

Emulating a large memory with a collection of smaller ones*

James Hanlon

June 2015

Sequential computation is well understood but does not scale well with current technology. Within the next decade, systems will contain large numbers of processors with potentially thousands of processors per chip in order to deliver continuing growth in performance. Despite this, many computational problems exhibit little or no parallelism and many existing formulations are sequential. It is therefore essential that highly-parallel architectures can support sequential computation by emulating large memories with collections of smaller ones, thus supporting efficient execution of sequential programs or sequential components of parallel programs. This paper demonstrates that a realistic parallel architecture with scalable low-latency communications can execute large-memory sequential programs with a factor of only 2 to 3 overhead, when compared to a conventional sequential memory architecture. This overhead seems an acceptable price to pay to be able to switch between executing highly-parallel programs and sequential programs with large memory requirements. Efficient emulation of large memories could therefore facilitate a transition from sequential machines by allowing existing programs to be compiled directly to a highly-parallel architecture and then for their performance to be improved by exploiting parallelism in memory accesses and the computation.

1 Introduction

The main component of a conventional sequential computer is a monolithic memory that provides a uniform address space. The memory is typically implemented with a collection of DRAM arrays that are integrated in one or more chips and connected with an interconnect

*This work is licensed under the Creative Commons Attribution 4.0 International License.
<http://creativecommons.org/licenses/by/4.0/>

specialised to transmit control, data and address information, to provide efficient random access.¹ The architecture of a DRAM system is therefore tightly coupled but inherently distributed. In this sense, a DRAM system is conceptually similar to a distributed-memory parallel computer, which is composed of an array of processors with local memories, connected by a communication network. A parallel computer can thus be viewed as a more general system that provides processing capability at each memory as well supporting arbitrary patterns of communications between processors.

The similarities between conventional monolithic memory systems and the features of a distributed parallel architecture lead to the question, and the subject of this paper, *how efficiently can a parallel computer emulate a conventional memory system?* The ability to do so efficiently would allow a parallel computer to support sequential execution, in the same way that sequential computers are able to efficiently emulate parallelism, an ability of considerable importance. For parallel computers, as well as providing a means to execute legacy software directly, efficient emulation of large memories would allow the degree of parallelism to be matched with that of the problem being solved, by decreasing the number of available processors and increasing the amount of memory available to each one.

The critical factor in determining the ability of a distributed-memory parallel computer to emulate a large memory is latency. DRAM systems have been subject to around 40 years of optimisation in their architecture and manufacturing processes to minimise access latency. Latency in DRAM systems is therefore determined primarily by the intrinsic physical characteristics of the component devices; specifically by the transmission delay on wires, the dimensions of the component memory arrays and the critical delay of control components. The hypothesis of this paper however, is that because of the architectural and physical similarities, the overheads for a more-general parallel computer system are in fact relatively small and that it can therefore deliver an efficient emulation with only a small constant-factor overhead in performance.

1.1 Contributions and outline of the paper

The contributions of this paper are: (1) a proposal for a parallel architecture that can support sequential programming techniques by emulating large memories; (2) a model of the implementation of the architecture, based on current production manufacturing technologies, that demonstrates its practicality and scalability; (3) parameters for the implementation model based on the state of the art; and (4) a performance evaluation of the proposed architecture with results that demonstrate that the efficiency of a memory emulation incurs just a small constant-factor overhead.

The proposed parallel architecture is presented in §2; §3 provides some background to the relevant technologies; §4 describes the implementation model; §5 presents the parameters

¹ The central component of a DRAM is an *array core*, a two-dimensional array of cells and associated peripheral circuitry. The array core is sized to trade-off well between density and delay and energy per activation and refresh. Array cores are limited to a modest size that grows very slowly with respect to technology scaling due to their intrinsic capacitance. For more details on the architecture of DRAM chips see Itoh [17] and for general information on DRAM systems see Jacob, Spencer and Wang [19].

and an analysis of the model's cost and scaling; §6 describes the experimental methodology undertaken to characterise the emulation performance; §7 presents the performance results; §8 discusses related work; and §9 concludes.

2 A general parallel architecture

A general-purpose parallel computer architecture is presented in this section that is able to deliver an efficient emulation of a large memory. Only the salient aspects relevant to the comparison are discussed; for more details the reader is referred to [9, Ch. 5].

The architecture consists of a collection of tiles, each containing a processor, a memory and an interface to the network. The tiles are connected by an interconnect that supports efficient communication between all pairs of tiles. The specific details of the processor architecture are not important² except that the latency of moving data into the network is minimised. As well as mediating the transfer of data between the processor and network, the network interface also supports direct memory access. This ability allows a processor to read and write the memory of a different tile without involving the processor of that tile, and thus incurring additional latency in the transaction. Conceptually, each tile can be viewed as a unit of processing, or a unit of memory.

The structure of the interconnect is a folded Clos topology [21, Ch. 3], which provides a logarithmic diameter and scalable bisection width. The former characteristic is essential to reduce the distance a message travels between tiles and therefore its latency, and the latter for maintaining low latency when there are many simultaneous communications due to a parallel workload. Folded Clos networks are part of a family of closely-related hypercubic networks [26, Ch. 3] but they have a particularly flexible structure that is well suited to practical implementations.³ Specifically, folded Clos networks can be constructed from switches of any fixed degree and their hierarchical structure allows a network to easily be expanded or for the bisection width to be adjusted to meet the constraints of an implementation technology.

It is beneficial to use high-degree switches to implement a folded Clos network because, by connecting multiple terminals to each one, the network diameter can be reduced. A degree-32 switch provides a good balance between connectivity and implementation complexity. With two stages it can connect up to 256 tiles, which fits well into the target economical chip area (see §5.0.1). A degree-64 switch can connect up to 1,024 tiles with two stages but this size of system far exceeds the target area. Furthermore, with a degree-32 switch it is practical to use

² Microarchitectural optimisations to improve sequential performance are employed at the expense of silicon area. The larger the processor, then fewer cores can be integrated into a single chip. A tradeoff must be struck based on the requirements of the processor.

³ A Clos network is defined with three stages: input, middle and output, each consisting of a number of nodes. Each of the middle-stage nodes is connected to each of the input and output nodes. The symmetry of a Clos network allows it to be *folded* around the middle stage and for the input and output networks to be merged using bidirectional links [21, Ch. 3]. This construction is similar to a fat tree [27]. A folded Clos has two practical advantages: it makes the network more convenient to package when connecting terminals that produce both input and output, and messages can traverse shorter paths between terminals.

half the links to connect tiles, leaving sufficient bandwidth in the set of remaining links. Figure 1 shows example folded Clos networks connecting 64, 256 and 1,024 tiles.

It is interesting to note that, in contrast with parallel programs, execution of sequential programs will not induce any concurrent communication traffic in the network, and unless additional processes are run in parallel, each message will travel without contention. When this is the case, the interconnect needs only to provide low latency and high bandwidth, and therefore, a simple routing scheme can be used where messages travel along shortest paths.

Following the production and packaging model of commodity DRAM,⁴ the proposed parallel architecture can be implemented with tens or hundreds of tiles on a single chip that is sized to provide the best economic trade-off between performance and system cost. Each processing chip contains a complete sub-folded Clos network with replicated processing tiles, switches and communication links. The hierarchical structure of the folded Clos network allows multiple chips to be combined to directly extend the network. Each chip contributes additional switches to form a core switching stage and communication links connect switches on different chips, in the same way they are connected on-chip. Thus, a particular number of chips are chosen to build a system to provide a certain memory capacity or processing capability.

2.1 Emulation scheme

Typically, the requirements of local storage, which includes the program, constant values and the stack, are small and can fit into the local memory of a tile. The remaining global storage, which is used for statically-allocated data and the heap for dynamically-allocated data, can be stored in an emulated memory. The emulation is performed with a collection of tiles that are managed by a controller process. The controller receives access requests over a contiguous address range from a sequential client program and distributes them over the tiles by sending direct memory access read and write messages to them.

In the sequential (client) program, accesses to the emulated memory are written as communication sequences, where an instruction to load a value at an address (`addr`) in memory into a destination operand (`dest`) has the abstract correspondence:

$$\text{LOAD } \text{dest}, \text{ addr} \quad \rightarrow \quad \begin{array}{l} \text{SEND } c, \text{ READ} \\ \text{SEND } c, \text{ addr} \\ \text{RECV } c, \text{ dest} \end{array}$$

⁴ A low cost-performance ratio is the primary driver in the production of commodity DRAMs. The basic memory cells and supporting circuits are replicated millions of times on a chip and consequently the designs are highly optimised. DRAM chips are produced in high volumes and are produced at a size which makes the best trade-off between cost and performance. Consequently, their size has remained virtually constant between technology generations. Current DRAM chips are manufactured with an area of around 100 mm² with a capacity of approximately 0.3 Gbits [15, pp. 84-104].




Figure 1: Example 32×32 crossbar switch topologies for various size folded Clos networks. In the 64- and 256-tile networks, edge switches (labelled s_i) connect 16 tiles and core switches (labelled c_i) use all 32 links to connect to the previous stage. The 1,024-tile network is constructed by replicating a 256-tile network, with twice the number of core switches than (b), four times and connecting it with 32 core switches, making it three-stage. This construction maintains capacity between stages and has a logarithmic diameter (2 or 3 for the examples) in the number of tiles.

In this, a READ request and an address are sent to the controller, and then when the value has been retrieved, it is received and assigned to the destination operand. Similarly, a write has the correspondence:

		SEND c , WRITE
STORE value, addr	→	SEND c , addr
		SEND c , value

In both, c is a reference to a channel of communication, by which data can be conveyed to the memory controller process.

If the requirements of local storage exceeded the capacity of local memory, then there are two potential resolutions. Either, the program or stack could be split between processors, and control would transfer from one processor to the other when using either portions of the stack. Alternatively, a second emulated memory could be used. For the purposes of this paper and without loss of generality, only the simple case is considered where local data does not exceed the local memory capacity.

3 Background

This section discusses background to the implementation model: the silicon interposer technology that is used to connect multiple chips, the memory technology, the characteristics of on-chip wires, and the processor and switch components.

3.1 Interposer

The most challenging aspect of the implementation of the parallel architecture is the density of connectivity between chips that is required to extend the folded Clos network and maintain a good bisection bandwidth. Conventional packaging technology poses significant challenges for interconnecting chips that contain large numbers of processors. A high-density ball grid array (BGA) package can have up to approximately 2,800 pins. The Xilinx XCVU440 package, for example, has 2,892 pins with 1,327 I/Os [47]. It is typical that approximately 40% of package pins are used for ground and power [15, Tab. ORTC-4]. However, for bidirectional byte-wide links, only 82 such links could be connected with the XCVU440 package; significantly less than the level of on-chip connectivity.

Printed circuit boards (PCBs) offer a minimum wiring pitch of about 0.15 mm and a high-density BGA package would require approximately six wiring layers to route tracks to each of the pins [36]. The resulting density of wiring on the board is also low compared to on chip; 16 adjacent links, for instance, would occupy 38.4 mm. Even for modest numbers of processors, PCB technology does not provide sufficient connectivity between chips containing hundreds of processors to obtain a scalable bisection width. Thus, with PCB connections, the communication properties on chip do not extend off chip, complicating the programming model.

The only current production technology that provides high levels of connectivity between chips are silicon interposers. (However, in the future, other forms of 3D integration would apply naturally to package these processing chips in dense stacks.) Silicon interposers are becoming an established technology that can be used as a substrate to interconnect multiple processing chips with high-density wiring. In essence, they can be thought of as a high-density PCB. The interposer can be manufactured using a larger process geometry than the chips it connects to provide a good tradeoff between yield and wiring density⁵.

As an example, Xilinx use a 65 nm 776 mm² silicon interposer in their Virtex-7 FPGA [39] to connect four 28 nm FPGA ‘slices’ as a single larger device. This was the first production device to do so, with the principal reason to for using a silicon interposer being an economic one. Compared to contemporary devices, 776 mm² is a very large chip size, but at the mature 65 nm technology, production costs are relatively low and yields are high. Producing a number of

⁵ The *process technology* or *process geometry* of a chip refers to the minimum half pitch of contacted metal lines, as defined by the ITRS [15, pp. 5-6], which is the minimum feature size of a circuit. This has historically been an indicator of integrated-circuit scaling and it has been DRAM that has exhibited the smallest metal pitch. Larger process geometries produce higher yield because the technology is better developed and so it is practical to produce larger die sizes.

identical smaller devices and connecting them with an interposer to produce a single larger device provides a good overall trade-off between cost and performance. The same approach is taken with the implementation of the proposed architecture, using the characteristics of the Xilinx interposer are used as a basis.

3.2 Memory technology

The choice of tile memory technology and capacity depends on obtaining a good trade-off between access latency, area and its integration with microprocessor logic. Ideally, the access time of a memory should be matched with the time to execute one basic operation, but access time is generally inversely related to density, therefore faster memories require more silicon area⁶.

There are two dominant forms of fast random access memory that can be integrated with CMOS logic: static RAM (SRAM) and embedded dynamic RAM (eDRAM). SRAM can be integrated directly, has fast access latency and good random access performance but it has a relatively low density since it uses six transistors per bit. eDRAM is integrated with CMOS logic by adding three to six additional steps to the fabrication process and is more dense than SRAM (by a factor of 2 to 3) since it uses a single transistor-capacitor pair to store a bit [16].⁷

In contrast, contemporary commodity DRAM is produced on a manufacturing process that is optimised for density and power. Volume devices also lag logic technology by several nodes due to a better cost-performance tradeoff. DRAM is typically packaged on PCB modules, called dual-inline memory modules (DIMMs), and connected with wire traces on a parent PCB.

3.3 Wires

The logic and wiring resources provided by very-large-scale integration (VLSI) devices are plentiful and they continue to scale, but chip utilisation in terms of computational performance is limited by two factors: the I/O bandwidth of packages (i.e. getting data in and out) and the performance of long wires relative to the size of a chip in terms of delay and power [11] (i.e. moving data around a chip).

There are techniques for implementing power-efficient global interconnections such as low-swing signalling schemes [12] and optimisation of wire width and spacing [28], but an unavoidable effect of the scaling of VLSI devices is that wire delays exceed a single clock cycle. This makes modular parallel architectures that can tolerate multi-cycle latencies an attractive way to build systems [10], and adds weight to the proposed approach.

⁶ Faster memories are by design less dense, with control logic overheads amortised over a smaller number of bit cells.

⁷ As eDRAM technology has matured, the access latency has approached that of SRAM [1], especially for large memories where the flight-time across the array can dominate access latency and the improved density of eDRAM can significantly reduce this. Its main application has been to build large on-chip caches, replacing SRAM, for example in the BlueGene/L [18] and Power7 [5] chips.

Wire delay is related to the square of the length of the distance a signal is transmitted. Long wires are therefore unattractive. A conventional technique to reduce wire delay is to insert repeating elements at regular intervals along a wire [10, Chap. 4]. This reduces the total delay to a multiple of the delay between repeaters, at optimal intervals, this produces a linear relationship between delay and length. The cost of this is the delay through the repeaters and the logic area and power required for them.

3.4 Processor and switch components

The performance model of the parallel architecture and memory emulation uses the XMOS XCore processor architecture [31], for which detailed information is available, and because it is well suited to be integrated into the proposed architecture because it provides low-latency communications. The execution of a send operation, to transfer data between one processor and another, moves data in a single cycle from the processor's registers into a message buffer. In the next cycle, the message can be transmitted on to the network, towards its destination. Additionally, the provision of hardware support for a number of processes with dedicated registers and a simple scheduler avoids context-switch overheads.

The size of the switch component is based on the INMOS C104 switch [32, Ch. 3]. The C104 is of degree 32 and provides the required capabilities to construct a folded Clos network, particularly with its routing scheme and ability to aggregate links.

4 An implementation model

A VLSI floorplan of the architecture is described in this section, with a level of detail to produce approximate, but not unrealistic, estimates of area and wire delays to characterise the parallel machine performance. The floorplan is used to produce figures for the dimensions of chips and their components and wire lengths. These figures are then combined and compared to show the relative proportions of processors and memory to switches and interconnect wiring. Building on this model, the specific technology parameters are presented in §5 and an associated model of performance is presented in §6.3.

To provide a baseline for the evaluation, an equivalent construction with a 2D mesh network is presented. 2D meshes are a popular choice for an interconnect because they are simple to implement in two dimensions on a chip. 2D meshes are not however suitable for interconnecting large numbers of processors because their diameter and bisection bandwidth do not scale well with increasing numbers of attached processors.

4.1 VLSI model

The following simplified VLSI model is used to produce a floorplan for the chip. It is based on those in Mead and Conway [33] and Ullman [41].

4.1.1 Metal layers

- A chip consists of a number of metal layers. Metal layer 1 (M1) is used for logic components and the remaining metal layers M2, M3, M4 etc. are used to route wires.
- A particular metal layer contains wires orientated in one direction and adjacent metal layers contain wires in perpendicular directions. This is to reduce crosstalk, where a signal in one circuit affects the signal in another circuit.

4.1.2 Wires

- Wires carry signals in one direction and all wires are half shielded to further reduce crosstalk. This is where a ground wire is placed either side of a wire pair. Minimum signal wire density consequently decreases by $\frac{1}{3}$.
- All wires have repeaters inserted to minimise latency.
- Wires with multicycle delays are pipelined by inserting flip-flops.
- For sets of wires, banks of repeaters and flip-flops are required. To simplify their insertion, interconnect wires are routed in dedicated channels.

4.1.3 Chip

- The area of a chip is taken to be the smallest rectangle in which the complete circuit fits.
- External connections are made to I/O pads with a fixed driver circuit component. The driver deals with the higher capacitance of external wires.

4.1.4 Further assumptions and limitations

Only the following essential aspects of the system are captured in the layout to avoid over complicating the model:

- the only components included in the layout are the processor core, memory, switch, pad driver circuitry and I/O pads (internal component wiring is not considered) and these are assumed to have a square footprint;
- only the routing of the inter-switch and switch to chip I/O wiring is considered since this is the most significant component of the system in terms of wire lengths and density;
- the remaining communication links between switches and processors are not explicitly accounted for in the area; it is assumed they can be routed over other resources, in the area not occupied by the interconnect wiring channels;

- links are connected between components and it is assumed they can be connected from any position within the footprint of a component;
- layout relating to wiring and I/O for power and clock signals is not considered, but provision is made for them in the available metal layers and chip I/Os.

4.2 Folded-Clos layout

A single chip contains a collection of tiles interconnected with a complete sub-folded Clos network. The chip presents a number of external links from its stage-two switches so that multiple chips can be connected with the links between switches in different chips to directly extend the network. Banks of additional switches are included to contribute to the third core stage of a larger multi-chip network (note that the number of switches per chip is constant with respect to the overall network size). The layout of the folded Clos network has a recursive structure and is based on an H-tree, an efficient embedding of a binary tree onto a 2D surface [33, Ch. 8]. An example layout for a two-stage 256-tile folded-Clos network relating to the topology of one of the four sub-folded Clos networks in Figure 1c is given in Figure 2a.

The core switches of the chip's folded Clos network are placed centrally and half of the links are routed to the off-chip I/O pads. The next stage of switches is divided into four groups and each are placed in the centre of a quadrant surrounding the core switches. Connections are made from all of the core switches to each of the next stage switches. This layout continues recursively in each quadrant for each additional stage. The minimum separation of groups of tiles is constrained by the placement of switches and the width of the wiring channels between them. All links from the contributed stage-three core switches are routed to off-chip I/O pads, even though they will connect to the chip's second-stage switches, because the connection pattern depends on the size of the multi-chip network.

Groups of switches are arranged in staggered sets to minimise the resulting size of their bounding box. Their placement relative to one another is constrained in two dimensions by the pitch of wires connecting horizontally and vertically. These are routed on the minimum of two metal layers, but additional layers can be used to reduce the dimensions of a group. Each group of switches has branching connections in both directions. This connectivity restricts the horizontal and vertical placement of switches. A switch arrangement is chosen to minimise the width of the group, subject to not exceeding the height of its quadrant, since it can extend into the wiring channel. An example of the wiring pattern for a switch group is illustrated in Figure 2a. The total area of the interconnect is calculated as the sum of the area of the wiring channels and all of the switch groups.

The pads and driver circuitry for off-chip I/O are positioned along one edge of the chip due to the wiring pattern on the interposer. A chip that contains N tiles requires I/O for $2N$ links to extend the network, N from the core switches and N from the contributed bank of system core switches.

4.3 2D-mesh layout

A 2D mesh network is laid out as an array of blocks of tiles where each block connects to a single switch. The blocks are separated by wiring channels accommodating the width and height of a switch, which is placed at the corner of its block. Adjacent horizontal and vertical links connect directly between switches. The layout for a 256-tile 2D mesh network is given in Figure 2b.

The pads and driver circuitry for the off-chip I/O are positioned at the edges of the chip so that the mesh can be extended directly between adjacent chips. A chip that contains N tiles requires I/O for $4\sqrt{N} - 4$ links to extend the network.


4.4 Silicon interposer

The processing chips are mounted on a silicon interposer using a flip-chip assembly. This process bonds the top side of the chip with an array of solder microbumps. The interposer provides wiring traces between the microbumps to connect links between switches on different chips. The interposer is connected to a package substrate with through-silicon vias (TSVs) that provide a connection through the interposer's substrate to controlled-collapse chip-connection (C4) flip-chip bumps on the package substrate. These connections only bridge the power, ground and clocking I/Os. The package substrate carries these connections out of the package, to the BGA balls.

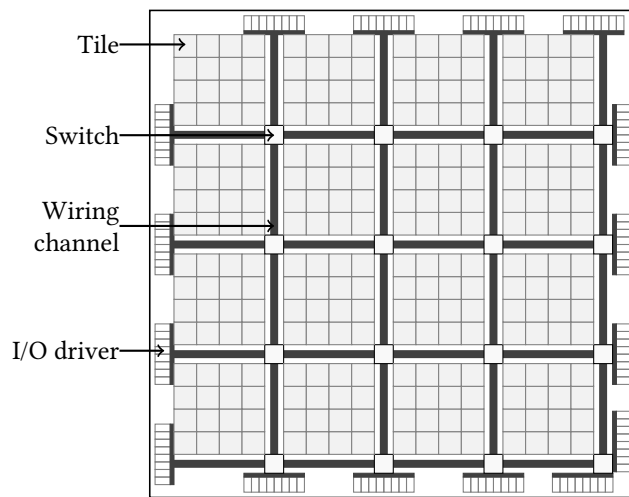
To extend the folded Clos topology, a set of chips are arranged in two rows, either side of a wiring channel on the interposer, orientated with their wiring area closest to this. Each wire in the channel is used to connect between pads on two chips with additional horizontal wires. The wiring channel contains sufficient wires so that every pair of connections between the chips can be made. The maximum height of the channel between two chips is therefore twice the total pitch of the wires connecting a chip. To extend the 2D mesh topology, connections are provided on the interposer between the links of adjacent chips, directly extending the mesh. In both networks, wiring for power, ground and clock signals is not considered since provision is made in the available metal layers, but the area requirement of the additional I/O pads required is accounted for directly. Figure 4 illustrates the packaging of a set of processing chips stacked on the silicon interposer for both networks.

5 Implementation technology parameters

This section presents and explains the parameters used to characterise current production technology for the implementation model. The International Technology Roadmap for Semiconductors (ITRS), which is a set of documents produced annually by a broad collection of industry members, provides many of these; others are taken from elsewhere in the literature or estimated. Although each parameter is given a specific value, the abstract nature of the implementation model means it is relatively robust to variations in them.



(a) Clos layout



(b) mesh layout

Figure 2: Block diagrams of the 256-tile folded Clos and 2D mesh network layouts for the processing chip. In (a), the topology corresponds to one of the four sub-folded Clos networks in Figure 1c. The layout is based on an H-tree organisation with groups of switches at each node organised in staggered sets to minimise area subject to the constraints of their vertical and horizontal connections. The driver circuitry and pad for each chip I/O is situated along the right-hand side of the chip. In (b), running around the edge of the chip is the driver circuitry for each of the chip I/O pads. Note that the connections from switches to processors and I/Os for power and ground are not included in either layout.




Figure 3: Cut-through view of a multi-chip package with a silicon interposer. A collection of chips are stacked using a flip-chip assembly onto solder microbumps, connecting them to the silicon interposer. The interposer provides connectivity between the chips, essentially as a high-density PCB. TSVs provide connectivity through the interposer's own substrate and these are connected with C4 solder bumps to the package substrate.




Figure 4: Chip layouts on the silicon interposer. For a folded Clos network, the chips are positioned beside a wiring channel that contains a common wire for every connection between two chips. A connection is made by connecting horizontal wires from each chip to the common wire; diagram (a) shows a wiring pattern for connections from each chip to every other chip. For the 2D mesh layout, in (b), chips are arranged in a grid and the network is extended directly between chips.

Parameter	Value	Note
Process geometry	28 nm	
FO4 delay	11 ps	See §5.0.1.
Economical chip sizes	80-140 mm ²	Cost-performance processor [15, Tab. ORTC-2C].
Metal layers	8	M1 logic; M2, 7 & 8 power & clock; M3-M6 wiring
Interconnect wire pitch	125 nm	For the global interconnect [14, Tab. INTC6 ('12)].
Repeated wire delay	155 ps/mm	See §5.0.1.
Processor area	0.10 mm ²	See §5.0.2.
Switch area	0.05 mm ²	See §5.0.2.
I/O pad dimensions	45 × 225 μm	See §5.0.1.
Wires per link	18	1 control, 8 data per direction, see §5.0.1.
Power and ground I/Os	40%	See [15, Tab. ORTC-4].
Clock rate	1 GHz	See §5.0.1.

Table 1: Implementation parameters for the processing chip model.

Parameter	Value	Note
Process geometry	65 nm	†
FO4 delay	24 ps	See §5.0.1.
Metal layers	4	†, M1 & M2 power & ground; M3 & M4 wiring.
Interconnect wire pitch	2 μm	†, 333 half-shielded wires/mm.
Repeated wire delay	89 ps/mm	See §5.0.1.
Microbump pitch	45 μm	†, 493.83 bumps/mm ² .
TSV pitch	210 μm	†, 22 TSVs/mm ² .
C4 bump pitch	210 μm	†, 22 bumps/mm ² .
Wires per link	10	1 control, 4 data per direction, see §5.0.1.

Table 2: Implementation parameters for the interposer model. † parameters based on Xilinx Virtex 7 FPGA package [37], which integrates four 28 nm FPGA slices on a 775 mm² interposer.

The remainder of this section is divided into the parameters of the processing and interposer chips, the main VLSI components and the memories. Table 1 summarises the parameters for the processing chip and Table 2 summarises the parameters for the interposer.

5.0.1 Processing chip and interposer

The parameters for the processing chip are based on a 28 nm logic process. Although 20 nm and 16 nm are also commercially available, they are relatively new and accordingly less mature, making them substantially more expensive. 28 nm can therefore be considered representative of current devices. The range of economical chip sizes, 80-140mm² is based on the ITRS data for a ‘cost-performance microprocessor at production’ [15, Tab. ORTC-2C].

The chip I/O pad area includes the contact and driver circuitry and its dimensions are estimated based on a 1-to-4 ratio between width and height, which is characteristic of conventional designs [46]. The width is taken to be the pitch of the interposer contacts (microbumps). An on-chip processor and interconnect clock rate of 1 GHz is chosen as this is typical of a high-end

Process geometry (M1 $\frac{1}{2}$ pitch)	Minimum global wire pitch (nm)	RC delay (ps/mm)	ITRS edition
150	670	not given	2001
90	300	96	2005
68*	210	168	2007
45	154	385	2010
37.84	114	621	2011
26.76*	81	1,115	2012

Table 3: ITRS data for global wires [14]. The rows marked with a * are the ones used to determine the wire delay for the processing chip and interposer.

embedded processor. The parameters for the silicon interposer and packaging are based on the passive interposer of the Xilinx Virtex 7 FPGA [22, 39, 37], however it is assumed that wires can be repeated.

All links between switch components on the processing chip have 9 wires in each direction, allowing up to one byte to be transmitted per cycle with an additional bit to signal control tokens. All off-chip links have 5 wires in each direction, allowing up to one byte to be transmitted every two cycles. The reduction in off-chip link wires is to reduce the off-chip I/O requirements and interposer wiring.

The delay of an optimally-repeated wire can be estimated with the equation

$$\tau = 1.47\sqrt{\text{FO4} \cdot \hat{R}\hat{C}}$$

where τ is the delay in picoseconds per unit length; $\hat{R}\hat{C}$ is a time constant equal to the product of the resistance per unit length and the capacitance per unit length; and FO4 (fanout-of-4) delay is a constant related to the particular process technology, defined as the delay of an inverter, driving four identical copies of itself (see [10, §2.2] for more details). This approximation is based on the derivation by Bakoblu [3] and follows the explanation by Ho [10, §4.2]. The FO4 delay is estimated for a particular process geometry using the heuristic $\text{FO4} = 360 \cdot f$, where f is the feature size in μm , producing a delay in ps [11].

The ITRS provide figures for wire RC delay, which are summarised for a range of process geometries in Table 3; their method of calculation for the RC delay is described in [14, §7.3 ('12)]. The M1 half pitches of 68 nm and 26.76 nm are the closest matching geometries of the processing chip and interposer respectively and the data for them are taken to estimate wire delay. Using these figures and the above formula for wire delay yields 155 ps/mm for the 28 nm processing chip and 89 ps/mm for the 65 nm interposer chip.

5.0.2 Processor and switch components

The area A of a component at a process geometry g is estimated for a different geometry h by applying a linear scaling $A_h = A_g/(g/h)^2$, where $g \geq h$. The area of the XMOS XCore

Type	Typical capacity (MB)	Cell area factor (F^2)	Area efficiency	Current process geometry (nm)	Density (KB/mm ²)	Cycle time (ns)
SRAM	<8	140	70%	28	778.51	0.5
eDRAM	1 - 64	50	60%	28	1,868.42	1.3
Comm. DRAM	>64	6	60%	40 ^a	7,629.39	30 ^b

Table 4: Comparison of contemporary memory technologies, with figures from the 2012 ITRS report [16, Tab. SYSD3b]. The area factor is a multiple of square half pitch units, the number given by the process geometry. The area efficiency is the proportion of area in a memory array that is occupied by storage cells.

^aHigh-volume commodity DRAM, in general, lags commodity logic by several process technologies.

^bRandom cycle time (t_{RC}) from a 1 Gb Micron DDR3 device [34], contemporary to 28 nm logic.

processor on a 90 nm process is conservatively estimated to be 1 mm²; on a 28 nm process it is estimated to be 0.10 mm². The area of the C104 switch on a 1 μ m process was approximately 40 mm²; on a 28 nm process it is estimated to be 0.03 mm².

These figures are consistent with the ARM Cortex-M0 processor [2] and the 32 \times 32 SWIFT switch [40]. The Cortex-M0 is a simpler processor than the XCore, having only a 3-stage pipeline and support for a single hardware thread. (The XCore has a 4-stage pipeline and eight hardware threads.) On a 40 nm process the M0 has an area of 0.01 mm² and on 28 nm an estimated area of 0.003 mm². The SWIFT switch has an area of 0.35 mm² on a 65 nm process and an estimated area of 0.06 mm² on 28 nm.

5.0.3 Memory

The area and access latency of memory for the processing tiles is estimated using data from the ITRS [16, Tab. SYSD3b]. Table 4 gives the key characteristics of the dominant forms of memory technology. Figures for SRAM and eDRAM are given for 28 nm, in line with the processing chip; 40 nm DRAM is included as a baseline, which is characteristic of contemporary commodity devices.

Although density, and in particular, cycle time (the time to perform a random access) will vary with capacity and process technology, typically, memory capacity increases with process technology and these remain roughly constant between generations. In general, eDRAM is 2 to 3 times the density of SRAM and 4 to 5 times less dense than commodity DRAM. SRAM cycle time is around three times faster than eDRAM and commodity DRAM cycle time is much higher because of the specialised process on which it is produced.

Although eDRAM can, in principle, be integrated with microprocessor logic, there are relatively few devices that use it as a technology, and the additional process steps that required significantly increase manufacturing costs. For these reasons, only SRAM technology is considered in the

implementation models. The tile capacities of 64 KB, 128 KB, 256 KB and 512 KB are selected since they have a similar area to the processor (0.08 mm² to 0.66 mm²).

5.1 Cost and scaling

The figures presented in this section are produced from the calculation of specific instances of the chip layouts for the chosen technology parameters.

5.1.1 Of the processing chip

Figure 5 shows how the total chip area scales with the number of tiles integrated. Horizontal lines are drawn to mark the range of economical chip sizes (80 mm² to 140 mm²) and marked 'Min' and 'Max'. Of the chips that fall in the economical range, the folded-Clos chip requires 13% to 43% more area than the interconnect in the 2D mesh chip with the same number of tiles and memory capacity. For example, the largest folded-Clos chip with 256 tiles with 128 KB of memory occupies 132.9 mm² (of which 44.6 mm² is occupied by I/O) and the corresponding 2D mesh occupies 87.9 mm².

For the folded-Clos chips, apart from the configuration with 128 tiles and 512 KB of SRAM, the wires connecting the tiles to the stage-one switch are less than 5.5 mm long, have sub-nanosecond delays and thus are single cycle. All other delays on wires up to 11.2 mm are less than two nanoseconds and thus have a two-cycle latency. For the 2D-mesh chips, wire lengths between switches vary from 1.7 mm to 3.5 mm with sub-nanosecond single-cycle delays.

5.1.2 Of the components of the processing chip

Figure 6 shows the total area of the switch, wire and I/O components of the chip as a percentage of the total area for 256 KB tile memories. These plots show how the component areas scale with respect to the number of tiles they connect, highlighting the difference between the folded Clos and 2D mesh in the rate that resources are invested with increasing size.

The total switch area for the folded Clos chips is calculated as the total area occupied by the switch groups. Although the number of switches and links per-tile is proportional to the number of stages and the growth of additional components is logarithmic in the number of tiles, the area grows more quickly than this due to the increasing inefficiency of larger switch groups. The I/O remains relatively constant per tile but occupies a large region of the die. For the smallest 64 KB memories it occupies approximately 40% of the die. Overall, for the economical chip sizes, the interconnect occupies between 5% and 8% of the die area.

For the 2D mesh, switch area remains constant per tile and the wire grows slowly because of the decreasing ratio of switches on the edges of the mesh to those in the middle. Similarly, the proportion of I/O diminishes as the number of tiles increases since the number of external

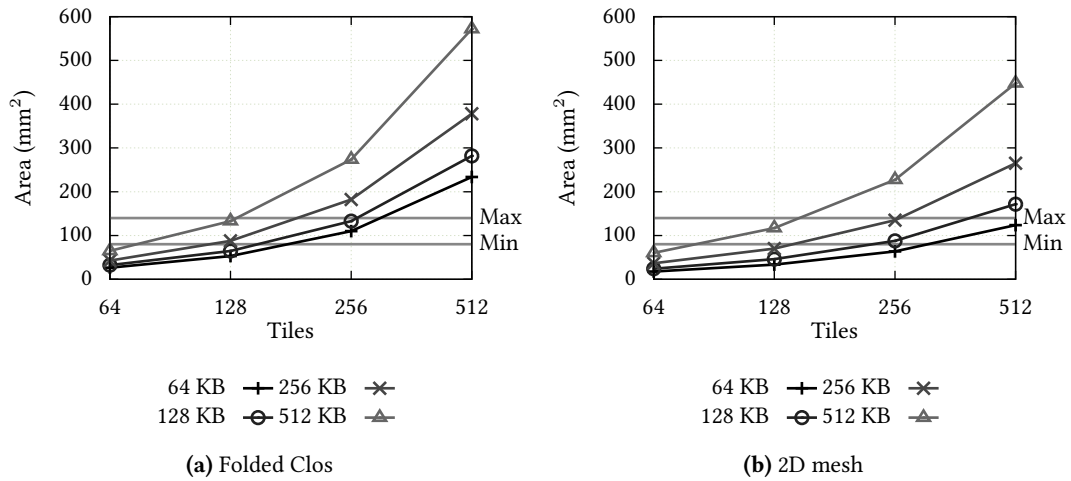


Figure 5: Log-linear plots of the total chip area as a function of the numbers of tiles. The grey horizontal lines indicate the range of economical chip sizes, from 80 mm² to 140 mm².

connections diminishes relative to the number of tiles. Overall, for the economical chip sizes, the 2D mesh interconnect occupies 2% to 3% of the die area.

5.1.3 Of the interposer

For a folded Clos system with multiple chips integrated on a silicon interposer, the percentage area occupied by the wiring channel varies between 2% for two 128-tile chips with 64 KB memory per tile (234 mm² total), and up to 42% for 16 512-tile chips with 128 KB of memory per tile (1,979 mm² total). Accordingly, the minimum and maximum wire delays range from 1 ns to 8 ns, which correspond to the width and the width plus the height of the channel. For the 2D mesh, the wire delay is a constant 0.09 ns for all systems since the chips are tiled and the distance between adjacent pads is constant.

The plots in Figure 7 summarise the total interposer area for a range of different system configurations, with varying numbers of different processing chips.




Figure 6: Log-linear plots of the total area of the switches and wiring components as a percentage of the total die area for chips with 256 KB of tile memory. Switch area is calculated as the sum of the switch group area in the folded Clos. This adds an overhead to the otherwise logarithmic scaling of switch and wire area. Both components remain constant in the 2D mesh, apart from a small convergent growth in the wire area due to a decreasing ratio of switches on the edges of the 2D mesh to those in the middle.




Figure 7: Log-linear plots of the total area of the interposer for different configurations of selected economically-sized processing chips.

6 Experimental methodology

This section presents the methodology by which the performance of the emulation of a large memory with the proposed parallel architecture is compared to that of a sequential machine.

6.1 Sequential machine model

The performance of a sequential machine is used to provide a baseline for the memory emulation. Performance is presented principally by the relative slowdown of the emulation compared to the sequential machine executing the same program. The modelled performance is based on latency since it is the most difficult aspect to scale; schemes for scaling bandwidth generally involve replication and hence may in principle be applied to both systems.

The sequential machine is modelled as a single processor of the same class as the parallel machine connected to a DRAM system. A cache is not modelled directly but memory accesses to areas stored in local memory in the parallel emulation incur only a single-cycle latency. The effect of this is similar to providing the sequential system with a fast cache memory with an 80% to 90% hit rate since global memory accesses constitute between 10% to 20% of executed instructions in the benchmarks. The comparison therefore directly compares DRAM accesses with emulated accesses. Indeed, any caching scheme for the sequential machine could equally be implemented as part of the the parallel emulation. The clock frequency of both systems is held at 1 GHz and it is assumed the sequential DRAM can operate at this speed (typical DDR3, which is contemporary to 28 nm logic, operates well in excess of this).

The access latency is estimated for a modern DRAM by simulation with DRAMSim2 [38]. Performance is measured by performing read and write accesses to addresses chosen uniformly at random over the address range and the fixed latency is calculated as the average of these accesses.⁸ Latency measurements are based on random accesses and accesses are issued only once the last has completed to restrict the memory controller to processing a single transaction at a time. Although random access is pessimistic because successive accesses to the same DRAM page exhibit less latency and schemes such as prefetching do also reduce latency, these are not considered in the model because they do in principle apply to both systems. For a system with 1 Gbit DDR3 chips [34], average random-access latency is measured at 35 ns for a single rank⁹ with a 1 GB capacity. For multi-rank system with 2 GB to 16 GB capacities, this increases to 36 ns due to a small overhead in switching between ranks. This choice is unimportant since the

⁸ There is an overhead in accessing a particular row of a DRAM array and successive accesses to the same row exhibit lower latency. Consequently, measurement of DRAM latency is based on two main components. The column address strobe (t_{CL}) latency, which is the time between specifying a column address and receiving the data in response, given that the row being accessed is already open. The row cycle time (t_{RC}), which is the minimum time between the activation of one row and another; there is a minimum period a row must be active for to perform a refresh and a non-active row must be precharged before it can be read from.

⁹ A DRAM rank is a set of one or more DRAM chips that are accessed simultaneously and provide a particular data width. For standard DRAMs, a channel typically has a 64-bit data bus and with error-correcting code (ECC) memory, this is expanded to 72-bits.

empirical analysis is concerned with obtaining results to within small factors and to demonstrate scaling behaviour.

Typically, DRAMs are packaged in DIMM cards but production processes are moving towards stacked packages integrated with through-silicon vias and it is expected that stacks with multiple DRAM chips will be available soon. Since DRAMSim2 does not model a particular packaging or wire delay, the above sequential machine model's performance is inclusive of a stacked DRAM configuration.

6.2 Benchmarks

The performance analysis is based on two benchmarks: synthetic instruction sequences and a simple compiler.

The synthetic instruction sequences contain a particular ratio of global memory accesses to local memory and non-memory operations, to coarsely characterise sequential programs. Non-memory instructions are, for example, arithmetic and branching; local-memory instructions include all accesses to the program, stack and constant data; and global memory instructions access static data and the heap region. The ratios are chosen based on the instruction mix of the Dhrystone benchmark [44] (which is intended to characterise integer general-purpose sequential programs) and at points over the range of potential different ratios to demonstrate the effect they have on performance. Figure 8a shows the proportions of different types of instructions executed for the benchmark. In all experiments, the proportion of local memory accesses in the synthetic sequences is fixed at 20%.

The compiler benchmark provides an example of a realistic general-purpose sequential application and a means of compiling sequential programs to the parallel architecture. A modified version of the compiler emits message-passing sequences in the place of global memory accesses in order to generate sequential programs that interact with an emulated global memory (the effect on code size is discussed in §7.3). Figure 8b shows the proportions of different types of instructions executed for the compiler benchmark.

6.3 Performance model

The system model presented in §4 provides values for the latency of signal transmissions along links and their bandwidth. The performance of communication depends on some additional parameters related to the latency of the switching elements in the network. The following describe these parameters and formulate a performance model for the folded Clos and 2D mesh networks.

Given an interconnection network represented by a graph $G = (V, E)$ whose vertices V represent nodes containing a switch and zero or more tiles, and edges E that represent communication links, the latency of a message sent from processor $s \in V$ to processor $t \in V$ depends on:




Figure 8: Instruction mix of the Dhrystone and compiler benchmarks, showing the proportions of executed non-memory, local memory and global memory instructions.

- t_{tile} , the latency of the link between the tile and the switch;
- t_{switch} , the switch latency;
- t_{open} , the additional latency to open a route through the switch;
- c_{cont} , the switch contention factor;
- $d(s, t)$, the length of the path, p , in the network ($d(s, t) = |p(s, t)|$);
- $t_{\text{link}}(u, v)$, the latency of the link between nodes $u, v \in V$;
- t_{serial} , the serialisation latency, which is determined by the message length and channel bandwidth, such that if s and t are on the same chip, then $t_{\text{serial}} = t_{\text{serial-intra}}$ and if they are on different chips, then $t_{\text{serial}} = t_{\text{serial-inter}}$.

When the route between s and t is not already open, message latency is calculated as

$$t_{\text{closed}}(s, t) = 2t_{\text{tile}} + t_{\text{serial}} + (d(s, t) + 1)(t_{\text{open}} + t_{\text{switch}} \cdot c_{\text{cont}}) + \sum_{\ell \in p(s, d)} t_{\text{link}}(\ell)$$

and when the route is open it is calculated as

$$t_{\text{open}}(s, t) = 2t_{\text{tile}} + t_{\text{serial}} + (d(s, t) + 1) \cdot t_{\text{switch}} \cdot c_{\text{cont}} + \sum_{\ell \in p(s, d)} t_{\text{link}}(\ell)$$

Both models are comprised of the sum of four latency components: tile-to-switch, serialisation, switch traversal and link traversal. With shortest-path routing, $d(s, t)$ is the minimum distance between s and t .

Parameter	Symbol	Value (cycles)	XMP-64 (cycles)
Tile-to-switch latency	t_{tile}	see §5.1	1
Switch latency	t_{switch}	2	2
Latency to open a route	t_{open}	5	5
Ser. latency intra-chip	$t_{\text{serial-intra}}$	0	0
Ser. latency inter-chip	$t_{\text{serial-inter}}$	2	4
Link latency	t_{link}	see §5.1	2 on-chip, 3 off-chip

Table 5: Parameters for the network performance model. The switch latencies are based on measurements made with the XMP-64; the other measurements are included for comparison.

Values for the parameters of the above latency are summarised in Table 5. The estimates are used to calculate the link and serialisation latencies and the switch latencies are estimated by fitting measurements taken with an XMOS device to the above latency model.¹⁰ Measurements taken for the other parameters are included for comparison, noting that there is no serialisation latency on chip because the tiles are connected to the switch with 8-bit links; off-chip, the link bandwidth is one byte every four cycles.

7 Performance results

In this section, the results of experiments with the modelled systems are presented. Absolute latency is considered first since this determines the performance of the emulated memory and is the most difficult characteristic to scale, after this, the performance of the benchmarks are presented.

7.1 Absolute latency

Figure 9 shows how the average access latency of random reads and writes in the emulated memory scales as the number of tiles is increased in the emulation. The baseline latency measured from the simulated DDR3 memory is included for comparison.

Performance of the folded Clos network clearly reflects the logarithmic growth of the network diameter and the latency incurred by additional stage in systems larger than 256 tiles can be seen. Latency in the 2D mesh increases linearly with the size of the emulation, with a change of gradient as communications occur between chips. Overall, the folded Clos delivers access latency that is within a factor of approximately 2 to 5, relative to a sequential machine with

¹⁰ The device was the XMP-64 [48], an array of 16 quad-core XS1-G4 chips [49]. For a general discussion of the methodology for taking performance measurements and other specific results with the XMP-64, the reader is referred to [8].




Figure 9: Log-linear plots showing how the absolute memory latency scales for 1,024- and 4,096-tile systems as the number of tiles in the emulation increases.

a DDR3 memory. The performance of the two networks is similar on-chip but the 2D mesh incurs a 30% to 40% overhead relative to the Clos for larger multi-chip emulations.

It is important to note that the results for absolute latency are based on the sequential and parallel systems having the same clock speed. An increase in clock speed for the sequential system relative to the parallel one would only improve bandwidth because the inherent latency of a DRAM cannot be improved. However an increase in clock speed for the parallel system would improve latency because the network would operate faster.

7.2 Benchmark performance

Figure 10 shows the performance of the synthetic Dhrystone and compiler benchmarks under a range of different system parameters. The general behaviour reflects that of Figure 9. The absolute access latency of the emulated memory is high, a factor of 3 to 4 times that of a specialised sequential machine for the folded Clos systems, but the effect of this is diluted by fast local accesses and other non-memory operations. The folded Clos systems can deliver an emulation with a slowdown of between approximately 2 to 3 up to 4,096 tiles over the sequential machine. The performance of the 2D mesh systems is similar to the folded Clos systems up to approximately 128 tiles. Beyond this, the performance of the 2D meshes deteriorate relative to the folded Clos, in which latency only increases slowly. With both systems, the execution of Dhrystone is less efficient due to the higher proportion of global accesses. It is also interesting to note that up to 16 tiles, there is a speedup over the sequential machine because the tiles are attached to a single switch.

In general, as the ratio of global memory operations to local and non-memory operations decreases, the slowdown does also, converging to a worst case of 1.5 to 2.5 overhead. This is the ratio between the sequential machine with the DDR3 memory and the parallel emulation shown

in Figure 9. This trend can also be seen in Figure 11, which shows the emulation performance for the synthetic benchmark with different proportions of global memory accesses (between 0% and 50%). For general sequential programs where there is a mix of operations and 10% to 20% global accesses.

7.3 Program binary size

Since each memory reference is written as a communication sequence (listed in §2.1), the size of the program binary increases. For loads, there is an overhead of two instructions and for stores, there is an overhead of three. For the version of the compiler that uses the emulated memory, when it compiles itself, the size of its executable binary increases by 8%. This however is a small price compared to the amount of extra memory that can be provided.

8 Related work

The need for scalable general-purpose systems has led to a number of proposals for single-chip tiled parallel architectures. Examples include the MIT Raw [43] and descendant Tile [45] architectures, Smart Memories [30], the PicoChip architecture [35], the UC Davis AsAP architecture [50], the Intel Xeon Phi processor [13], the Kalray MPPA [24] and the Cavium ThunderX [7]. However, despite an established theory of general-purpose parallel computation which requires low-diameter high-capacity networks [27, 42, 26], these systems neglect this theory and generally employ a 2D mesh-based interconnect. The results in this paper indicate that at zero load, 2D meshes may be suitable for memory emulations with modest numbers of tiles. However, it will be difficult to maintain efficiency with parallel workloads because the effects of congestion will increase latency.

Clos [6] networks and the related fat-tree [27] network have been widely studied and used for interconnecting parallel systems. They are well established as interconnects for large systems comprising thousands of processors in datacentre- and supercomputer-class systems. On chip, various proposals for electrical, e.g. [4, 29, 25, 20], and optical, e.g. [23, 51], Clos/fat-tree networks have been made in order to reduce latency, improve capacity and simplify programming. These proposals however only consider relatively small systems up to 64 processors and low-degree 2×2 to 16×16 crossbar switches.

9 Conclusion

This paper has proposed a general-purpose parallel computer architecture and a practical implementation for it using current production technologies. It is able to efficiently emulate large memories with collections of small-memory processing tiles, thus supporting the execution of sequential programs with large memory requirements. The resulting machine is capable

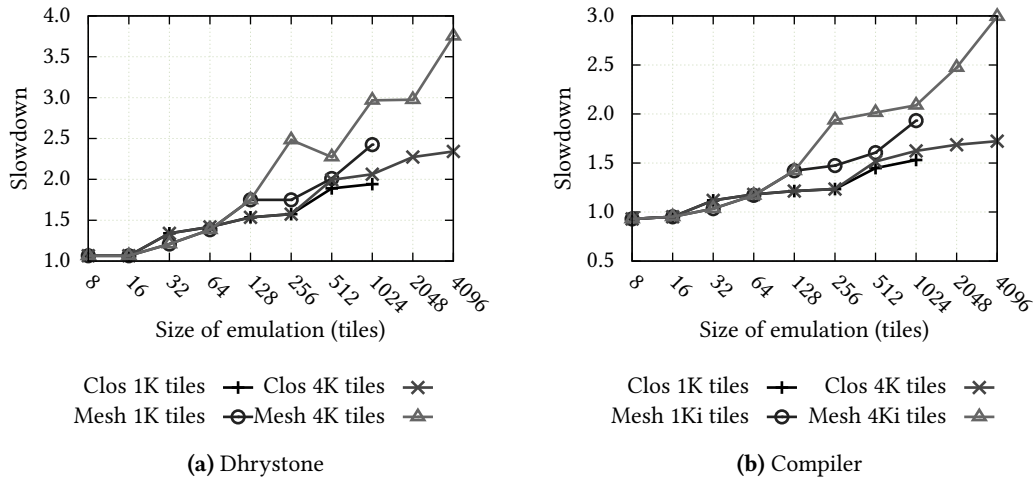


Figure 10: Log-linear plots of the performance of the synthetic Dhrystone and compiler benchmarks, relative to the sequential machine, using an emulated memory on 1,024- and 4,096-tile systems.

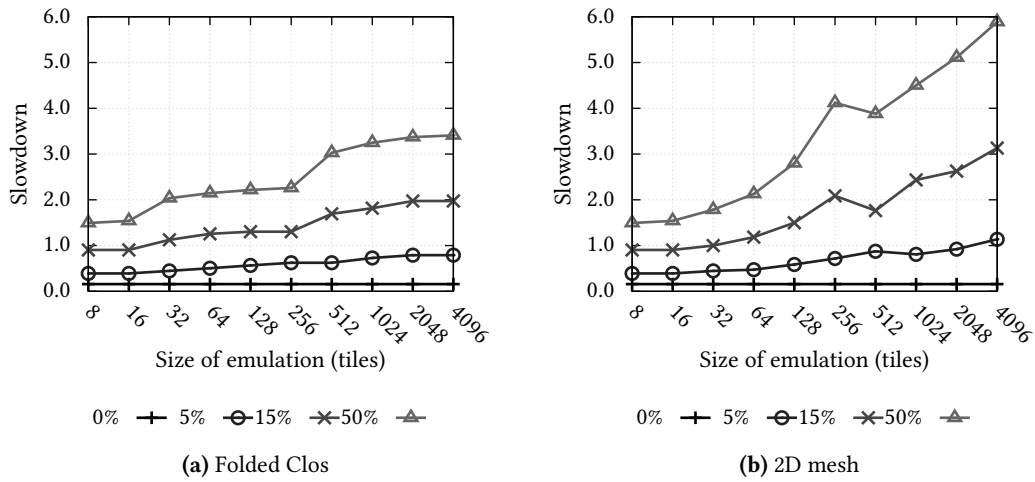


Figure 11: Log-linear plots showing the emulation slowdown, relative to the sequential machine, over a range of instruction mixes, with proportions of global accesses varying between 0% to 50%, for 1,024- and 4,096-tile systems. The proportion of local memory access is fixed at 20%, based on the Dhrystone and compiler instruction mixes.

of switching between executing highly-parallel and sequential programs, or allowing the programmer to compose parallel programs with large-memory sequential components.

The empirical performance results presented demonstrate that the emulation can be performed efficiently, with just a factor of 2 to 3 overhead for general sequential programs, compared to a conventional sequential machine. This small penalty however can be mitigated by exploiting parallelism in the program. The cost of the network, the central component of the system, is also relatively low. An on-chip folded-Clos network occupies approximately 7% of the die, and off chip, using silicon interposers to provide high density wiring, it occupies approximately 30% of the interposer die. Overall, this represents a 20% to 30% investment in the interconnect, compared to around 5% with a 2D mesh network.

The presented performance and implementation-cost results from the high-level modelling demonstrate what can be expected from the systems examined. This level of investigation is sufficient to reveal the scaling behaviour and the approximate constant factors involved.

References

- [1] Darren Anand et al. Embedded DRAM in 45-nm technology and beyond. *IEEE Design & Test of Computers*, 28:14–21, January 2011.
- [2] ARM Ltd. Cortex-M0 processor, 2012. Technical specification.
- [3] H. Buman Bakoglu. *Circuits, Interconnections, and Packaging for VLSI*. Addison-Wesley, 1990.
- [4] J. Balfour and W. J. Dally. Design tradeoffs for tiled CMP on-chip networks. In *Proceedings of the ACM/IEEE conference on Supercomputing*, ICS '06, pages 187–198. ACM, 2006.
- [5] J. Barth et al. A 45nm SOI embedded DRAM macro for POWER7 32 MB on-chip L3 cache. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pages 342–343. IEEE, 2010.
- [6] C. Clos. A study of non-blocking switching networks. *Bell System Technical Journal*, 32(2):406–424, March 1953.
- [7] Linley Gwennap. ThunderX rattles server market, June 2014. Linley Group Microprocessor Report.
- [8] James Hanlon. XMP-64 performance measurements. XMOS Ltd., Bristol, UK, 2009.
- [9] James Hanlon. *Scalable abstractions for general-purpose parallel computation*. PhD thesis, Department of Computer Science, University of Bristol, March 2014.
- [10] Ron Ho. *On-chip wires: scaling and efficiency*. PhD thesis, Department of Electrical Engineering, Stanford, 2003.

- [11] Ron Ho, Ken Mai, and Mark Horowitz. The future of wires. *Proceedings of the IEEE*, 89(4):490–504, 2001.
- [12] Ron Ho, Ken Mai, and Mark Horowitz. Efficient on-chip global interconnects. In *Symposium on VLSI Circuits, Digest of Technical Papers*, pages 271–274. IEEE, 2003.
- [13] Intel Corporation. Intel Xeon Phi coprocessor, June 2013. Datasheet 328209-002EN.
- [14] International Technology Roadmap for Semiconductors. Interconnect, 2001, 2005, 2007, 2010, 2011, 2012.
- [15] International Technology Roadmap for Semiconductors. Executive summary, 2012.
- [16] International Technology Roadmap for Semiconductors. System drivers, 2012.
- [17] Kiyoo Itoh. *VLSI Memory Chip Design*. Springer, 2001.
- [18] S. S. et al Iyer. Embedded DRAM: technology platform for the Blue Gene/L chip. *IBM Journal of Research and Development*, 49:333–350, March 2005.
- [19] Bruce Jacob, Spencer Ng, and David Wang. *Memory Systems: Cache, DRAM, Disk*. Morgan Kaufmann, 2007.
- [20] Yikun Jiang and Mei Yang. On circuit design of on-chip non-blocking interconnection networks. In *27th IEEE International System-on-Chip Conference (SOCC)*, pages 192–197, September 2014.
- [21] A. M. Jones, N. J. Davies, M. A. Firth, and Wright C. J. *The Network Designer’s Handbook*. IOS Press, 1st edition, 1997.
- [22] Handel H. Jones. Technical viability of stacked silicon interconnect technology. Technical report, IBS Inc., October 2010.
- [23] Ajay Joshi et al. Silicon-photonic Clos networks for global on-chip communication. In *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip, NOCS ’09*, pages 124–133. IEEE, 2009.
- [24] Kalray. Kalray MPPA Manycore processors, September 2012. Product brief.
- [25] Yu-Hsiang Kao et al. Cnoc: High-radix clos network-on-chip. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(12):1897–1910, December 2011.
- [26] Frank T. Leighton. *Introduction to parallel algorithms and architectures: array, trees, hypercubes*. Morgan Kaufmann, 1992.
- [27] Charles E. Leiserson. Fat trees: universal networks for hardware-efficient supercomputing. *IEEE Transactions on Computers*, 34(10):892–901, October 1985.
- [28] X. C. Li, J. F. Mao, H. F. Huang, and Y. Liu. Global interconnect width and spacing optimization for latency, bandwidth and power dissipation. *IEEE Transactions on Electron Devices*, 52(10):2272–2279, 2005.

- [29] D. Ludovici et al. Assessing fat-tree topologies for regular network-on-chip design under nanoscale technology constraints. In *Proceedings of the Conference on Design, Automation and Test in Europe, DATE '09*, pages 562–565. European Design and Automation Association, 2009.
- [30] K. Mai et al. Smart Memories: a modular reconfigurable architecture. *SIGARCH Computer Architecture News*, 28(2):161–171, 2000.
- [31] D. May. *The XMOS XS1 Architecture*. XMOS Ltd., October 2009.
- [32] D. May, P. W. Thompson, and P. H. Welch, editors. *Networks, Routers and Transputers: Function, Performance and Applications*. IOS Press, 1st edition, 1993.
- [33] Carver Mead and Lynn Conway. *Introduction to VLSI systems*. Addison-Wesley series in computer science. Addison-Wesley, 1980.
- [34] Micron. 1Gb: x4, x8, x16 DDR3 SDRAM: MT41J128M8JP-125 features, 2012. Datasheet.
- [35] Gajinder Panesar et al. Deterministic parallel processing. *International Journal of Parallel Programming*, 34(4):323–341, August 2006.
- [36] Charles Pfeil. BGA breakout challenges. *OnBoard Technology*, pages 10–13, October 2007.
- [37] Suresh Ramalingam. Stacked silicon interconnect technology (SSIT) qualification – requirements and tools, July 2011. Presentation given at the 3D Stress Workshop.
- [38] P. Rosenfeld, E. Cooper-Balis, and B. Jacob. DRAMSim2: A Cycle Accurate Memory System Simulator. *Computer Architecture Letters*, 10(1):16–19, January 2011.
- [39] Kirk Saban. Xilinx stacked silicon interconnect technology delivers breakthrough FPGA capacity, bandwidth, and power efficiency. Technical report, Xilinx Inc., October 2011.
- [40] S. Satpathy et al. SWIFT: A 2.1tb/s 32x32 self-arbitrating manycore interconnect fabric. In *Symposia on VLSI Technology and Circuits, Koyoto, Japan*, pages 138–139, June 2011.
- [41] Jeffery D. Ullman. *Computational Aspects of VLSI*. W. H. Freeman & Co., 1984.
- [42] L. G. Valiant. General purpose parallel architectures. In *Handbook of theoretical computer science (vol. A): algorithms and complexity*, pages 943–973. MIT Press, 1990.
- [43] E. Waingold et al. Baring it all to software: RAW machines. *IEEE Computer*, 30(9):86–93, September 1997.
- [44] Reinhold P Weicker. Dhrystone: a synthetic systems programming benchmark. *Communications of the ACM*, 27(10):1013–1030, 1984.
- [45] D. Wentzlaff et al. On-chip interconnection architecture of the Tile processor. *IEEE Micro*, 27(5):15–31, September 2007.
- [46] Neil H. E. Weste and David Money. *CMOS VLSI Design*. Pearson/Addison-Wesley, fourth edition, 2005.

- [47] Xilinx. Ultrascale architecture: Packaging and pinouts, March 2015. Product Specification User Guide, UG575 (v1.3).
- [48] XMOS Ltd., Bristol, UK. *XK-XMP-64 Hardware Manual*, February 2010.
- [49] XMOS Ltd., Bristol, UK. *XS1-G04B-FB512 Datasheet*, October 2012.
- [50] Z. Yu et al. AsAP: An asynchronous array of simple processors. *IEEE Journal of Solid-State Circuits*, 43(3):695–705, March 2008.
- [51] Jing Zhang, Huaxi Gu, and Yintang Yang. A high performance optical network on chip based on Clos topology. In *2nd International Conference on Future Computer and Communication (ICFCC)*, volume 2, pages V2–63–V2–68, May 2010.