A Neural Network Scheduler for Packet Switches

R. P. Webb, A. J. Waddie, K. J. Symington, M. R. Taghizadeh and J. F. Snowdon*
Heriot-Watt University, Edinburgh, EH14 4AS, Scotland.
** BT Laboratories, Martlesham Heath, Ipswich, IP5 3RE, UK.*

At the core of some of the latest generation of internet routers is a hardware switch that transports packets between the line cards. A central scheduler is required to select a set of packets from queues on the line cards that can be connected to the correct outputs simultaneously without blocking [1]. The larger the set chosen, the greater the throughput, but the decision must be made within the cycle time of the switch. This assignment of outputs to inputs subject to constraints imposed by the switch fabric is an example of a resource allocation problem which can be solved by a Hopfield neural network [2, 3]. We have implemented a Hopfield network as a parallel optical system incorporating a diffractive optical element (DOE) and measured its performance as a scheduler for both crossbar and self-routing switch fabrics.

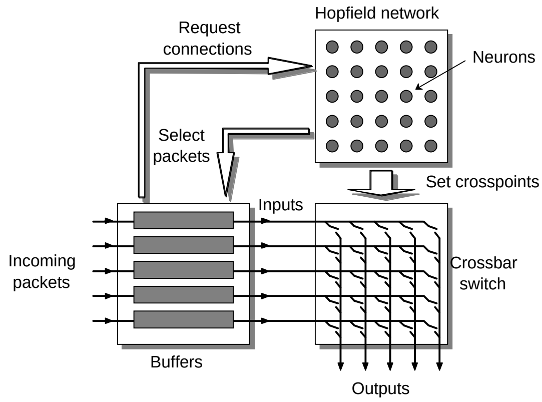

Fig. 1. *Packet scheduling and switch control.*

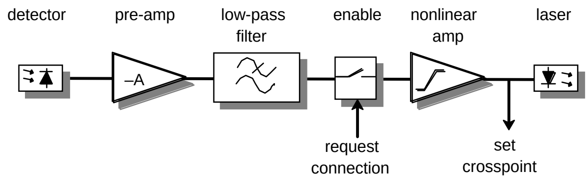

Fig. 2. *Neural circuit.*

A two-dimensional array of neurons represents all possible input to output connections. In the case of a crossbar, the neurons correspond directly to the crosspoints of the switch (Fig. 1). The neuron outputs can vary continuously between the *off* and *on* levels. In order to choose a set of connections, the neurons representing all the requested connections are enabled simultaneously and set to the same intermediate level. Each has a bias input that tends to increase its output, but also receives inhibitory inputs from those neurons which represent blocking connections. Crossbar switches can be blocked at their inputs and outputs only, so the neurons are arranged to be inhibited by others in the same row or column. Other types of switch can be blocked internally, for these switches additional inhibitory connections are provided between the appropriate pairs of neurons. The dynamics of the network resolve the conflicts between all the mutually excluded neuron pairs, leaving a valid set of neurons in the *on* state and the remainder *off*.

In this implementation, each of the 48 neurons has an input detector followed by a capacitor-coupled inverting amplifier chain and a low-pass filter, and the output drives a vertical-cavity surface-emitting laser (VCSEL) (Fig. 2). Initially all the lasers are set to a fixed output level, slightly higher than the *off* level. This sets a stable total power for the array and effectively biases the neurons towards the *on* state. When the network is enabled, the lasers of all the requested neurons are connected to their amplifier outputs and the others are set to the *off* level. Between the laser and detector arrays are a pair of lenses and a DOE that divide the light from one neuron's laser and focus it onto the inputs of the other neurons in the same row and column, but not its own input (Fig. 3). Because of the inversion in the amplifier chain, light falling on a detector inhibits that neuron, decreasing its output.

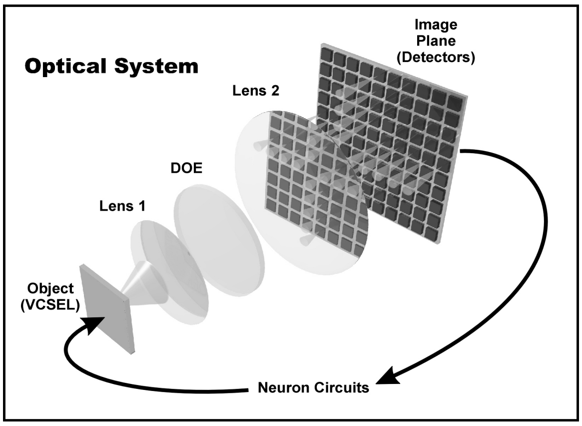

Fig. 3. *Optical system.*

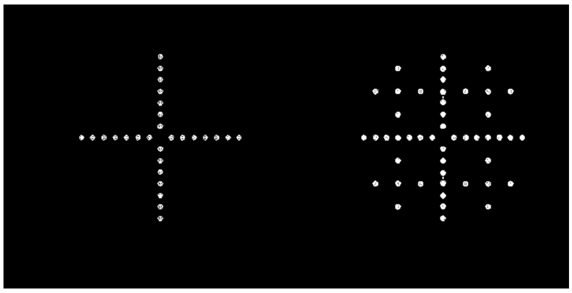

Fig. 4. *DOE far-field patterns for crossbar (left) and SR switch schedulers.*

For a crossbar switch, the pattern of neurons inhibited by a given active neuron is shift invariant. That is, it remains the same relative to the position of the active neuron. (The pattern overlaps the array.) Shift invariance makes optical implementation especially efficient because the same optical system provides independent paths from all lasers to the appropriate detectors. An electrical system would require a separate wiring network for each output. Self-routing (SR) switches can also be scheduled by neural network if additional inhibitory paths are provided to prevent internal blocking at intermediate switch stages [4]. These paths can also be made shift invariant if the switch outputs are rearranged, for example, in bit-reversed order.

DOEs were fabricated for scheduling both crossbar and self-routing switches and their far-field patterns are shown in Figure 4. In order to ensure the correct operation of the neural network, the far-field outputs were constrained to have a non-uniformity of less than 2% *after fabrication* [5]. This low non-uniformity could only be achieved by sacrificing efficiency during the design optimisation stage [6]. The degree of freedom in the design procedure was further limited by the small (~96μm) DOE period required to produce the correct order separation in the detector plane. The theoretical efficiencies (η) and non-uniformities (Δr) of the resulting designs shown in Table 1. The DOEs were binary resulting in a single fabrication step reducing the non-uniformity caused by multiple fabrication steps (estimated at ~1-2% per step).

	η	Δr
Crossbar	50.0%	0.81%
SR switch	49.9%	0.83%

Table 1. *DOE theoretical efficiencies and non-uniformities.*

The scheduler was first tested with the crossbar DOE in position. Twelve sets of requested neurons representing different combinations of waiting packets were prepared. For each test, the appropriate neurons were enabled and Fig. 5 shows typical waveforms observed at the outputs. Initially the outputs were at about half the maximum power, so most of the detectors were receiving large amounts of light from conflicting neurons (Fig. 6). The outputs began to turn *off* at a rate governed by the low-pass filters. When their inputs were low enough, some neurons were able to turn *on* again, increasing the inhibition of those in conflict with them. After approximately 200 μs all the neurons were clearly either *on* or *off*. Each request set was presented 10 000 times to the scheduler. The validity of every output state was checked and the number of neurons *on*, representing the number of packets routed through the switch, was counted. The results for the request sets where the maximum allowable number of neurons *on* is 6 are shown in the form of a histogram (Fig. 7). The experiment was repeated with the self-routing DOE and the same request sets. In some cases the maximum allowable number of neurons *on* was reduced by the extra blocking mechanisms in the switch, but those results for which the maximum was still 6 are shown in Fig. 8.

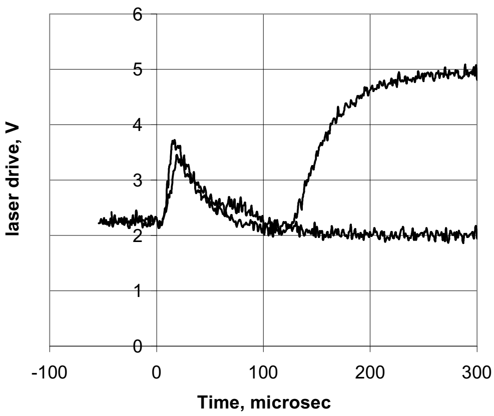

Fig. 5. *Outputs of neurons turning on and off.*

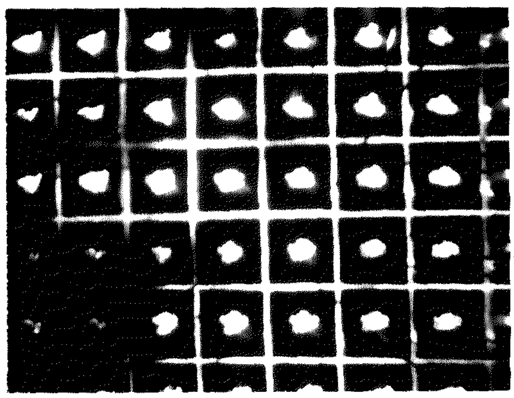

Fig. 6. *Detector array (part) with all VCSELs on.*

The scheduler never produced an invalid result. Most times it found an optimal result except with the requests *TRIAL6* and *TRIAL8* (crossbar only). With these it usually routed one fewer packets. Thus the switch would have near maximum throughput and never block. No attempt has been made to make this demonstration system run fast. The decision time was 6x the time constant of the low-pass filters (33 μs). By reducing this time constant, it should be possible to obtain results in tens of ns and achieve scheduling decisions at a rate compatible with the latest router requirements. Designs have now been produced for a smart-pixel based system to lower this time constant and such a system is currently under construction. This system will have increased functionality in that it is able to prioritise input queues. Details of the new system will be presented at the conference.

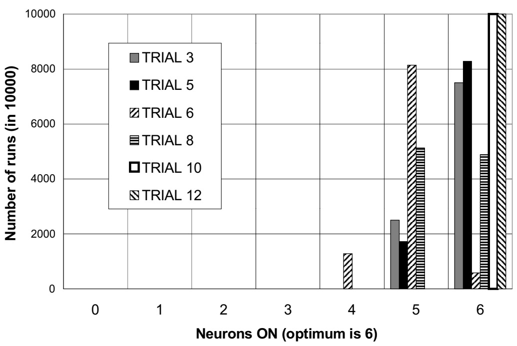

Fig. 7. *Histogram of packets routed in the crossbar switch.*

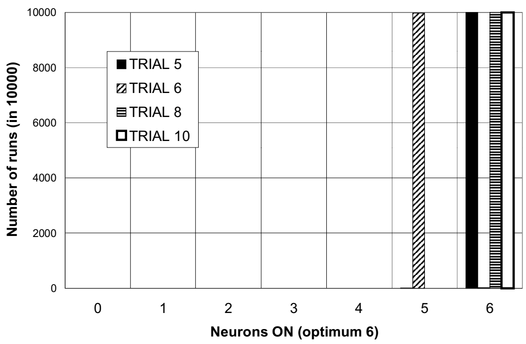

Fig. 8. *Histogram of packets routed in the self-routing switch.*

1. "Tiny Tera: a packet switch core," N. McKeown et al.,
2. "Neural computation of decisions in optimization problems," J. J. Hopfield and D. W. Tank, *Biological Cybernetics,* **52**, pp. 141-152, 1985.
3. "A neural network approach to studying the switching problem," S. J. Amin, S. Olafsson and M. A. Gell, *BT Technology Journal,* **12**, No. 2, pp. 114-120, April 1994.
4. "Neural networks for switching," T. X. Brown, *IEEE Communications Magazine,* pp. 72-81, November 1989.
5. "Synthesis of High Efficiency Diffractive Optical Elements", J.M.Miller, Ph.D. Thesis, Heriot-Watt University, 1993.
6. "Binary surface-relief gratings for array illumination in digital optics", A.Vasara, M.R.Taghizadeh, J.Turunen, J.Westerholm, E.Noponen, H.Ichikawa, J.M.Miller, T.Jaakkola, S.Kuisma, *Applied Optics* **31** No.17 pp. 3320-3336, 10 June 1992.