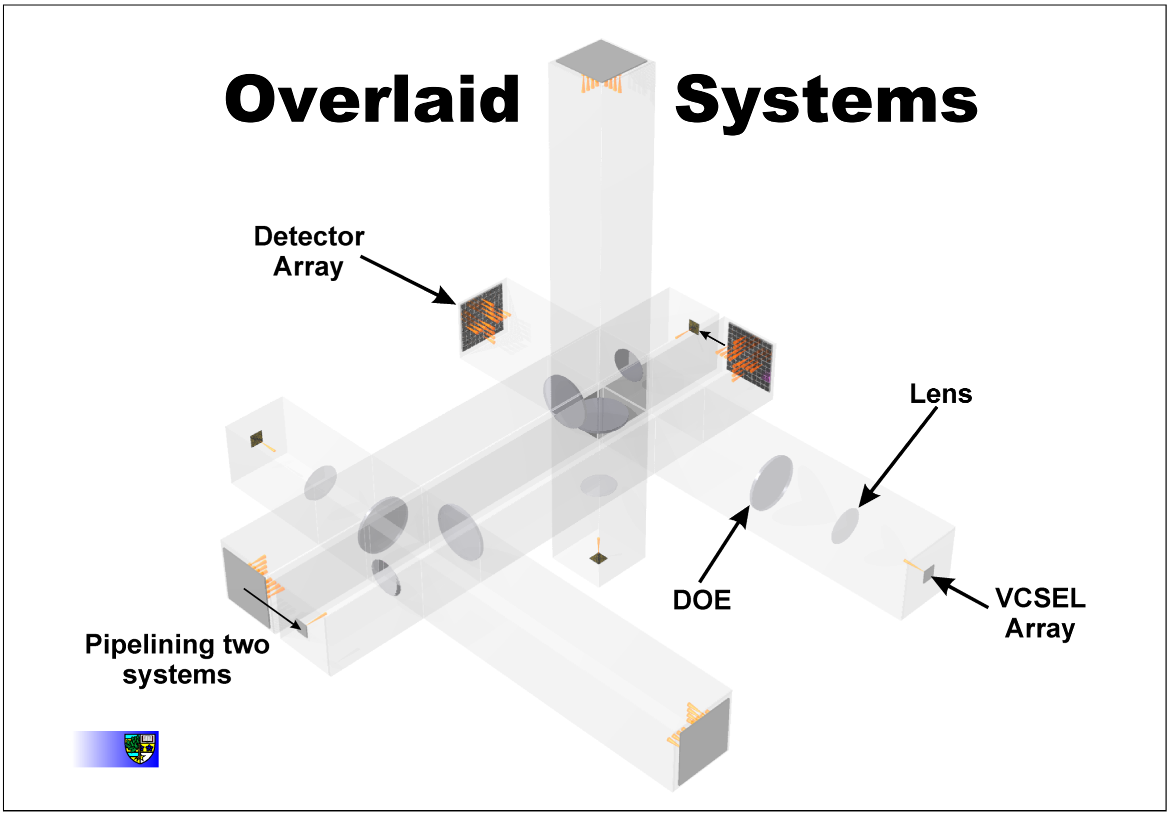

The use of optics allows multiple systems to be combined in the same space.

Wild construction of systems.

Can create pipelines of systems.

Folded System

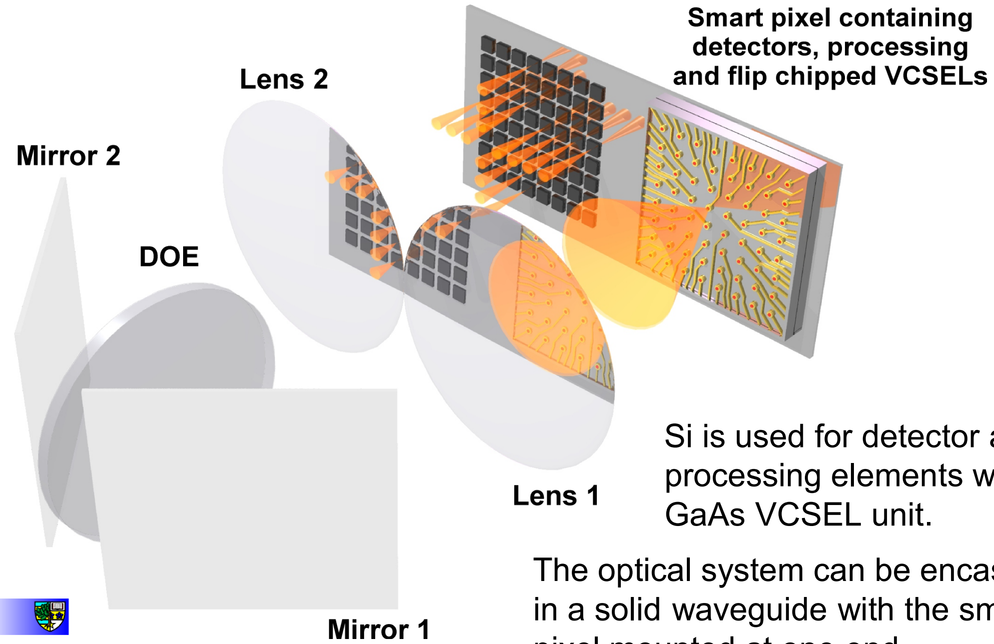

Si is used for detector and processing elements with a GaAs VCSEL unit.

The optical system can be encased in a solid waveguide with the smart pixel mounted at one end.

Could be MSM level or chip level.

Conclusions

- Second generation builds on first in that it supports prioritisation.
- Generalisation of interconnect scheme simply by replacing DOE element.
- Reconfigurability of neuron functionality by simply reprogramming DSP.
- Further work:
  - Smart pixel implementation and packaging.
  - FPGA or custom ASIC implementation using optical interconnects.
  - Novel neural algorithms and learning.

Simply change the DOE to perform a different allocation problem.

Any transfer function that you can program in C is a valid function.