

Analysis and Modelling of a 2D Optical Interconnect for Parallel Computing

AMOS Project

G. A. Russell, J. F. Snowdon, T. Lim, A. C. Walker, J. J. Casswell, K. Symington


(Department of Physics, Heriot-Watt University)

P. M. Dew, I. Gourlay, K. Djemame

(Informatics Research Institute, University of Leeds)

Why Optics?

- Inherent parallelism
- Higher temporal and spatial bandwidths
- Higher interconnection densities
- Less signal cross-talk
- Immunity from electromagnetic interference and ground loops
- Freedom from quasi-planar constraints
- Lower signal and clock skew
- Lower power dissipation
- Larger number of fan-ins fan-outs
- Potential for reconfigurable interconnects





‘**Optical Highway**’ carrying hundreds of thousands of channels linking non-local nodes with thousands of channels

Technologies - Free Space Optics

Polarisation Beam Splitters (PBS)

Used to route data streams into and out of Optical Highway.

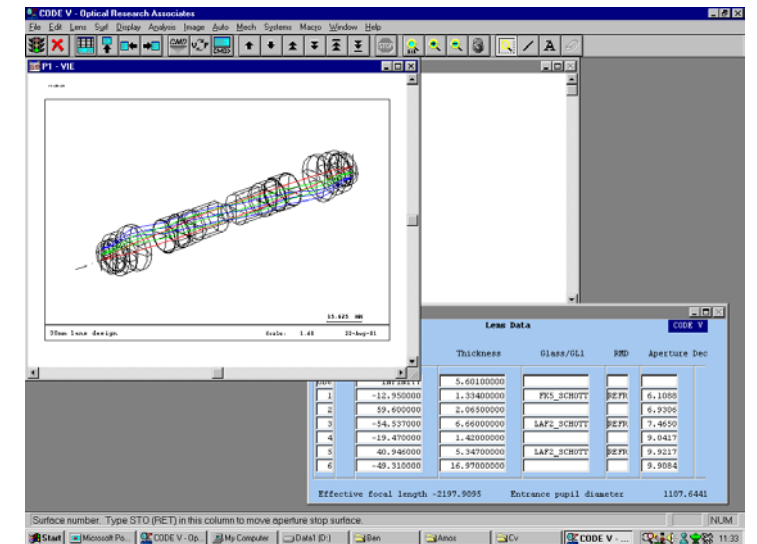
Patterned Birefringent Plates are used to selectively flip polarisations of the beams to be routed out of the Optical Highway.



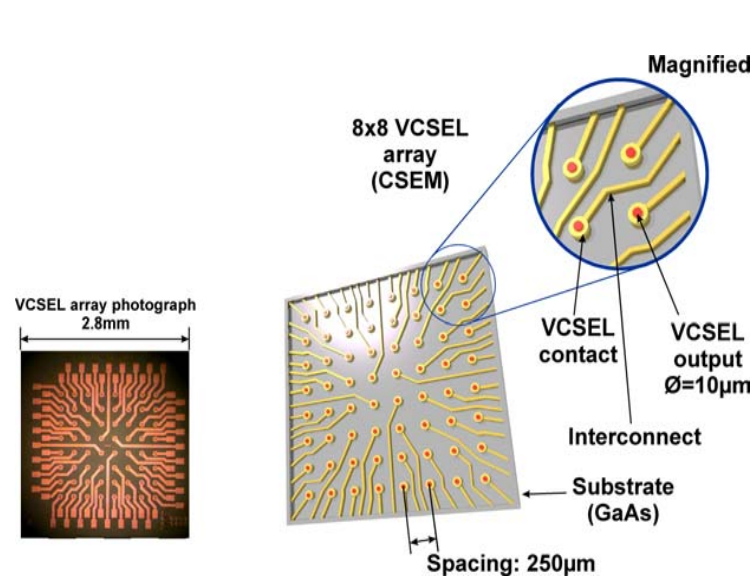
Lens

Combination of bulk and micro lenses could be used

Code V simulations to insure aberration limited spot size small enough to eliminate cross-talk between channels. Model developed to generalise to multiple relay stages.




Technologies - Smart Pixel Array




VCSEL - Light Source

Arrays of optical components fabricated in a suitable optically active material are 'flip-chip' bonded to a silicon processing layer.

Modulators have a number of advantages for bonding to silicon especially in thermal dissipation.




Modulator - Data Transmission




Detector - Data Reception

Silicon layer may be conventional VLSI silicon or FPGA (Field Programmable Gate Array) for reconfiguration.



Silicon can be used for routing or fine grained processing




Physical Layout for a 2D Interconnect



- Data channel hard 'wired'
- Polarisation used to control optical path
- Wrap-around links used at ends of each dimension to condense package and reduce optical path length
- 18 node layout shown below



Experimental Optics Set-Up



Proof-of-concept set-up to show polarisation routing used to send signal beams from VCSEL array to CCD cameras.

Used to help develop models of the optics and obtain parameters used in later modelling of the system.

Modelling



- Analytical model
- Equal Bandwidth Methodology
 - Longer length links require more channels as bandwidth decreases with optical loss
 - Requires multiple clock speeds across optoelectronic interface chip
- Link bandwidth assume to be receiver limited
- SPA silicon only used as routing / traffic controller
- Simulation / Experimentation required to validate model

$$B_{optical} \leq \frac{1}{4} f_o N \left[\frac{\xi^X (1 - \xi)}{1 - \xi^X} \right]$$

$$L_{optical} \approx \frac{q\sqrt{p}}{2c} + 40 \times 10^{-9}$$



Comparison of Results




2D Optical Interconnect compared with analysis of 1D Optical Interconnect and Aspect limited electronic interconnect assuming equal silicon area as basis.



Computational Issues of 2D Architecture

Architectural Issues

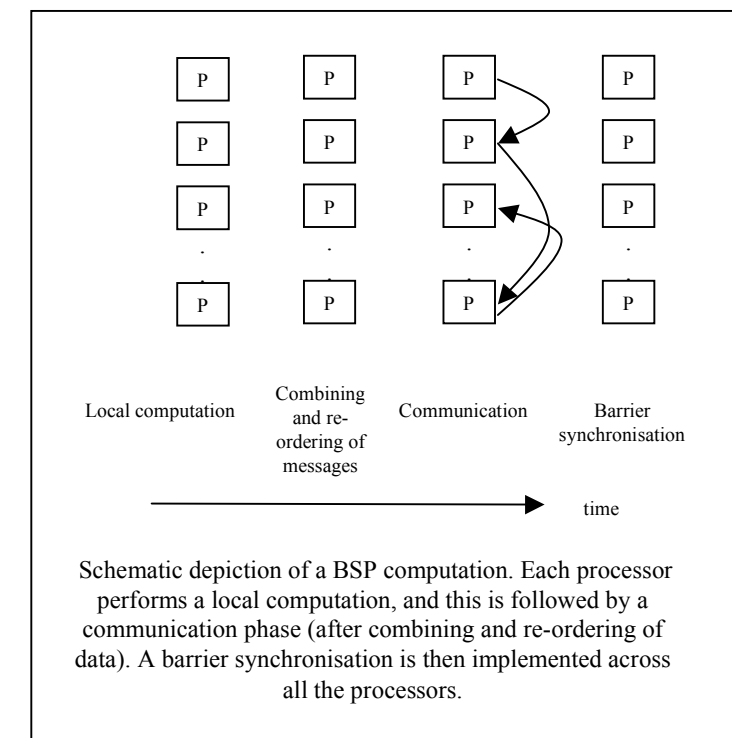
- High connectivity (e.g. completely connected network or hyper-cube like), below
- Large buffers in smart-pixel layer



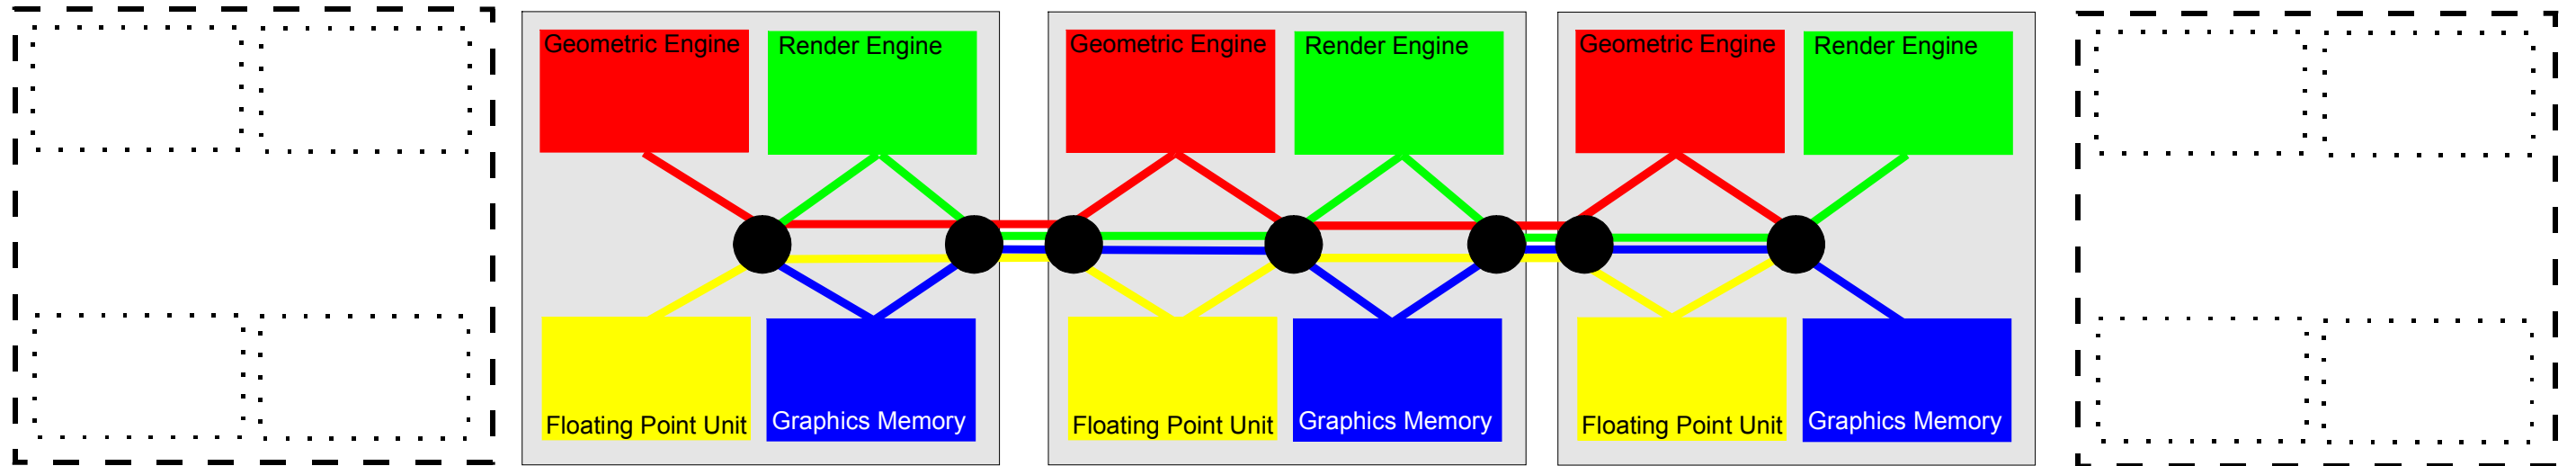
Initial performance results, based around an implementation of the Bulk Synchronous Parallel model indicate potential for substantial performance enhancement for communication intensive problems. Examples of such problems are parallel sorting, graphics and FFTs.

Algorithmic issues

- Algorithms designed to allow communication of large messages (to exploit optical bandwidth)
- Messages collected in smart-pixel layer prior to inter-processor communication.

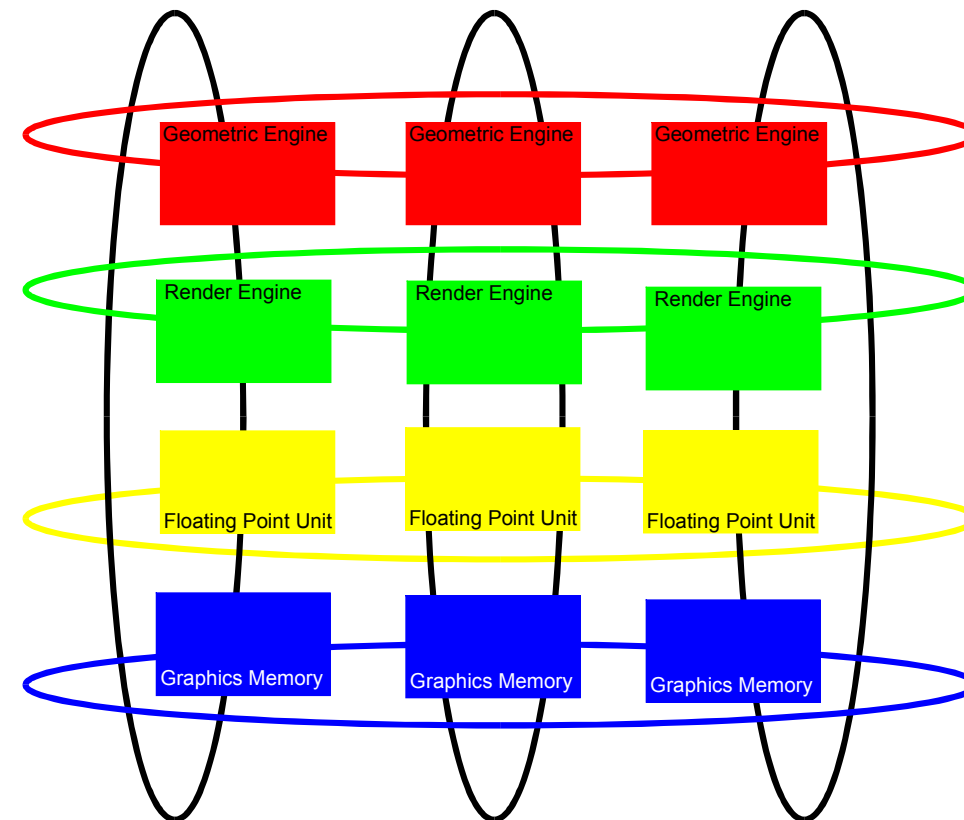


Graphics Subsystem - An Application



Processor Nodes are now the Geometric Engine, Render Engine, Floating Point Unit and other specialised processors. Inter-processor interconnect is also 'memory bus'.

Optics connects all nodes / engines of each type in one dimension and all nodes with in same unit (one of each node) in the other.



High Bandwidth Memory Bus - A Second Application


Memory interleaving

Memory interleaving is a standard technique to increase memory bandwidth by using a large number of memory banks and utilising the aggregate bandwidth. This becomes very important as processor speeds increase or if multiple processors share memory.

Problem

The large number of channels required makes this impractical using an electronic solution.

Using an out-of-plane optical interconnect the channels can be supported.

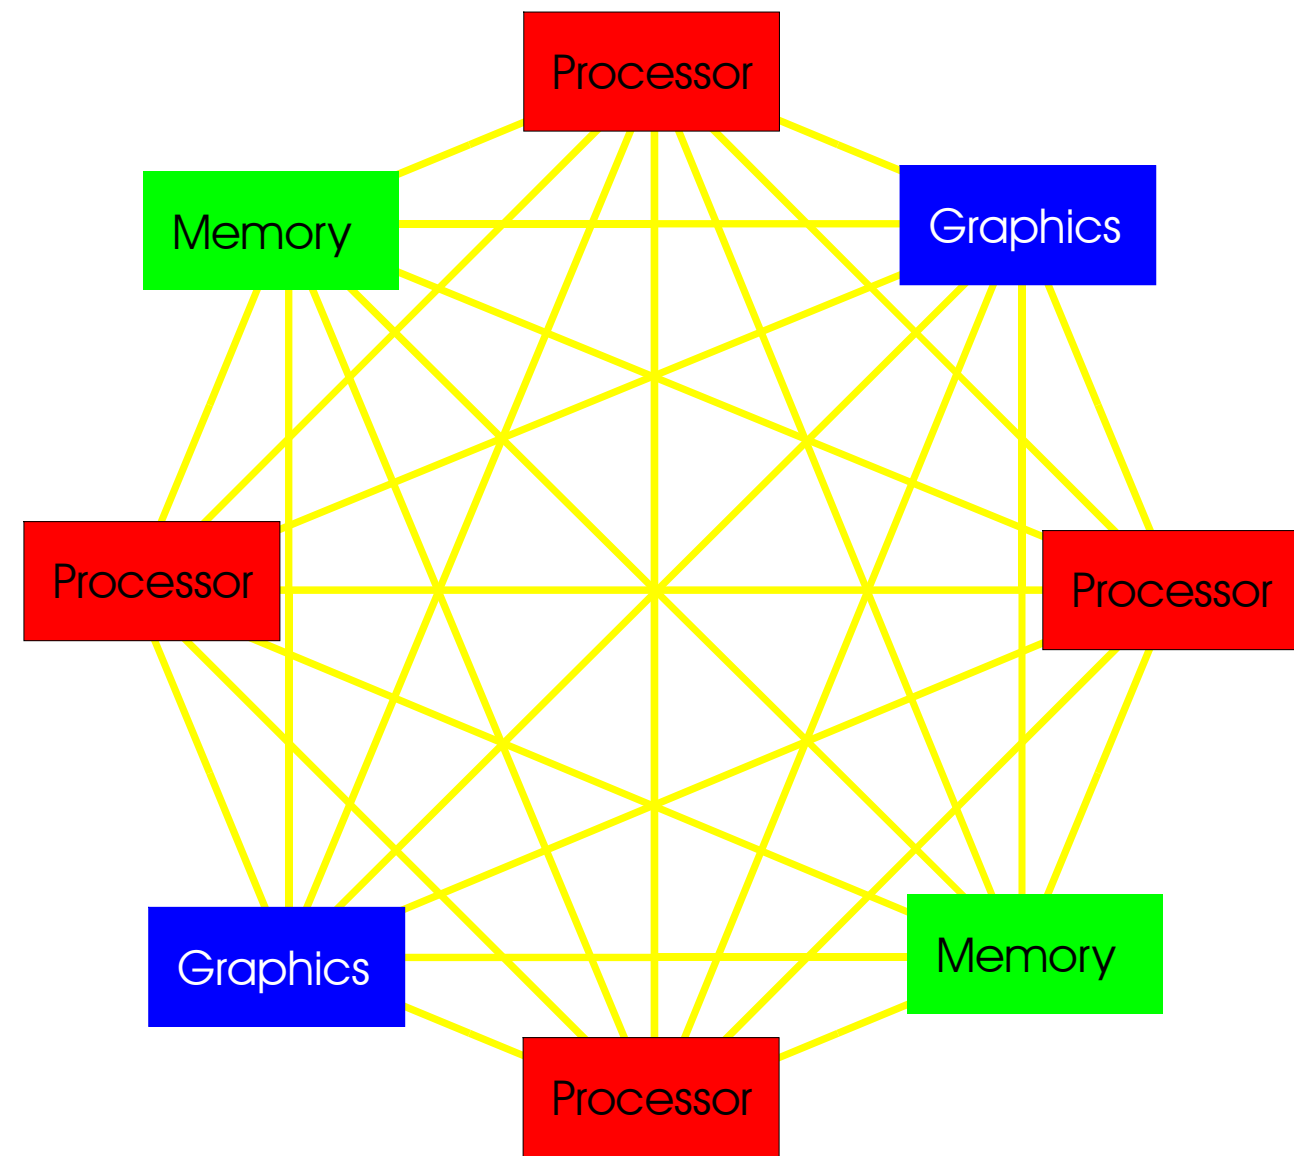


A Massively Parallel, Optically Connected Computer - Bringing it all Together

Parallel computer architecture including possibly thousands of processors, memory arrays, graphics pipelines and disk arrays.

The individual nodes maybe internally optically connected, e.g. the graphics nodes based on the design previously.

Optical back-plane directly interfaces with internal component optics and external telecommunications optics. This would create the processing resources for projects such as the 'Grid'.



Conclusions, Research Areas and Further Work

- 2D interconnect can provide high bandwidth, highly scalable communications for massively parallel computers.
- Latency penalties from optics to electronics domain changes are offset by smaller routing costs of the more direct connections available optically.
- High bandwidth required for communications bounded problems and to service the data requirements of faster processing.

Research Areas and Further Work

- Large, reliable arrays of optoelectronic components ($>128 \times 128$).
- Integration with silicon technologies (bonding techniques or silicon devices). OFPGAs.
- Efficient thermal engineering.
- Simulation of real traffic patterns on interconnect.
- Topology of interconnect and components.



*Department of
Physics*

