

# A programmable optoelectronic neural network packet switch scheduler

K. J. Symington, Y. Randle, A. J. Waddie, M. R. Taghizadeh and J. F. Snowdon

*School of Engineering and Physical Sciences, Heriot-Watt University, Riccarton, Edinburgh, EH14 4AS, Scotland, UK.*

*kjsymington@iee.org, (A.J.Waddie, M.R.Taghizadeh, J.F.Snowdon)@hw.ac.uk*

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

**Abstract:** A programmable optoelectronic neural network architecture is presented that has been optimized to make routing decisions in both crossbar and banyan packet switch fabrics. Simulation has indicated excellent scalability in this particular application with only a minimal increase in decision time even when problem set size grows by an order of magnitude. Experimental results are presented that demonstrate a high tolerance to both noise and component inconsistencies. An assessment of system performance is made using the common metric of connections per second (CPS).

©2003 Optical Society of America

**OCIS codes:** (200.4700) Optical neural systems, (200.4260) Neural networks, (060.4250) Networks.

## 1. Introduction

This paper discusses a scalable and programmable optoelectronic neural network which is capable of solving the assignment problem. The assignment problem optimizes task allocation to ensure best use of available resources and is in essence analogous to the traveling salesman problem (TSP). It is specifically applied here to a packet switch, in this case of crossbar and banyan types, to optimize switch throughput. Indeed, this neural network solution can, in terms of mean packet delay, outperform digital schedulers such as iSLIP [1] at higher load levels.

Two generations of this system have now been completed. The first generation system was designed to prove the optical system and had a simple electronic control system consisting of an amplifier chain. The second generation demonstrator attempted to build on the first by adding prioritization, reducing system size, curtailing non-linearity, improving performance and adding reprogrammability. Its operation is based around previous hardware minimization [2] which removes a digital to analogue conversion stage thereby significantly increasing scalability.

## 2. Hardware Description

A neural network consists of a set of neurons, interconnected in an application specific manner, that perform a transfer function on the summation of a set of inputs. Here the optical system is used to perform fixed weight interconnection and summation with the transfer function calculated by a digital signal processor (DSP). Fig. 1 examines the second generation optoelectronic neural network [3].

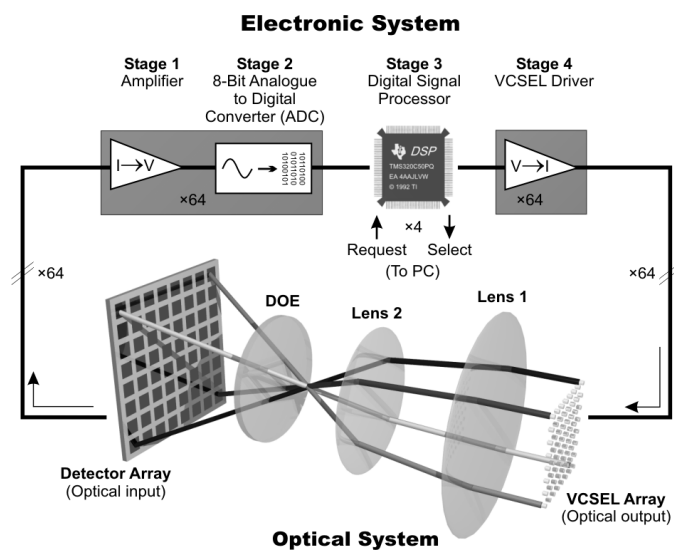


Fig. 1. System overview. This system is designed to handle an 8 input/8 output packet switch ( $N=8$ , 64 neurons in total).

There are two measures used to describe this system. The first is switch size: if a crossbar switch is of size  $N=8$  then it has 8 inputs, 8 outputs and  $N^2=64$  neurons. The second is the number of iterations to convergence and is a measure of performance. A neural network iterates until a decision is reached which can only be as fast as the slowest component. If the slowest device is capable of 10MHz then  $10 \times 10^6$  iterations are performed per second.

Each neuron requires its own optical input (detector) and output (vertical cavity surface emitting laser - VCSEL) with associated electronic hardware. The use of diffractive optics enables the scaling of complex neural interconnection patterns in a manner not possible solely in electronics. It is worth taking into consideration how different neural networks could be constructed in this way, particularly with regard to exploiting shift invariant interconnection patterns. Such patterns remain fixed regardless of where in space they are formed. The approach taken here separates this system from optical matrix vector solutions [4] by allowing two-dimensional arrays of optical emitters and detectors, thereby reducing hardware size and improving optical spatial bandwidth.

### 3. System Scalability

Preliminary design studies [2] indicated that digital operation (with corresponding component savings) was feasible. Simulation was used to build on experimental results to determine the relative performance of both analogue and digitally driven VCSELs (analogue and digital neurons respectively). Fig. 2 plots the results.




Fig. 2. Rates of convergence for varying switch sizes in both digital and analogue implementations of the system.

As network size  $N$  is increased, the number of iterations required increases linearly in the analogue system. When  $N$  becomes greater than 21, the network starts to produce invalid solutions so this was considered to be the limit of analogue scalability. With a digitally driven system, the number of iterations required remains roughly constant over its entire range until the system started to reach its limit at  $N=62$ . The final limit appears to be ADC discrimination: i.e. at low signal levels it is no longer effectively possible to tell whether a VCSEL is active or not. However, this digital thresholding adds an impetus to convergence that reduces solution optimality by 6%.

### 4. Experimental Results

To ensure that the network is operating within expected tolerances, calibration of the optical system is performed at startup. This is done by the DSPs whose flexibility enables examination and compensation for minor inconsistencies such as weak optical signal, excessive background noise and minor imperfections in the DOE. In real world situations, this process is invaluable since more serious system errors such as failed and failing components, or even complete misalignment of the system, can be rapidly diagnosed.

As with all parallel systems, synchronization is a serious issue. The first cause of de-synchronization is skew between the times that the DSPs begin calculation. If skew is high, then the time required for a single iteration increases. This is because all DSPs must reach the same point in their program which requires an arbitrary waiting time to be in-built. The second source of de-synchronization is drift across the DSP modules. Since each DSP has its own oscillator, there are unavoidable frequency discrepancies introduced during component manufacture which result in a gradual de-synchronization. An investigation of this drift showed that after one second the greatest difference between two modules was  $9.0\mu\text{s} \pm 2.7\mu\text{s}$ . Although this is not currently an issue, its effects will become increasingly pronounced as single iteration times decrease. Eventually it will become the dominant timing issue.

Experimental results with the  $N=8$  crossbar DOE show that as long as the load remains greater than 15% validity will be higher than 95%, with the number of sub-optimal results staying below 10%. Note that this corresponds to within 1% of values predicted during simulation [2]. However, an invalid solution is a situation unacceptable in real

hardware. Observation indicated that the unexpected input values occurred most frequently at the outer edges of the optical array, so the edges were isolated by using only the central array of  $6 \times 6$  ( $N=6$ ) neurons. Unfortunately this did not provide a solution and the problem was eventually tracked down to low signal-to-noise ratios that are compounded at the edge of the array. The inherent programmability in this system allowed the problem to be solved by simply reducing neural bias. Further test runs showed that validity now remained equal to 100% regardless of load. Unfortunately, validity is traded off against optimality, with an increased number of results that should have 6 neurons on having 5 instead. A simple analogy is that a rushed decision due to increased bias may result in the network missing a slightly better solution. The results are shown in Fig. 3.




Fig. 3. Although not all results are not optimal (5 or 6), more importantly they are all valid.

Analysis of the signal-to-noise ratios gave 4 for the crossbar interconnect and 2 for the banyan. Nevertheless, it is encouraging that the neural network was able to provide successful results even under such adverse conditions.

To evaluate the potential neural network performance, the accepted metric of connections per second (CPS) was used as defined by Holler in 1991 [5]. Given a maximum optical throughput of 250kHz, we can calculate the CPS rating of our neural network demonstrator as  $224 \times 10^6$  CPS in crossbar switch configuration,  $304 \times 10^6$  CPS in banyan switch configuration and  $1 \times 10^9$  CPS in fully interconnected configuration. The best performance achieved to date is 110μs for one iteration where 150 iterations are required for evolution. This equates to 16.5ms for each optimization and therefore around 60 decisions per second. In a banyan switch configuration, this gives a respectable  $11 \times 10^6$  CPS. Unfortunately, synchronization issues and a lack of computational power through multiplexing have artificially lengthened iteration times. However, simply improving synchronization will increase the number of solutions per second by at least one order of magnitude without any change at all to the current hardware. In fact, modifying the electronics to operate at the 100MHz that the optical system is currently capable of will result in an increase in performance of four orders of magnitude.

## 5. Conclusions

The second generation system presented here suffered unexpectedly from poor signal to noise ratios. This was due to a combination of electronic amplification noise, ADC conversion error and crosstalk in the optical system. All of these could potentially be minimized by design in any subsequent system, but reprogrammability has allowed adaptation of the neural network to operate well even in such an environment. Simulation indicates that the system unsurprisingly breaks down as signal-to-noise approaches 1, however we have experimentally demonstrated successful operation with a signal to noise ratio of 2.

Finally, an important measure that has not been made is quality of service (QoS) using accurate network traffic models. Such analysis will conclusively define the future of the optoelectronic neural network scheduler.

## 6. References

- [1] R. P. Webb, A. J. Waddie, K. J. Symington, M. R. Taghizadeh, and J. F. Snowdon, "An optoelectronic neural network scheduler for packet switches", *Applied Optics*, vol. 39, no. 5, pp. 788-795, Feb. 2000.
- [2] K. J. Symington, A. J. Waddie, M. R. Taghizadeh and J. F. Snowdon, "A neural-network packet switch controller: scalability, performance, and network optimization", *IEEE Trans. On Neural Networks*, vol. 14, no. 1, Jan. 2003.
- [3] A. J. Waddie, K. J. Symington, M. R. Taghizadeh and J. F. Snowdon, "An Optoelectronic Neural Network: Implementation and Operation" *Optics in Computing*, Quebec City, Canada, R. A. Lessard and T. Galsitan, vol. 4089 of OSA Proceedings Series, pp. 304-310, 2000.
- [4] K. Ballüder, and M. R. Taghizadeh, "Optimised phase quantisation for diffractive elements using a bias phase", *Optics Lett.*, vol. 24, no. 23, pp. 1576-1578, 1999.
- [5] M. A. Holler, "VLSI Implementations of Learning and Memory Systems: A Review", *Advances in Neural Information Processing Syst.* 3, San Mateo, CA, Morgan Kaufmann, pp. 993-1000, 1991.