

**FEASIBILITY STUDY ON IMPLEMENTING  
THE “BALL COMPUTER”**

Richard Hind BSc, MCSE, MBCS

**Submitted for the Degree of:  
MASTER OF SCIENCE (BY RESEARCH)**

University of York  
Department of Computer Science

April 2013



FEASIBILITY STUDY ON IMPLEMENTING THE  
“BALL COMPUTER”

Richard Hind BSc, MCSE, MBCS

Submitted for the Degree of: Master of Science (by research)

University of York

Department of Computer Science

April 2013

## **Abstract**

As processor cores become ever smaller and more power efficient, coupled with the growth of multicore packages, the ratio of data interconnect power to core power is set to invert. While current technologies for on-chip wireless data communication offer an alternative to hard wired buses, the designs of multicore processors and large array machines still rely on conventional power rails to supply the cores.

This thesis will discuss a range of technologies with the potential to enable the wire-free transmission and efficient storage of power within large processor arrays and consider how this could be extended to a 3-dimensional array of self contained computing nodes. This discussion will also include a review of commercially available CPU cores with sufficiently low power consumption and deal with the issue of cooling, taking into consideration the overall system efficiency. There will also be a brief discussion of wireless communication technologies which could be employed to allow high speed data communication between nodes and a review of a small range of power storage technologies which potentially offer high power density and a rapid charge cycle.

# Contents

<b>Acknowledgements.....</b>	<b>6</b>
<b>Author's Declaration.....</b>	<b>7</b>
<b>1 Introduction.....</b>	<b>8</b>
1.1 High Performance Computing and Power Efficiency.....	9
1.2 CPU Power Consumption vs. Computational Performance.....	11
1.3 CPU Power as a Proportion of System Power Consumption.....	13
1.4 Comparison of a Multicore CPU vs. Array of Low Power CPUs.....	15
1.5 Manufacturing Costs – Multicore CPU vs Single Core RISC CPU.....	17
1.6 Improved Power Efficiency Beyond CMOS Technology.....	17
1.7 Summary.....	18
<b>2 Low Power CPU Cores Suitable for a Large Array.....</b>	<b>20</b>
2.1 Choice of CPU Core.....	21
2.2 Intellectual Property Cores.....	23
<b>3 Wire-Free Power Transmission Technologies.....</b>	<b>26</b>
3.1 Near Field Inductive Power Transmission.....	27
3.2 Near Field Capacitive Power Transmission.....	32
3.3 Far Field (Radio Frequency) Power Transmission.....	33
3.4 Optical Power Transmission.....	36
3.5 Mechanical Power Transmission.....	38
3.6 Chemical Power Transmission.....	39
3.7 Summary.....	40
<b>4 High Density Power Storage Technologies.....</b>	<b>42</b>
4.1 Silicon Nano-wire batteries.....	42
4.2 Air Fuelled Lithium Ion Batteries.....	43
4.3 Supercapacitors.....	44
<b>5 Cooling Techniques for a Large Array.....</b>	<b>46</b>
5.1 Power Management in Current Commercial CPUs.....	46
5.2 Estimating Cooling Requirements for a Large Array.....	48
5.3 Considerations for Packaging Materials.....	51
5.4 Dissipating Waste Heat.....	53
5.5 Summary.....	55

<b>6</b>	<b>Wireless Data Communication Technologies.....</b>	<b>56</b>
6.1	Wireless Communication with Microwaves.....	56
6.2	Wireless Communication with Infra-red.....	58
<b>7</b>	<b>Modelling Wire-Free Power Transmission.....</b>	<b>60</b>
7.1	Physical Measurement of Inductive Transmission.....	60
7.2	Software Modelling of a 2-D Array.....	62
7.3	Initial Results of 2-D Software Modelling.....	66
7.4	Software Modelling of a 3-D Array.....	74
7.5	Initial Results of 3-D Software Modelling.....	77
7.6	Summary.....	85
<b>8</b>	<b>Conclusions and Future Work.....</b>	<b>87</b>
8.1	Conclusions.....	87
8.2	Future Work.....	89
8.2.1	RF Power Transmission.....	89
8.2.2	Optical Power Transmission.....	90
8.2.3	Data Communications.....	91
8.2.4	Cooling System.....	93
8.2.5	Node Internals.....	94
8.2.6	Node Packaging.....	96
8.2.7	Containment Vessel .....	97
8.2.8	Software.....	99
8.2.9	Summary.....	100
	<b>Appendix.....</b>	<b>101</b>
A.	CPU Statistics.....	101
B.	Example of Raw Array Simulation Data.....	102
	<b>References.....</b>	<b>104</b>

## List of Tables

Table 2.1: Low Power Processor Cores Compared with a GPU.....	20
Table 3.1: Limits for General Population/Uncontrolled Exposure.....	35
Table 3.2: Comparison of Solutions.....	41
Table 7.1: Comparison of 12x12 Array with Single and Multiple Antennae.....	73
Table 7.2: Comparison of 8x8 Array With/Without* Optical Power Sharing.....	79
Table 7.3: Comparison of 12x12 Array With/Without* Optical Power Sharing....	80
Table 7.4: Comparison of Arrays for Optimal Power.....	82
Table 9.1: Comparison of CPU power/performance.....	101

## List of Figures

Figure 1.1: A Large Array of Computing Nodes (Balls).....	8
Figure 1.2 : CPU Power/Performance Ratio Trend.....	12
Figure 3.1: Taxonomy of Wire-Free Power transmission Technologies.....	26
Figure 3.2: Efficiency of Inductive Power Transfer.....	29
Figure 3.3: Employing “Power Rods” for Inductive Power Transfer.....	30
Figure 3.4: Close Coupling of Windings.....	31
Figure 3.5: Test Circuit.....	33
Figure 4.1: Silicon Nano-wire Battery.....	42
Figure 4.2: Lithium-Air Battery.....	43
Figure 4.3: Supercapacitors Compared to Batteries / Fuel Cells.....	44
Figure 5.1: Temperature Rise Against Coolant Flow Rate for Water.....	49
Figure 7.1: Inductive Power Transfer Test Rig.....	60
Figure 7.2: Tx/Rx Coil.....	60
Figure 7.3: Schematic of Improved Test Rig.....	61
Figure 7.4: 8x8 Array in the Process of Charging.....	66
Figure 7.5: Performance of a 12x12 Array.....	68
Figure 7.6: Performance of an Inner Node.....	69
Figure 7.7: Performance of an Outer Node.....	70
Figure 7.8: Power and CPU Utilisation (%) with Single Antenna.....	71
Figure 7.9: Power and CPU Utilisation (%) with Multiple Antennae.....	72
Figure 7.10: Comparison of System Stability in 12x12 Arrays.....	73
Figure 7.11: Numbering Nodes in the Array.....	74
Figure 7.12: Temperature Rise of Coolant.....	78
Figure 7.13: 12x12 Array Running at 100% CPU Utilisation.....	84
Figure 8.1: Test Rig for Optical Data Communication.....	92



## Acknowledgements

York College (Employer) for co-sponsoring this MSc work.

Amir Mansoor Kamali Sarvestani, Ph.D. Student, University of York, working on the wireless communication technique for the Ball Computer referred to in Chapter 1.

Mike Hume, Tutor of Engineering, York College, for assistance with the Solidworks 3-D modelling software for figure 1.1.

Richard A. Clarke, Lecturer in electromagnetic theory, University of Surrey for producing a simple software model of adjacent windings predicting a very low coupling, referred to in Chapter 3.

Richard Russell, Trustmarque Solutions, York, for advice on RF theory and antenna design covered in Chapter 3.

Prof. Ben Allen, Head of Centre for Wireless Research, University of Bedfordshire, for advice on RF power harvesting techniques covered in Chapter 3.

Dr. John Slattery, Dept. of Chemistry, University of York, for advice on the chemistry of Redox Flow batteries in Chapter 3.

Tim Tozer, Dept. of Electronics, University of York, for background information on long range data communication via IR laser – the HAPCOS project referred to in Chapter 5.

Stephen Pateman (1969-2013), York College, for boundless encouragement and inspirational line management.

Mr P. F. Hind, Retired Engineer, for proof reading final draft.

## **Author's Declaration**

I declare that, except where explicit reference is made to the contribution of others, that this thesis is the result of my own work and has not been submitted for any other degree at the University of York or any other institution.

Elements of this thesis, notably the discussion of wire-free power transmission technologies, will appear in a paper to be submitted to the IEEE Design and Test of Computers in December 2013 / January 2014.


# Chapter 1

## 1 Introduction

The concept of the “Ball Computer” as proposed by Prof. Jim Austin [1],[2] is a machine constructed from a cluster of processing elements or nodes. Each node is completely self-contained with a processor core and local memory, with wireless data communication to allow interaction with close neighbours.

Each node will receive power by wire-free transfer. It is suggested that these nodes will each be packaged in a spherical housing to facilitate efficient packing and the circulation of coolant. The cluster will be loosely contained in a 3-D structure (figure 1.1) and there will be no restriction on the nodes moving position within the array. The patent application suggests one method of supplying power is via light and using chilled water as the coolant. The feasibility of these ideas will be discussed in this thesis.

Figure 1.1: A Large Array of Computing Nodes (Balls)



It is envisaged that the nodes will be of a size 5 mm diameter or upwards, with 10mm to 50mm suggested as a reasonable size range to aim for. Throughout this discussion a diameter of 10mm will be assumed for the nodes.

With this very compact size for individual processing elements, it is anticipated that clusters could be as large as a million nodes, providing highly scalable, massively parallel, supercomputing capability with savings in overall cost and improvements in energy efficiency over conventional, hard wired arrays and multicore CPUs.

This thesis will investigate current technologies which could potentially be employed in a practical implementation of the Ball Computer concept: suitable (low power) CPU cores, wire-free power transmission, high capacity power storage, cooling systems and wireless data communications.

To begin with it is important to consider the trends in high performance computing and the constraints of power consumption, along with existing techniques for delivering multi-core solutions for the desktop in order to evaluate the potential benefits of the Ball Computer concept.

## **1.1 High Performance Computing and Power Efficiency**

According to Horst Gietl and Hans Meuer [3] supercomputers performing at 100 PetaFLOPS will be a reality by the year 2016. They predict high performance computer (HPC) systems will incorporate hundreds of thousands of CPU cores, maybe a million (although this has already been beaten by IBM, see next page). This will result in an inevitable increase in power consumption. However, the main problems will be non-uniform memory access and the reliance on parallel code to benefit from the highly parallel nature of the machines, as characterised by Amdahl's law [4].

There is a class of problems referred to as “embarrassingly parallel”, such as complex graphics rendering or simulation of particle physics which naturally lend themselves to massively parallel machines. The more processors available to the task, the quicker it is completed, a concept familiar to project managers, known as effort driven scheduling.

Massively parallel arrays are well suited to this class of problem, with the main strength of the Ball Computer concept being that it would be very easy to scale up a machine as required, simply by adding more nodes as required, in theory without even having to power down the machine, i.e. fully hot swap processing units.

However, most problems have a sequential component which reduces the effectiveness of parallel machines and the speed-up factor that can be achieved. As the number of processors becomes very large, Amdahl's law shows that the actual speed-up will tend to a limit of  $1/(1-P)$  where P is the proportion of code that will execute in parallel. Even with a very small proportion of “non-parallelisable” code, the speed-up factor will be poor in relation to the number of processing elements available. For example if just 1% of code cannot be run in parallel the maximum speed-up in a massively parallel array will be limited to 100, even with 1000 processors available. Again, the flexibility of the Ball Computer means that redundant processing nodes will be able to power down completely or simply be removed from the array with minimal disruption.

Gietl and Meuer mention current techniques employing GPUs for specialist Single Instruction, Multiple Data (SIMD) computations due to their relatively low power consumption and high performance. This is an area that is rapidly gaining in popularity for a range of applications [5] such as medical research simulation. Since the parallel computing capability of the graphics card is already available, where high performance, 3-D graphics rendering is not required, that capacity can be redeployed to execute parallel code. Any viable solution for the Ball Computer should be competitive in terms of price/performance with a GPU engine.

In July 2012 “Emerald” became the UKs fastest, GPU based supercomputer [6]. This supercomputer, built by Hewlett Packard, uses 372 NVIDIA Tesla T20A GPU and is ranked at 159 in the June 2012 Top500 list with a performance rating of 114.4 TFLOPS. The Tesla T20A GPU [7] contains 512 “CUDA” cores, giving “Emerald” over 190 000 cores.

Gietl and Meuer also discuss the increasing requirements of cooling and the development of green technologies such as recycling heat using liquid cooling.

They suggest that current HPC systems use anywhere between 50% and 70% of their total power for cooling. Again, this is an area where the Ball Computer concept offers a very neat solution since the sealed, self-contained nodes are intended to be immersed in liquid coolant. The spherical design provides maximum surface area for heat transfer and the way in which the spheres naturally align allows plenty of space for coolant circulation.

The TOP500 List, which first appeared in 1993, provides a current list of supercomputers worldwide based on a standard benchmark (Linpack). According to Scientific Computing [8] twenty nine of the systems on the November 2011 list used more than one Megawatt of power, while the number one system, the Japanese “K Computer”, is quoted as having a power consumption of 9.89MW, delivering 10.51 PFLOPS using 705,024 SPARC64 processor cores.

However, since the original research was carried out for this work, Top500 has announced [9] that the number one position has now gone to an IBM supercomputer “Sequoia”, a BlueGene/Q system. It delivers 16.32 PFLOPS using 1,572,864 cores (1.6 GHz, 16 core, PowerPC processors).

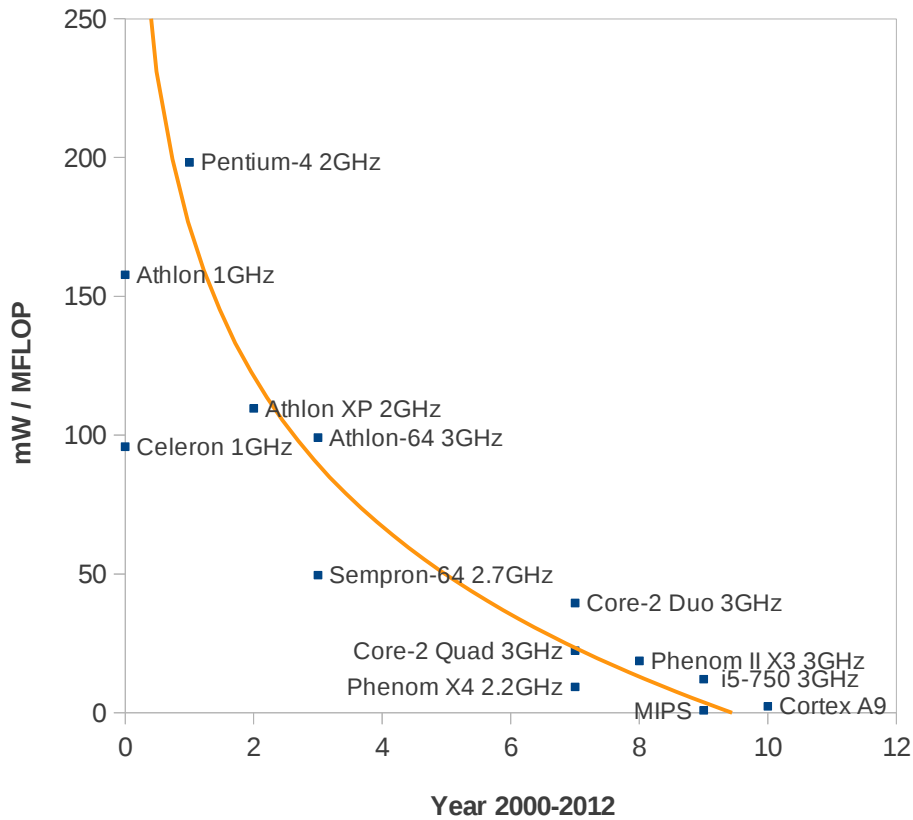
These are the benchmarks against which an implementation of the Ball Computer will be measured, capable of being scaled to a million plus nodes delivering PFLOPS of processing power while using a few Megawatts of electrical power.

## **1.2 CPU Power Consumption vs. Computational Performance**

A survey of a range of thirteen popular CPUs (including Intel, AMD, ARM and MIPS32 24K) shows that between 2000 and 2010 the ratio of computational power to electrical power consumption has dropped steadily as shown in figure 1.2, and it is reasonable to assume this trend will continue.

Looking at two CPUs from opposite ends of the trend line: in 2001 the Intel Pentium-4 2GHz required 91W of power to deliver 459 MFLOPS or 5402 DMIPS and in 2010 the ARM Cortex A9 running at 830MHz required 4mW of power to deliver 2075DMIPS.

Figure 1.2 : CPU Power/Performance Ratio Trend



The 2001 Pentium-4 has performance figures available in both DMIPS and MFLOPS, this was used to calculate an approximate conversion ratio for the purpose of creating the chart above [10]-[13]. The source data is included as appendix A.

The ratio trend is logarithmic ( $\log_n$ ), so it seems reasonable to assume that by 2014 a micro-controller delivering the same performance as a Cortex A9 based chip, could require as little as  $250\mu W$  of power. Ultimately the power/performance ratio will tend to a limit, with leakage current being the most significant factor.

Calculating the trend of Energy per Instruction (EPI) for Intel processors from 130nm to 65nm technologies (486 to Pentium 4) Grochowski and Annavaram [14] conclude that:

$$PowerConsumption = ScalarPerformance^{1.75}$$

This suggests that as CPU performance increases, the ratio of power consumption to performance will in fact gradually increase. So for example a CPU delivering 100 MFLOPS drawing just over 300mW of power would be operating at an efficiency of 3.1mW/MFLOPS. A similar CPU delivering 200 MFLOPS would then be expected to draw 1 Watt, delivering an efficiency of 5.3mW/MFLOPS.

Of course, modern CPUs are able to deliver higher performance through the use of pipelines and other architectural improvements without a proportional increase in power consumption. In fact the Intel Core Duo manages to achieve about the same EPI as an Intel 486 [14] which explains the results plotted in figure 1.2.

In terms of high performance computing, according to Scientific Computing [8] the average power efficiency of the Top500 supercomputers was 248MFLOPS/W in November 2010, an improvement from 195MFLOPS/W the year before.

The average power efficiency of the top ten supercomputers was 464MFLOPS/W. To compare this with the data shown on the previous page, that can be expressed as 2.2mW/MFLOPS, which rivals the estimated performance of the ARM Cortex A9 processor at 2.3mW/MFLOPS.

### **1.3 CPU Power as a Proportion of System Power Consumption**

When considering power consumption and efficiency, there are a number of factors to consider besides the power consumption of the CPU core. As previously stated, cooling systems currently account for 50% to 70% of total system power consumption. Conventional cooling systems which use refrigeration (heat pumps) require power to drive a compressor unit and air circulating fans in both heat exchangers. This is a very inefficient method of cooling as air has a low heat capacity and a low density, making it far less efficient than water. However, conventional server hardware uses forced air cooling, so this method is widely used in server rooms throughout the world.

The power efficiency of air conditioning units is quoted as the EER (energy efficiency ratio) which, for new systems, should be at least ten [15]. The EER is a ratio of the cooling, measured in BTU, to the power consumption of the unit, measured in KW.



The BTU (British Thermal Unit) is a standard way of expressing heating and cooling, where 1 BTU per hour is equivalent to  $2.931 \times 10^{-4}$  kW.

A commonly used approximation [16] for calculating the heat output of server equipment is:

$$\text{Formula 1.3.1} \quad \text{EquipmentBTU} = \text{TotalWattage} \times 3.5$$

Cooling will be investigated in more detail in Chapter 5. In that chapter the example used for calculations is an array which consumes 100W of power, generating approximately 3500 BTU of heat.

A typical small, portable air conditioning unit that can handle 8500 BTU requires 720W of power, which suggests approximately 300W would be required to handle the heat from this example array, confirming the estimates given above, in this case the cooling system would potentially be using 75% of total system power! Clearly there is a big opportunity to save power with alternative cooling techniques based on water as the primary coolant.

Shoaib Kamil et al [17] assume that typically a further 10% loss is due to power conversion i.e. from AC to DC. Switched mode power supplies are used as they are more efficient than linear (transformer based) power supplies as well as being more compact. According to Power Electronics Technology [18], switched mode power supplies can achieve efficiencies above 90%, with certain configurations achieving as high as 97%.

In multicore and multi-processor arrays, data communication becomes a significant factor in power consumption. It is also suggested that simply maintaining a “global clock” can consume as much as 25% of the core power [19].

As fabrication technology improves, the losses due to leakage (static power consumption) will continue to increase [20] as we move from 45nm, where there is typically a 50:50 static/dynamic split. H. Iawa [21] forecasts that by 2014, 22nm will be the standard gate size and it is expected that static power consumption will approximately double from the 45nm to the 22nm technology, while dynamic power consumption will not change (as clock frequencies remain the same).

The static (or leakage) power is exponentially proportional to the switching threshold of the transistors  $V_{TH}$  and the temperature  $T$ .

The switching threshold is the point at which transistors switch off and is decreasing as transistors shrink:

Formula 1.3.2 [22]  $Static\ Power \propto V e^{-k(V_{TH}/T)}$

These are all important considerations when reviewing alternative methods of power distribution where quoted efficiencies are low, since significant efficiency savings elsewhere in the system may well offset any losses due to non-conventional, wire-free power transmission techniques.

#### **1.4 Comparison of a Multicore CPU vs. Array of Low Power CPUs**

While multicore CPUs offer increased processing power without the need for higher clock speeds, due to the problems of heat dissipation it is found that 21% of a 22nm chip will be powered off [23] with this expected to rise to 50% at 8nm. This issue of so called “dark silicon” could be avoided with an array of low power, single core CPUs - if that offers equivalent (or better) performance in terms of computation, power consumption and heat generation.

An alternative view of this problem is - what could be done with the otherwise powered-down silicon? Of course it could then be dedicated to power harvesting, power sharing and wireless data communications in a self-contained node as proposed.

One of the key considerations has to be price versus performance. There are few direct comparisons between the performance of a single core, 32-bit ARM based CPU and a 64-bit, 4-core Intel CPU. However, a simple comparison of performance in decompressing and compressing a ZIP archive [24] suggests that an Intel i7 3.2 GHz CPU is 30 to 60 times faster than an 800 MHz, ARM A8 based CPU, from the two benchmark scores quoted (the A8 CPU performs far worse on the compression algorithm, hence the dramatic variation in the figures).

The current best prices for the two CPUs in this comparison are: an Intel i7, 3.4 GHz at £208 [25] and a Freescale i.MX515 at £15.99 [26] in the UK. A very rough calculation suggests the i7 offers twice the performance of an array of thirteen i.MX515 processors, for the same price. The i7 also has the clear advantage of being in a single package.

The strength of an ARM core, however, is the very low power consumption, for the A8 core it is 300 times less than an Intel i7 according to the data available.

In theory then, an array of thirty ARM core processors could deliver similar computational power using only one tenth of the the electrical power and no powered off silicon in active nodes.

Another issue with multicore processors is the shared memory bus (and level-3 cache in many cases). Samuel Moore [27] states that “At the heart of the trouble is the so-called memory wall—the growing disparity between how fast a CPU can operate on data and how fast it can get the data it needs.”

The results Moore presents suggest that performance begins to fall off at eight cores and above, at sixteen cores the performance is no better than a dual-core CPU. This is because the cores are unable to access the shared memory fast enough to keep them fully utilised. The solution suggested is “memory stacked on processor”, that is integrating the memory with the processor, which was being explored by Sandia National Laboratories [28].

An array of individual processing nodes, with local memory and high speed (wireless) connections to the nearest neighbours, is not going to be constrained by this “memory wall” issue as each CPU core will have a dedicated local memory bus. However, implementing shared memory through the node to node wireless connections may limit the applications of a system based on such an array since the throughput is going to be constrained by the bandwidth of the wireless communication channel used.

## **1.5 Manufacturing Costs – Multicore CPU vs Single Core RISC CPU**

It has not been possible to obtain precise statistics on current silicon yields for multicore CPU fabrication and according to the Australian Tech News web site [29] “Intel is believed to hit over 90%, although numbers are never talked about”. It would be helpful to understand the cost implications of low yields for multicore fabrication in order to better understand the relative merits of employing simpler CPU cores in an array.

The difficulty in obtaining this information is further illustrated by the Wikipedia entry [30] on multicore processors, saying “Integration of a multicore chip drives chip production yields down and they are more difficult to manage thermally than lower-density single-chip designs.”

One workaround for the lower yields of quad-core CPUs being employed by AMD is to disable a faulty core and release the CPUs as triple core instead. The fact that this is done quite openly suggests that failed cores are indeed common in the production process of commercial multi-core CPUs.

It seems reasonable to assume that a solution based on multiple, RISC or at least less complex CPU cores, is going to have advantages in terms of manufacturing costs (better yields), requirements for cooling and avoiding dark silicon.

## **1.6 Improved Power Efficiency Beyond CMOS Technology**

Each new generation of processor has seen a switching speed increase of 41% [22] simply due to further miniaturisation of integrated transistors, but there is a practical limit to how fast conventional CMOS technology can operate because the transmission speed on the interconnects do not scale at the same rate.

In terms of the power efficiency for the current generation of CPU cores, this too is dependant on the gate size and chip complexity which increases as the transistors shrink. It has recently been announced [31] that researchers at Northwestern University (USA) have developed complete logic circuits using quantum mechanical principles, called “spin logic circuits” or “spintronic” logic which makes use of the so called spin property of fundamental particles e.g. electrons.

As this technology uses single Magnetic Tunnel Junction (MTJ) devices [32] that provide the same operation as each pair of MOS transistors, they are much more efficient, in fact it is claimed [31] that this technology could be up to a million times more power-efficient than standard CMOS transistors.

As MTJ technology is already in use for non-volatile memory – Magnetic Random Access Memory (MRAM) – it opens up the concept of “logic in memory” [32] bringing the actual processing into the memory and reducing the communication bottleneck.

## **1.7 Summary**

The first step for the Ball Computer concept is to identify a CPU core with a sufficiently low power requirement as to work with a wire-free power transmission system and at the same time deliver an adequate level of computational power to make the resulting machine competitive with contemporary, massively parallel HPCs. This is the subject of chapter 2.

It has been shown that the trend over the past ten years has been a steady improvement in the ratio of power consumption to computational power, with recent developments aimed at the smart phone market yielding some ultra low power CPU cores and a range of novel developments aimed at further improvement in efficiency.

However, the CPU core is not the only consideration in terms of power efficiency, any losses due to wire-free power transmission will have to be offset by improvements elsewhere in the system – data communications and I/O as well as local memory and level-1 / level-2 cache.

It has been suggested that an array of single core CPUs could offer many benefits over multi-core solutions from power and heat management to ease of manufacture. Clearly the manufacture of a self contained node will bring it's own challenges, but techniques such as stacking of silicon could be employed.

However, as the main aim is to improve power efficiency this is not an issue.

Going beyond the current technologies, alternatives to CMOS switches offer further improvements to power efficiency in both CPU core design and memory. With processor cores requiring only a few tens of micro-watts, potentially delivering TFLOPS of computational power, this will change the performance profile of HPC systems and make the Ball Computer a very feasible proposition. There still remains the issue of designing software to make optimal use of such massively parallel systems, but this thesis will focus primarily on the hardware implementation of such an array.

The main emphasis of this work is to recommend a solution for delivering power without conventional physical wiring, which will be covered in Chapter 3. A range of potential solutions has been investigated and that chapter will explain how each could be applied to the Ball Computer and highlights any short-comings, concluding with some recommendations for further investigation.

Chapter 4 considers a solution for local storage of power within each node, due to the power requirements and limited physical space, any solution will have to offer both high power density and a fast charge cycle (discharge-recharge time).

CPU power management and system cooling will be covered in Chapter 5, considering the issue in terms of achieving the most efficient use of power and investigating the potential of a water based cooling system.

Chapter 6 will give a brief overview of the two wireless data communication technologies that could be employed for node to node communication. This area of the project is currently the subject of PhD research at the University.

In Chapter 7 this information will all be drawn together into modelling and simulation of power transmission within a full array, in order to investigate the overall performance of a Ball Computer implementation based on the technologies identified as having potential.

Finally, Chapter 8 will summarise the findings and recommend areas for further research and development, and relevant lines of enquiry.

## Chapter 2

### 2 Low Power CPU Cores Suitable for a Large Array

Before power supply and cooling requirements can be calculated, a suitably low powered CPU core must be identified. A range of candidates has been considered for the Ball Computer CPU core and some of the more noteworthy are summarised in Table 2.1 below:

Table 2.1: Low Power Processor Cores Compared with a GPU

Attribute	MPIS32 24K	ARM Cortex M0	ARM Cortex A5	Nvidia G80 GPU
Bus width	32 bits	32 bits	32 bits	64 bits (x6)
Power consumption	100mW @ 1GHz	4.25mW @ 50MHz	0.12mW/MHz	100 W + 780mW/core
Max Clock Speed	1468 MHz	50 MHz	600 MHz	612 MHz
FLOPS	734 MFLOPS*	50 million Mult/S**	1.3 Mflops/MHz	384 GFLOPS
DMIPS	1.51 / MHz	0.9 / MHz	1.57 / MHz	Unknown
Fabrication Technology	40nm	180nm	40 nm	90nm
Architecture	RISC, 8 stage pipeline	RISC, 3 stage pipeline	SIMD Extensions In-order pipeline	SIMD, Scalar pipeline
DMIPS/mW	Unknown	12.5	13	Unknown
Transistors	Unknown	12 000	Unknown	681 million
Notes	n/a	No FPU!	Up to 4 cores	128 Cores

\*A 32 by 32 bit divide operation is quoted as taking 2 clock cycles, from this a worst case estimate of FLOPS has been made for the quoted clock speed (also given as worst case).

\*\*A 32 bit hardware multiplier is quoted as completing one operation per clock cycle.

The DMIPS measure is a Dhrystone benchmark based on the performance of a VAX 11/780, which is taken as the baseline for a 1 MIPS machine.

According to ARM Technical Support [33] the DMIPS score is the Dhrystone score

divided by 1757 which is the Dhrystone per second score from the VAX 11/780, which gives a good comparative score independent of the instruction set and architecture.

Further research lead to the Tiler [34] Multi-core CPUs. The Gx3100 provides 100 general purpose cores in a single package, with a quoted power consumption of 20-48 Watts for standard application - this chip offers their best performance per Watt. It also incorporates a 1/10 Gbps Ethernet interface.

An interesting feature is the “intelligent mesh”, or iMesh technology, which is a two dimensional switched network that interconnects the cores at terabit speeds. Another benefit of this architecture is the distributed level-2 cache. These are two concepts which are a key part of the Ball Computer, which introduces a third dimension to such a network.

Specification data is limited and no useful benchmark data is provided on the official web site. Using a very rough and ready calculation gives an estimated power consumption of 480mW per core, operating at 1.5 GHz, which will include power for inter-core communications. This makes it an unsuitable candidate as low power consumption is critical.

## **2.1 Choice of CPU Core**

The proposed choice of CPU core for initial development is an ARM Cortex-A5. The single core A5 running at 530-600MHz is quoted as delivering up to 13 DMIPS/mW. Using the same conversion ratio as in the table in appendix A, this is something to the order of 0.95mW/MFLOPS.

Running at full speed of 600MHz, an A5 could deliver 180MFLOPS for 72mW of power (assuming correct conversion from DMIPS). It should be noted that to rival the best computational performance of the current TOP500 supercomputers an array based on the A5 would require some twenty million cores!

The Cortex-A5 is currently available in one to four core configurations, using the standard 32-bit ARM architecture, which incorporates standard RISC features and Java acceleration technology (known as Jazelle).



There is an optional “VFP”, floating point unit, which is IEEE 754 compliant (VFP version 4). There are separate 64KB data and instruction caches and an integrated memory management unit.

There is also a range of System IP Components which provide (system on chip) support for a level-2 cache controller, a dynamic memory controller and a direct memory access (DMA) controller. The Memory Management Unit works with two levels of cache and implements a full virtual memory architecture – known as VMSA (ARM v7 Virtual Memory System Architecture).

Good power management at the node level is very important to the Ball Computer and the Cortex-A5 incorporates four levels of power management [35]. To save local power reserves the core can enter a standby mode while waiting for large data transfers (as the remainder of the system on chip remains fully powered). In this state most of the clocks are disabled so the only power drawn is due to leakage (i.e. static power). The standby state is entered in response to a Wait For Interrupt (WFI) instruction, before completing the transition all current memory accesses are completed (using a synchronisation barrier instruction).

The processor will provide an external signal to indicate its status to external support circuits (called STANDBYWFI). Run mode is resumed when an interrupt occurs, or an event arrives, or there is a reset. It also happens in response to a debug request, but this wouldn't be relevant to the live system. In the case of a large data transfer it would be the DMA controller that would trigger the interrupt to wake the processor.

Where a node has completed and/or terminated a process, the processor could then enter into the dormant mode. In this state only the two caches and translation lookaside buffer are left powered up. This allows the processor to resume run mode much quicker than if it were in the shut-down mode. Run mode is again resumed by the assertion of a reset signal.

Since dormant mode powers down much more of the core, all of the registers must be saved to external memory (this will be local to the node of course), along with any debug information (if applicable).

Once again a data synchronisation barrier instruction is issued to ensure that all memory access is complete before a WFI instruction is executed.

The A5 processor physically implements two “power domains” which enable the core to be powered down while leaving the caches etc. powered up. With a system on chip (SoC) implementation there is a third power domain for the support circuits. The SoC remains powered up when the processor is in dormant and standby (and shut-down) modes.

The ARM Intelligent Energy Management unit is another System IP Component available for the Cortex-A5 which provides the dynamic clock and the dynamic voltage scaling to run the processor cores.

Besides the excellent power efficiency and the reasonably good computational performance, the Cortex CPUs are fully customisable IP Cores, as explained in the following section.

## **2.2 Intellectual Property Cores**

Intellectual Property Cores or IP Cores provide a quick and effective way to design a custom SoC using proven technology [36]. This would enable a prototype to be developed very rapidly with an easy route to custom production.

The first step would be to design a Ball Computer node using simulation software and a development board. MIPS [37] provide a range of such development kits. “The Open Virtual Platforms (OVP) initiative ([www.OVPworld.org](http://www.OVPworld.org)) has high speed models available for most of the MIPS family of 32-bit processor cores”.

The next step would be a so called “soft IP core”, which offers greater flexibility during the early stages because the functionality is implemented in software running on a programmable chip rather than actual hardware units, which keeps the core hardware independent and much easier to modify during the development cycle.

The final step is to go for a “hard IP core”, whereby the core is actually manufactured in silicon using the same design as the soft core.

Moving to an actual hardware core will result in an improvement in terms of efficiency, greater speed (therefore processing capacity) and a reduction in power consumption and as a result heat generation.

The ARM IP solution is branded as Artisan [38] and claims to provide “efficient design of complex SoC solutions”. This offers a wide range of technologies from 250nm down to 20nm with so called “Advanced Physical IP” for 28nm to 20nm.

The ARM “Physical IP” provides a range of libraries for a full range of SoC components and processor design kits. The “Design Start” web portal offers free access to some of these tools for evaluation purposes. The ARM Electronic Design Automation (EDA also referred to simply as ECAD) deliverables include Verilog simulation models and place-and-route abstracts.

Verilog [39] is a widely use hardware description language (HDL) which has an associated international standard IEEE 1364. This would be used to design and verify the core / SoC, with the various IP components provided as ready built Verilog libraries (with timing and power models and advanced power format support ). The simulation models can be run as entirely software simulated or on an FPGA (the so called “soft core” stage). At this point the design can go through a number of cycles without major investment being needed.

Once the core / SoC has been thoroughly tested, the place-and-route abstracts are used to produce an actual integrated circuit layout. These layouts are then converted to masks, using the standard GDS II (Graphic Database System) file format. This would allow actual physical devices, which behave identically to the software model, to be fabricated.

There are currently fifteen foundries (fabricators) throughout the world manufacturing ARM (IP) cores and SoCs, which would not tie the Ball Computer project into one specific manufacturer. It is worth noting that one of the companies listed is IBM, which has already links established with the University.

The Life-cycle of IP core / SoC implementation is not a trivial matter but it is suggested [40] that the development of a Soft IP Core could be a suitable project for undergraduate students working in small groups.

The Soft IP Core development project would involve five key tasks:

First, an instruction set analysis is carried out, referring to the vendor's documentation. This allows the students to become familiar with the ISA and the IP components.

This leads into the register transfer level (RTL) design - decisions about registers and cache, pipeline and memory control are made. After an initial review of the ARM documentation it is not clear how much flexibility is available within the ARM IP core framework for custom RTL design.

Once an initial design has been completed it can then be tested in the development environment (provided as part of "Design Start" from ARM) to verify functionality. This is then turned into a Verilog model – but as has already been mentioned, the ready built Verilog library components are provided by ARM.

Using a software model allows the design to go through a number of iterations in order to optimise the performance and functionality of the core / SoC.

Once the Verilog model has been verified it can be implemented as a soft IP core on an FPGA or ASIC (application-specific integrated circuit) for testing.

Milenkovic and Fatzer [40] emphasise that such a project (theirs is based on a PIC18 micro-controller rather than an ARM core and uses VHDL instead of Verilog) will allow students to "gain real-world experience in developing IP cores." This is a very strong indication that the development of a soft IP core for the Ball Computer would be a realistic student project, once a suitably detailed specification has been created.

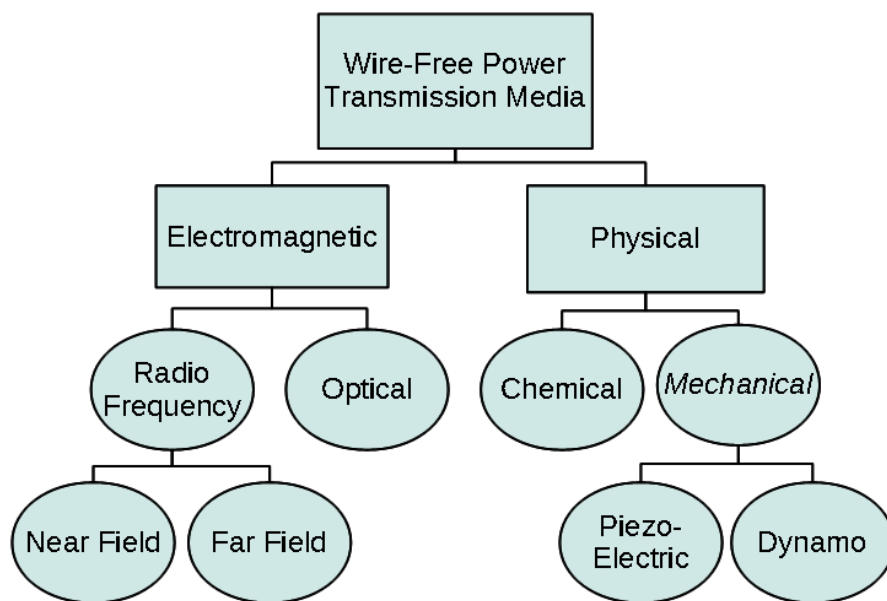
The specification of the Instruction Set Architecture and other SoC features is beyond the scope of the work carried out for this thesis. Since the design of the node CPU and support hardware is critical to the very nature of the final machine it could become a complete project in itself as it will have to include the I/O and data communications hardware.

## Chapter 3

### 3 Wire-Free Power Transmission Technologies

Assuming that a fully connected array of processing nodes using wireless communications is realistic then is there a solution for wire-free power transmission too? As the power requirements of CPU cores continues to decrease, as illustrated in chapter 1, wire-free power transmission does indeed become a more realistic solution as additional losses due the lower efficiency of a wire-free method will be compensated by the much lower power requirements for both the processor core, data communication between nodes and cooling system.

Figure 3.1: Taxonomy of Wire-Free Power Transmission Technologies



This chapter will consider eight methods for wire-free transmission of power, some of which are already used successfully for a wide range of other applications. These have been categorised as electromagnetic and physical as shown here in figure 3.1.

### 3.1 Near Field Inductive Power Transmission

Near field inductive transfer of power for rechargeable devices (from electric toothbrushes to mobile phones) is a successful technique when working at very close range. The Wireless Power Consortium web site [41] provides much useful background information, technical data and useful formulae.

This method uses the same principle as a transformer, where two sets of windings are magnetically coupled, allowing the current flowing in the primary winding (transmitter) to induce current in the secondary (receiver) via an alternating magnetic flux. The rate of change of the magnetic flux determines the EMF induced in the secondary windings, which means that working with higher frequency AC currents results in more efficient power transfer.

In a conventional transformer, the primary and secondary windings share the same iron core that tends to be laminated rather than solid to avoid eddy currents, which reduce efficiency. There is a simple formula to determine the secondary voltage  $V_s$ , based on the ratio of the number of turns  $N$  in each set of windings and the primary voltage  $V_p$ :

$$\frac{N_p}{N_s} = \frac{V_p}{V_s}$$

Using this method for wire-free power transmission involves separating the primary and secondary windings into a transmitter unit and receiver unit. If the windings have no ferrite core, referred to as “air cored”, then the magnetic coupling is only one quarter of an equivalent ferrite cored coil, due to the permittivity (magnetic conductance of materials) of iron being four times that of air. With a weaker coupling, the efficiency of power transfer is also reduced.

There is a lot of published research on the use of magnetic induction for close range power transfer: Xun Liu et al [42] demonstrate how a perpendicular magnetic flux can be generated with a hybrid of spiral windings and conventional coils, providing a recharging surface which allows multiple devices (placed directly on the surface) to be charged simultaneously.

Using conventional coils generates an uneven magnetic (flux) field which would affect the efficiency of power transfer depending on where the receiving device was placed on the charging surface. The addition of the spiral windings evens out the field. This technique can be used for recharging devices such as mobile phones.

An alternative approach was taken by Jianbo Gao [43] using multiple windings, in a similar configuration to a linear motor. The conventional windings are flattened to give a more even magnetic field across the surface. While the magnetic field strength is good at the surface, as distance between the surface and the receiver doubles from 8mm to 16mm the output voltage reduces to one third, as predicted by the calculations shown in the paper.

As far as commercially available solutions go, Texas Instruments [44]-[46] produce a range of wireless power transmitter / receiver chips which are Wireless Power Consortium compliant. The “TESLA” chips provide intelligent monitoring – scanning the surroundings for compliant receiver devices and will adjust output power to suit. The chips operate in the 110 to 210 KHz range, providing up to five Watts of power transfer at just over 50% efficiency (which was calculated based on the quoted transmitter supply of 19 Volts at 0.5 Amps).


The “TESLA” chipset is expensive, costing over £200 for a full evaluation kit and they are physically large (approximately 50 x 50mm) requiring bulky external, passive components, including the primary/secondary windings.

These methods all have limitations in the context of a 3-D array of processing nodes, as the coupling efficiency falls off rapidly with separation [47]. The plots in figure 3.2 on the following page show the power transfer efficiency (y-axis) for five different pairs of windings: 1:1, 3:1, 10:1, 30:1 and 100:1 ratios of diameters.

As the ratio of separation distance to primary winding diameter (x-axis) increases, the efficiency of the transfer falls off rapidly beyond about one tenth i.e. a pair of 5mm diameter coils would need to be no more than 0.5mm apart. The plots are based on coils with a Q-Factor (See section 7.1) of 100.

The Q-Factor (Quality) is a ratio of the inductive reactance to the ohmic resistance of the winding at a given frequency – reactance is the effective resistance to alternating current due to an inductor or a capacitor. Inductive reactance increases with frequency whereas ohmic resistance remains constant (other than increasing slightly with temperature) and is proportional to the cross section area and length of the wire used to form the inductor i.e. thinner, longer wire has a higher resistance.

Figure 3.2: Efficiency of Inductive Power Transfer [41]




Alignment of primary and secondary windings is also critical to efficient power transfer, with the magnetic coupling also being proportional to the cosine of the angle between the axes of the windings' cores [48],[49].

An early solution which was proposed involved the use of “power rods” to distribute the power within the array. An external winding (possibly surrounding the entire array of nodes) would be a neat solution, but the formulae clearly show that this would be ineffective due to the high ratio in primary and secondary winding diameter, as well as the distance between them.



So instead, a series of windings, of the same diameter as the nodes, would be housed in cylinders (power rods) and positioned so as to form part of the hexagonal packing of the layers of spherical nodes (see figure 3.3 below).

Figure 3.3: Employing “Power Rods” for Inductive Power Transfer




This arrangement would overcome to the two main factors – relative diameters of windings and proximity. If the rods were also hollow they could perform a secondary function of circulating the liquid coolant.

There would also be a requirement for the secondary windings (in the nodes) to align in the same plane as the primary windings in the power rods. To achieve this the nodes would be designed to have neutral buoyancy in water, but weighted so they float in the same vertical orientation. The windings would be positioned at 90 degrees to the vertical, thereby ensuring the correct alignment.

In the left hand side of figure 3.3, a cross section of a power rod is shown, with eight layers of (ball) nodes arranged around it. The lower left node has the secondary windings indicated as a black band and the adjacent primary windings are shown as black stripes. The broken line indicates coolant flow through the centre of the power rod (from beneath the base of the containment vessel).

The right hand side of figure 3.3 shows a view from above, with the power rods shown as solid black circles and the nodes as the grey filled circles. Careful consideration would have to be given to the layout in order to maximise the number of nodes available. Here there is one node which is adjacent to two power rods, it would be desirable to have nodes only adjacent to one power rod (although this does result in an array where a small proportion of nodes are not adjacent to any power rod and would therefore receive proportionally less power). Unfortunately, it was found by software modelling (c/o Mr Richard A. Clarke) and by physical experiment, that coils placed adjacent in the same horizontal plane, have minimal coupling while maximum coupling is achieved when the windings are coaxial. In an array of spherical nodes, each of which could lie in any orientation, this critical alignment could not be guaranteed. This result lead to the abandonment of the power rod idea. However, further discussion lead to the idea of using inductive power transfer to allow nodes to share power with immediate neighbours.

Figure 3.4: Close Coupling of Windings



To achieve the best coupling between nodes, flattening two opposite faces of each spherical node to give a 2.5mm radius face (on a 10mm sphere), the windings can be placed as close as possible. Figure 3.4 shows how a pair of ferrite coil formers could also be used to maximise the magnetic coupling. In theory this could be as high as 90-95% (ignoring copper losses and hysteresis losses in the ferrite core/former) using coils with a suitably high Q-Factor.

In order to maintain alignment (hexagonal packing), the inner sides of the containment vessel would need to be moulded to allow the spheres to form into hexagonal close packing, while the front and back can be flat to mate with the flat faces of the nodes. A planar charging surface [43] could be used to transfer power to the outer layer of nodes these then transfer power node to node into the array.

### **3.2 Near Field Capacitive Power Transmission**

In 2010 Murata [50],[51] announced the release of a practical wire-free power transfer system based on capacitive coupling. The capacitive effect relies on a good dielectric material, a large surface area for the capacitor plates and close proximity of plates. Rather than using a magnetic field (inductive) it relies on the electric field between the plates. The earth is then used to complete the circuit by providing a return path for the current.

Theoretically, the electrical field strength could be increased to enable capacitive coupling over a distance, but this would involve much higher voltages, possibly into thousands of volts (although at a proportionally lower current).

By introducing a second pair of plates (instead of relying on the earth to complete the circuit as in the Murata system) it could prove feasible to use a close range, capacitive coupling to power pairs of nodes.

Calculations based on a 10mm diameter node with six curved plates (arranged at right angles) on the surface, using distilled water (which has a high relative permittivity of 80) as the dielectric; a capacitance of approximately 100pF could be achieved between a flat transmission plate and the surface of a node. With the nodes being spherical, the capacitance between two adjacent nodes would only be around half that. Theoretically a low voltage DC power supply switched at 100KHz would lose about 20% of the power due to the reactance of the capacitors. As the frequency increase these losses will reduce, however there will inevitably be increased losses due to electromagnetic radiation, so an optimal switching frequency must be identified.

The simple circuit shown in figure 3.5, modelled as a network of capacitors and resistors, this appears to be too complex for the available circuit analysis/simulation software to cope with. In a simple test rig, a switchable clock generator was used to provide square waves at 250Khz, 500Khz, 1MHz and 2MHz. However, no substantial power transmission could be measured when working with low voltages (i.e. 12V and under) when measuring the voltage drop across each load resistor.


One interesting possibility does arise from this potential solution since the nodes are routing power between them instead of harvesting it from a general source. With intelligent power control circuitry, nodes would be able to divert power to their neighbours if their local power reserve was fully charged. This could easily be implemented with simple switches to bypass the local power store (high efficiency battery or supercapacitor) allowing more power to be routed to neighbours.

Taking this one step further, nodes could negotiate with neighbours to receive extra power when they have to perform a power intensive task. This is an idea that will be developed further in Chapter 7.

### 3.3 Far Field (Radio Frequency) Power Transmission

Power transmission via radio frequency is widely used for short, medium and long range, low power applications such as RFID tags and contactless smart cards.

Figure 3.5: Test circuit



Passive RFID tags operate at a range of a few centimetres to 10 metres [52] operating at frequencies of 124KHz, 125KHz, 135KHz and 13.56MHz for near field (inductive) transmission and 860MHz to 960MHz for far field transmission i.e. propagating waves.

The lower frequencies can penetrate solid materials while the higher frequencies offer a greater range – although it must be line of sight i.e. no obstructions. The low frequency tags use inductive coupling to transfer power whereas the high-frequency tags use propagation coupling i.e. RF power transmission.

Another application of far field RF power transfer is power harvesting (also known as energy scavenging), a technique used in conjunction with very low power devices where RF power is collected from the surrounding environment [53],[54] from sources such as radio and TV transmissions.

The Powercast P1110 chip [55] is a dedicated RF power harvesting chip which quotes greater than 70% efficiency, giving 4.2V at 50mA output for an input of 902-928MHz at 20dBm. At 14mm x 11mm (in the ceramic package) it is a realistic size for the proposed Ball Computer nodes.

Tajiu Suzuki et al [56] demonstrate how easy it is to achieve 20% efficiency over a distance of 10mm using RF at 250KHz with off the shelf components. They were able to demonstrate almost 50% efficiency at a distance of 4mm, transmitting over eight watts of power.

In considering substantial power transmission by RF it is essential to bear in mind safe exposure levels according to FCC guidelines [57] which are summarised in table 3.1 on the following page:

Table 3.1: Limits for General Population/Uncontrolled Exposure

Frequency Range (MHz)	Electric Field Strength (E) (V/m)	Magnetic Field Strength (H) (A/m)	Power Density (S) (mW/cm <sup>2</sup> )	Average Time  E  <sup>2</sup> ,  H  <sup>2</sup> or S (minutes)
0.3-1.34	614	1.63	(100)*	30
1.34-30	824/f	2.19/f	(180/f <sup>2</sup> )*	30
30-300	27.5	0.073	0.2	30
300-1500	--	--	f/1500	30
1500-100,000	--	--	1.0	30

f = frequency in MHz \*Plane-wave equivalent power density

A typical RFID tag requires around 50μW to operate, in order to supply that power at a distance of one metre, a 2.4W transmitter operating at 915 MHz would be required according to Lehpamer [58]. The power density will follow an inverse square law, so theoretically at a distance of only 100mm the transmitter power required falls to 24mW.

Assuming that each node in a large processor array requires only 100mW of power and the furthest a node will be from the transmitter is 100mm (ten nodes deep) then the required transmitter power would be at least 48W.

If 70% efficiency was achieved then presumably 330-340 nodes could be powered this way, to deliver enough power for an array 10x10x10 a transmitter of around 150W would be required, providing a strong enough field for the size of the containment vessel.

At centimetre wavelengths the safe exposure limit is just 1mW/cm<sup>2</sup>, so clearly some form of Faraday Cage would be required to contain the RF field, but one side effect of this could be further efficiency improvement as no power would be wasted outside of the array. Another issue to consider is protecting the active components from the intense electromagnetic field.

One factor that has not been considered so far is the attenuation due to the coolant, water being the proposed option. Since water is a dielectric (as mentioned in the previous section) there will be significant losses.

Any RF energy absorbed by the water will be converted to heat (the principle of cooking with microwaves) and so add to the requirements for cooling.

However, the single biggest challenge will be to construct a suitable antenna that is compact enough to fit into a 10mm diameter node. Consider a simple quarter wave dipole: the wavelength for 900MHz is going to be 333mm, so one quarter wavelength would be just over 83mm therefore it would require two arms each almost 42mm long. However, with a ground plane it only requires one half of the dipole, resulting in a single 42mm long antenna. While this is longer than the proposed node (10mm diameter) it could be achieved using a trailing wire antenna, with the body of the node forming a ground plane, alternatively a loop antenna may be more effective.

### 3.4 Optical Power Transmission

Conventional photovoltaic cells will produce 0.4 – 0.6 volts, converting up to 20% of solar energy into electricity. The low efficiency is because only a limited bandwidth of light will result in the photoelectric effect, longer wavelengths such as infra red, pass through the cell. The power density of solar radiation is typically quoted as 1.4 KW/m<sup>2</sup> [59] so a square metre of solar panel produces 280 Watts peak. Scaling down the cells to the 10mm spherical nodes discussed here, it is theoretically possible to generate up to 90mW at 0.6 Volts in full spectrum illumination. This is calculated from the standard formula: Surface Area=4  $\pi$  r<sup>2</sup>. Since the node is 0.01m diameter, r<sup>2</sup> is 2.5x10<sup>-5</sup> giving a surface area of  $\pi$  x10<sup>-4</sup> m<sup>2</sup>. Therefore a spherical PV cell of 10mm diameter, in full solar radiation, operating at 20% efficiency will provide:

$$\pi \times 10^{-4} \text{ m}^2 \times 1.4 \times 10^3 \text{ W/m}^2 \times 20\% = 90\text{mW}$$

Recent developments in Photovoltaic technology [60] using quantum dots or semiconducting crystals could improve the efficiency to as much as 50%. This could improve the above example to 250mW, which is theoretically enough power required to drive a low power CPU.

A successful application of power transmission via laser was recently announced [61],[62] when a Lockheed Martin Stalker Unmanned Aerial System (drone) was flown non-stop for 48 hours, using a ground based UV laser to supply power. This gave an improvement of 2400% compared to the flight duration when the drone is powered only by internal storage. On landing it was found that the drone had more energy stored internally than it did at the start of the test flight.

Another promising development in optical power transmission [63] is designed to deliver power to remote devices via a fibre optic (glass) cable in hazardous environments, such as aircraft fuel tanks. A laser acts as the source and the receiver is a gallium arsenide (or similar) chip which can be as small as 1mm x 1mm. The efficiency is claimed to be as high as 50% because the laser is tuned to the most efficient frequency for conversion to electrical energy. While the information available doesn't clearly state, it is assumed that the 50% efficiency include both the laser LED and the PV cell. However, according to RP Photonics [64] "Diode lasers can reach high electrical-to-optical efficiencies – typically of the order of 50%, sometimes even above 60%", which suggests that even with a very efficient PV cell and the system tuned to the optimal wavelength, the quote efficiency seems very optimistic.

However, while it might be viable to power a 2-D array optically, illuminating the PV cells in the middle of a large 3-D array will prove difficult. So far the only viable solution identified would be to use so called side-emitting (leaky) fibre optics [65] to carry the light into the array. Although further losses would be expected, reducing the overall power transmission efficiency even further.

Using artificial light sources would also provide an opportunity to use selective wavelengths, further improving the efficiency, although working in the UV range could be problematic as some wavelengths of UV are attenuated by glass.

A more promising approach could be to use optical power transmission for node to node power sharing (as mentioned previously), using surface mounted LEDs to transmit the power.



A development which could prove workable is a printed LED [66] which would allow the surface of each node to be covered with LEDs, this would make alignment less critical as nodes could negotiate which LED to use to achieve the best line of sight to each neighbour. Each node would have twelve, small photovoltaic cells mounted around its surface to collect the power.

### **3.5 Mechanical Power Transmission**

According to Dipen N. Sinha [67], power can be harvested from the mechanical vibrations in gas pipelines using piezoelectric materials, which is sufficient to drive small sensors (a few mW) using a 3cm x 3cm area of a 100mm diameter pipe.

The figure quoted for power generated by piezoelectric materials in this configuration is approximately  $200\mu\text{W}$  per  $\text{cm}^2$ . Theoretically with a 1cm diameter spherical node the amount power that could be harvested from its surface is using low frequency vibration is  $4\pi \times 0.5^2 \times 200\mu\text{W}$  or 0.6mW. Sinha states “Today’s vibration energy harvesters are so sensitive and efficient; they can generate electricity from vibrations that are barely noticeable to the human touch. ”, but clearly this solution does not provide nearly enough power for the proposed CPU core.

However, Tripathi and Khan [68] show how piezoelectric materials can be used to transfer power point to point, working at ultrasonic frequencies in the range 10 – 20 MHz, concluding that this method offers “a convenient way of wireless electricity transmission with negligible loss ”, although efficiency is not stated. In discussions with Jeff Neasham of Newcastle University, it was suggested that the efficiency of such a system using closely matched (tuned) piezoelectric crystals operating in the range 10 – 20 MHz and working in close proximity, could transfer power at an efficiency of 50%.

An alternative method investigated as part of this research is using micro-fans [69] as combined turbine/dynamos. Since these fans use brush-less motors it is not possible to generate a current from mechanical rotation (as in a conventional motor), however, if the driver circuit were to be bypassed then the alternating current generated in the fixed windings should be available as a power source. This would of course require custom manufactured devices, but the principle can be verified by modifying a standard sized, CPU cooling fan.

The advantage of this solution is that it could harvest energy from the circulation of the coolant fluid. However, to deliver this energy a pressurised containment vessel for the nodes is likely to be necessary as would correct aerodynamic design of the packaging to ensure that the fans/turbines align correctly in the flow. The question is whether a device physically small enough could deliver sufficient power.

### **3.6 Chemical Power Transmission**

IBM has recently announced the use of a “redox flow” battery to power an array of processors [70]. A redox (re-oxidation) battery uses two Vanadium electrolytes which are pumped through two “half-cells”, the oxidation that occurs results in a flow of electrons [71]. The open circuit potential is 1.4V at 25 °C and these batteries can respond very quickly to large changes in load. However, it should be made clear that the electrolyte, which is made from vanadium pentoxide dissolved in sulphuric acid, remains very acidic making it a hazardous substance.

IBM is developing this idea one step further by proposing to also use the liquid electrolytes for cooling, so as well as refreshing the electrolytes as they circulate a heat exchanger is also required. All of this being done with highly corrosive chemicals.

This idea does offer some interesting possibilities for the Ball Computer, if one electrolyte could be contained within in each node (and refreshed on a cyclical basis) the other electrolyte could be circulated as the coolant thus avoiding the need to pipe the two electrolytes separately. However, the energy/volume ratio is low compared with standard batteries such as lead acid cells.

Rather than using chemical reactions to generate power, could there be a way to extract energy from heat? Recently, researchers in China have found a way to harvest thermal energy using graphene [72]. They report an output of 0.35 Volts sustained for 20 days from a (toxic) copper chloride solution.

A graphene “device” 7 x 7mm, with gold and silver electrodes, is able to harvest thermal energy from the solution. By connecting six of these in series they were able to power a standard red LED ( $V_{\text{fwd}} = 2 \text{ V}$ ). Given the surface area of a 10mm diameter sphere ( $315 \text{ mm}^2$ ), in theory six such devices could be accommodated.

However, the down side to both of these techniques is the use of hazardous substances in a system which is intended for users to open up in order to add and remove nodes as necessary, which would expose them to these chemicals.

### **3.7 Summary**

Clearly there are a number of potentially viable options for powering a large array of self-contained computing nodes, although each method has its drawbacks. It is possible that a hybrid solution which employs two complimentary technologies is the way to proceed.

Two of the methods discussed that really stand out as potential solutions for the Ball Computer are near field (inductive) and far field (propagating waves) which both work on a similar principle, using electromagnetic waves at radio frequencies.

These are widely used in analogous applications and inductive power transmission in particular has proven to be safe and reliable in a range of domestic and industrial applications. However, there is a question of efficiency at small scale.

One potential approach is a hybrid solution using RF to supply power to all nodes within the containment vessel and then inductive or optical power transmission for nodes to transfer power to each other at close range. The nodes would have to be aware of their own power reserve and the likely requirement for a particular task. It would also require a communication protocol to be developed for nodes to negotiate requests for power from their immediate neighbours.

This protocol would have to take into account the relative orientation of a pair of nodes in order to achieve the best efficiency of the power transfer. This idea of node to node power transfer will be developed further in Chapter 7.

Below is a summary of the solutions discussed in this chapter:

Table 3.2: Comparison of Solutions

<b>Power Tx Method</b>	<b>Strengths</b>	<b>Weaknesses</b>
Near Field Inductive	Very widely used, safe	Only works at short range and relies on precise alignments of windings, efficiency falls at small scales
Near Field Capacitive	Appears to be efficient and simple to implement	Only works at very short range or requires high tension supply
Far Field Propagation	Widely used for low power devices	Limitations on power levels, RF screening required
Optical	Widely used on a large scale, safe	Currently inefficient, requires additional hardware to conduct light into the array
Piezo-electric	Efficient, simple to implement	Only efficient at short range
Dynamo	Widely used on a large scale	Requires pressurised coolant, may prove inefficient on a small scale
Chemical	Efficient, used in a range of different configurations	Redox requires electrolytes to be supplied separately. Other solutions may require toxic chemicals.

## Chapter 4

### 4 High Density Power Storage Technologies

It is unlikely that it will be possible to deliver continuous power to all nodes in a large array, considering the shortcomings of the various solutions discussed in the previous chapter. This introduces a requirement for some form of local power storage at each node. With the proposed size of nodes, this puts a constraint on the physical size of any power storage cell.


This chapter will consider some of the technologies currently available and technologies which may prove viable in the near future.

There are some interesting new developments [73] in battery technology, some of which may be applicable to the Ball Computer nodes: Silicon nano-wire and air fuelled Lithium Cells.

#### 4.1 Silicon Nano-wire batteries

In Silicon Nano-wire batteries the carbon anode of a conventional Li-ion (Lithium ion) cell is replaced with a stainless steel one which has silicon nano-wires covering its surface, in theory [74] this should store ten times as much charge.

Figure 4.1: Silicon Nano-wire Battery [74]



At present this is still in the experimental stage, but on the scale of the proposed

10mm diameter nodes it may be more viable than other storage technologies. This could work in conjunction with a wire-free power transmission system to maintain the charge, the quoted capacity is around 3500mAh/g.

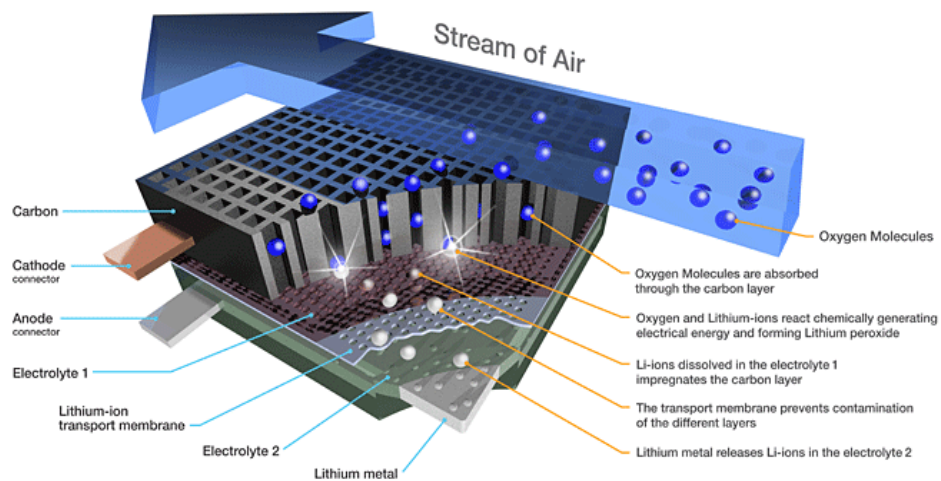
## 4.2 Air Fuelled Lithium Ion Batteries

In an air-fuelled battery, the cathode of a conventional Li-ion battery is replaced [75] with a porous carbon electrode (normally a lithium cobalt oxide electrode). Oxygen is drawn into the cell to react with lithium ions (within the porous carbon cathode), which means the usual chemicals are not required. It is claimed these batteries also have up to ten times the storage capacity of current batteries. By infusing oxygen into the liquid coolant, or by using a forced-air cooling system, the Ball Computer could potentially employ this type of battery technology if it will scale to work within the volume of 10mm nodes.

Figure 4.2: Lithium-Air Battery [75]

### Battery 500

The Battery 500 technology is an open system using common air as a reagent which upon recharge releases oxygen back into the environment.




For conventional lithium-based batteries the highest capacity [76] is 180Wh/Kg at 3.6V, which is equivalent to 50mAh/g, although it is suggested [75] that the theoretical capacity is as great as 372mAh/g. If it is assumed the air fuelled battery does indeed offer a tenfold increase in capacity, it could achieve 3700mAh/g which is even an improvement on the figure quoted for silicon nano-wire batteries.

### 4.3 Supercapacitors

An alternative to a conventional battery or fuel cell is a supercapacitor. Miller and Simon [77] suggest that a supercapacitor is comparable to a lead acid battery in terms of energy density, with a much better discharge cycle, as shown in the chart figure 4.3 below. A supercapacitor could be used in conjunction with wire-free power transmission to deal with varying power demands within nodes.

The physical constraints of a 10mm diameter node would currently limit the choice of supercapacitor to a (physically compact) “stacked coin” electric double-layer, 0.33F device as a maximum [78].

Figure 4.3: Supercapacitors Compared to Batteries / Fuel Cells [79]



A 0.33F supercapacitor fully charged at 5V (maximum rating 5.5V) will store just more than 4 Joules of energy. If a node requires 100mW at full power, this charge would provide around half a minute's worth of power. This means that nodes could be recharged in bursts instead of continuously, which may offer a little more freedom in the design of the wire-free power transmission system.

The lifespan of a typical electric double-layer supercapacitor is quoted as 15 years, although this will be reduced if the device is worked hard or kept in unsuitable conditions. For example this can be as low as 2.8 years [78] if it is kept at 40°C, in relative humidity of 50% and charged to 10% less than the quoted voltage. Considering the design of the Ball Computer, maintaining low temperature could prove problematic and as water is to likely to be used as the coolant, humidity could be an issue if the units are not well sealed.

It should also be noted that the physically smaller supercapacitors have a limited discharge current, rated 1mA maximum, and this would be insufficient to power a node. However, the more bulky radial type supercapacitors are much higher rated in terms of discharge current, but take up much more space and at present would be too big for the proposed node size, the smallest being 12.3mm diameter and 23mm in length.



## Chapter 5

### 5 Cooling Techniques for a Large Array

Clearly the advantage of both wireless data and wire-free power transmission is that the processing nodes can become totally self contained, allowing them to be completely sealed units and therefore totally immersed in a liquid coolant. This would provide a large surface area for cooling and being sealed units ordinary tap water could be used for example. Water is actually much more effective than air as a coolant, by a factor of 3000. This is due to the higher specific heat capacity (x4) and the greater density (x772) of water over air. As has already been mentioned, forced air cooling using a conventional air conditioning unit is very inefficient and typically contributes up to 70% of the total system power consumption.

While a pump will be required to circulate liquid coolant, the power required will be less than that required by the compressor in a forced air / refrigeration unit. This simplification of the cooling system will therefore produce further power savings, increasing the overall system power efficiency.

It is also important to consider the power saving techniques being used in current commercial CPUs and the associated heat management problems, which are addressed by such intelligent power management solutions. As has already been mentioned in Chapter 2, the proposed ARM Cortex-A5 incorporates such power management systems.

#### 5.1 Power Management in Current Commercial CPUs

As chip fabrication techniques have produced smaller gate sizes and more densely packed silicon, the heat density in CPU cores has been increasing, putting a greater reliance on the cooling system. According to Intel [80] their recent generations of CPUs have a power density between 10W/cm<sup>2</sup> (Atom) and 100W/cm<sup>2</sup> (Pentium 4), the current trend appears to be keeping within this range.

In order to manage heat there are a number of power management strategies

which are currently being employed within CPU cores. Dynamic voltage and frequency scaling (DVFS) is normally under the control of the operating system and is able to reduce power consumption by reducing the clock speed and/or core voltage of idle cores [81].

The Intel Core 2 Duo [82] has four power-down states or C-states (also known as sleep states) which are APCI (Advanced Power and Configurations Interface) compliant. In stage C1 a core is idle and execution is halted, by stage C4 the Core will also reduce its voltage to reduce the static power consumption.

AMD has developed Cool n' Quiet technology [83] which manages power in multi-core CPUs by gradually reducing clock frequency and powering down unused sections of the CPU depending on the computational power required at any time. AMD has also introduced "Smart Fetch Technology" which claims to "use fewer processing cycles to locate information since data storage is streamlined and stored in the shared L3 cache". The cached data is shared between cores which are switching between these different states under operating system control.

It is estimated that these techniques will reduce average power consumption by 30% with less than a 3% drop in performance [81]. One limiting factor with the dynamic voltage scaling is that all cores must be operating at the same voltage, so power down is limited by the most active core.

Another issue to consider is that of latency, i.e. the delay in getting back up to full power after a core has been in a powered-down state. It takes a finite length of time for the phase-locked loops to stabilise at the new clock frequency before operation can resume, the deeper the sleep state the longer the latency, which is another trade-off.

The ARM cores [84] use a similar principle, they have three reduced power states (as covered in Chapter 2). In the first stage (standby) the global clock is switched off. The next state (dormant) keeps hold of the program state and the core powers down but the cache is still powered. The final state (shut-down) requires the program state to be saved and the cache is flushed.

These techniques cannot react to changes in power any faster than tens of milliseconds, so to get very precise control over power new techniques are needed. Thread Motion [81] moves active threads between cores that are running at different clock speeds and voltages depending on their computation demands. Since this happens at the micro architecture level it can react in microseconds. However, as the proposed design for the Ball Computer envisages independent processing nodes which communicate via a wireless data channel, the overhead in moving a thread between two nodes is likely to far out-weigh any benefit from this approach.

Another novel approach comes from the GreenDroid project [85] which proposes to build specific software structures into hardware units to improve performance while reducing power consumption, moving away from the usual, general purpose CPU core. This technique is reported [86] to be eleven times more efficient than current architectures. With 120 specialist cores in a single package, it is claimed that the actual computational efficiency can be as much as 10 000 times greater than general purpose cores. The c-cores (conservation cores) are said to require eighteen times less energy per instruction [87] by translating code into specialist circuits.

While the Ball Computer concept is based on general purpose processing nodes, specialised nodes could be introduced to support specific computational tasks or the nodes could include a secondary, specialised core. Alternatively, since the ARM A5 supports up to 4 cores, each node could contain a general purpose and three specialised cores which could be independently powered up depending on what sort of processing task the node is required to carry out. It is not clear from the ARM documentation whether this level of customisation is possible though.

## **5.2 Estimating Cooling Requirements for a Large Array**

To get an idea of the cooling requirements, first consider a 3-D array of 1000 ceramic packaged computing nodes arranged in a cube of 10x10x10 (each node being 10mm in diameter) running at a maximum of 100mW per node, resulting in a total heat output of near enough 100W.

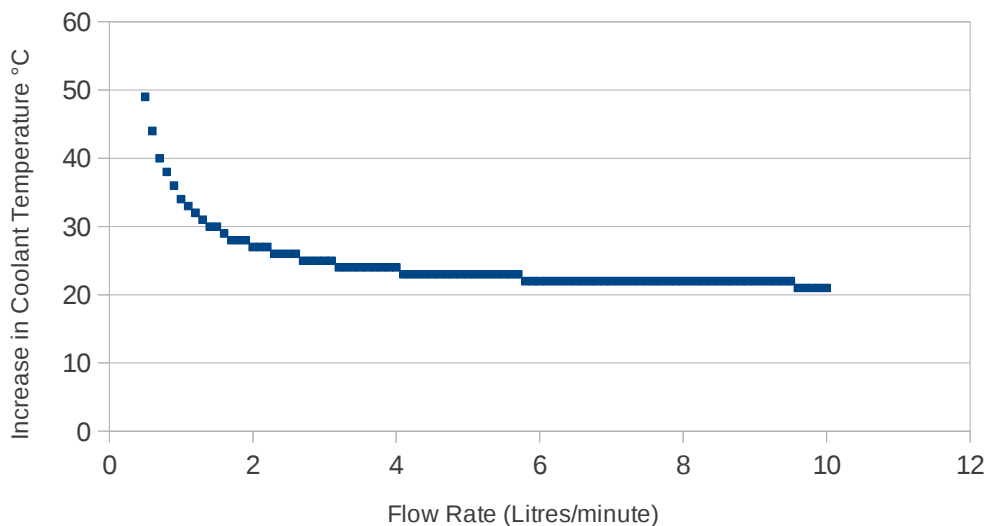
If the spheres are arranged into hexagonal close packing then the total cubic volume to contain the spheres will be approximately 87.5mm each side or 670cc, while the total volume of the spheres will be 524cc, leaving 146cc for coolant to circulate. (According the 18<sup>th</sup> Century mathematician Gauss, spheres in close hexagonal packing take up approximately 74% of the overall volume, which implies 174cc of coolant in this scenario, but this is a generalised solution).

Each node has an effective surface area of 315 mm<sup>2</sup>, so it will be assumed to have similar thermal resistance to a FBGA (Fine-Pitch Ball Grid Array) CPU in a 208 pin (15 x 15 mm) package, of approximately 20°C/W [88], giving a surface temperature of 1°C above the coolant temperature.

Using a simple formula [89] for calculating coolant flow, in this scenario the water temperature could be kept below 50°C (assuming incoming water temperature of 15°C) with a flow rate of just 0.5 litres per minute to deal with 100W of power consumption. To put this in context, a central heating pump capable of circulating just over 36 litres per minute requires only 35W of power [90].

Formula 5.2.1 calculates temperature rise as: 
$$\delta T = \frac{Power}{Density \times SHC \times Flowrate}$$

Figure 5.1 : Temperature Rise Against Coolant Flow Rate for Water



A basic aluminium heat-sink placed on an equivalent FBGA package (15 x 15mm) is calculated to run at 1.1°C above the ambient air temperature, while natural convection would dissipate this heat for a single device in a well ventilated enclosure, in a large array of closely packed elements forced air cooling would be required.

Assuming an equivalent array of 1000 such CPUs arranged over ten PCBs, with 5mm separation between the CPUs on a PCB and 30mm between PCBs to allow clearance for the heat-sinks, the total cubic volume of a suitable enclosure would be 12 000cc or 12 litres. Using the same formula for coolant flow rate, to keep the temperature below 50°C (assuming ambient air temperature of 20°C) would require an air flow rate of at least 28 litres per second i.e. more than the total volume of the enclosure. To put this in context, a standard cooling fan capable of an air flow of 17 litres per second runs at about 1.1 Watt, however, this does not take into account the resistance to air flow imposed by such a densely packed enclosure.

These illustrations do not take into account the method of removing the excess heat from the coolant, which will be carried out by some form of heat exchanger, to ensure that the returning coolant is at the quoted starting temperature.

Currently IBM is using distilled water as the coolant in its micro-channel architecture [91] but since the processors are in a hard wired array, the coolant is piped through each processor's heat sink. In the Aquasar project [92],[93] the waste heat is then used to heat the building instead of using conventional, costly air conditioning systems to remove the waste heat, supplying 9KW of thermal power to the building's heating system. It has been found that the Aquasar actually consumes 40% less power than an equivalent, air cooled system. The same idea could also be used to recycle the heat from the Ball Computer.

The other advantage of taking liquid coolant directly to the chips is that there is much less thermal resistance, so the coolant can run at the higher temperature of 60°C while keeping the components well within the 85°C maximum operating temperature.

Since the rate of cooling is proportional to the temperature gradient, cooling the water from 60°C when working at room temperature or outdoor temperatures can be done with simple heat exchangers much like a car radiator.

### 5.3 Considerations for Packaging Materials

In designing the packaging for the Ball Computer nodes it will be important to consider the thermal properties [94] of the materials used:

**Maximum temperature** – the maximum temperature indicates the point at which the material begins to break down, through chemical reactions such as oxidation or even mechanical failure of the material. If a maximum operating temperature of 85°C is assumed for the silicon core, this will set a lower limit for the temperature performance of the packaging material.

**Creep** – a measurement of how much a material will deform as it heats up (expansion). With regular heating and cooling the material could eventually fracture. In a densely packed Ball Computer system, the expansion of the material due to heating must also be allowed for in the tolerance of the containment vessel. With thousands of nodes even a small change in the dimensions of each unit could scale up to a significant change.

**Thermal conductivity** – a measure of how well the material conducts heat. For the nodes it is essential to dissipate the heat as quickly as possible to the circulating coolant. The temperature gradient i.e. the difference in temperature between the interior of each node and its surface will depend on the thermal conductivity. If the interior is allowed to become too hot then components will begin to fail.

It is also important to take into account the electrical and optical properties [94] which can have significant impact on the propagation of electromagnetic radiation used in communications and power transmission:

**Dielectric constant or relative permittivity** – a measurement of the insulating properties of the material in terms of its “tendency to polarize under the influence of an electric field.” The packaging for the nodes must be a good electrical insulator if the coolant is conductive, such as water.

**Dielectric loss factor** – a measurement of how much RF energy will be absorbed by the material. This is a critical factor if wireless data communication is going to be based on microwaves and if RF is used to transmit power to the nodes. If water is used as a coolant it will also attenuate microwaves. Any RF absorbed by the material will be dissipated as heat, so not only reducing the efficiency of wire-free power transmission but also adding to the problem of removing the waste heat.

**Dielectric breakdown potential** – a measurement of the potential (voltage) required to cause the insulator to fail. At the low voltages used within digital systems this is not a critical factor, but if capacitive power transmission was used then this would become a significant factor.

**Transmission** – a measure of how well a material allows light to pass through, which is critical if infra-red is chosen as the communication medium or node to node optical power transmission is employed. It may be that just small areas of the outer package are transparent to allow the passage of light or the entire package could be transparent. This offers an interesting option when considering the overall aesthetic of the Ball Computer – an area where it could really stand out from competitors.

**Refraction** – a measure of the variation in the speed of light as it passes through a particular material. If the package is spherical then there are going to be lens effects to take into account, which could prove beneficial in terms of making the angle of contact less critical. This is potentially another line of investigation in terms of modelling the behaviour of infra-red and ultra-violet light through a spherical lens.

Good electrical insulators with good thermal conductivity are primarily the ceramic materials. While glass provides good insulation, it is a worse thermal conductor than ceramics by a factor of a hundred or more. The only potential reason for using glass encapsulation is if optical is chosen as the medium for inter-node communication channels or if an optical power transmission method proves viable. Some ceramics do have a refractive index close to that of glass, and offer similar dielectric properties.

One approach to manufacture could be epoxy encapsulation, a standard technique

for producing hybrid devices (so called “Glob Tops”). The polyester resin that is widely used [94] has a high dielectric strength, an operating temperature of up to 165°C and a thermal expansion factor of just 0.11mm/m/°K, for a 10mm diameter sphere operating at up to 60°C above room temperature this equates to a change in diameter of 0.07mm. In an array of 100x100x100 nodes, a tolerance of at least 7mm in each dimension should therefore be allowed.

Polyester resin can also be finished so that it is transparent to light. However, the main appeal is cost, being cheaper than other types of epoxy resin. It is also relatively immune to the effects of UV, which can cause some resins to break down. If UV frequencies were used for optical power transfer this would be critical. On the downside, the mechanical properties are not quite as good as other resins and care has to be taken during manufacturing due to the chemical emissions. Another factor to take into account is that there is a significant reduction in volume as the epoxy cures – which is a very rapid process (introducing its own problems to the manufacturing process).

For a pilot production or proof of concept, epoxy encapsulation appears to offer the best solution, while ceramic packaging may prove most effective for large scale production.

#### 5.4 Dissipating Waste Heat

While heat calculations are fairly complex, an approximate calculation can be made for a heat exchanger using some basic heat transfer formulae [95]. The heat transfer rate Q (referred to as Power in formula 5.2.1) can also be expressed in terms of heat dissipation via a heat exchanger:

Formula 5.2.2 heat exchange:  $Q = U A \delta T$

Where: U=Heat transfer coefficient, A=Surface area,  $\delta T$ =Average temperature difference, calculated from 5.2.3:

$$\delta T = \frac{(T_{InHot} - T_{OutCold}) + (T_{OutHot} - T_{InCold})}{2}$$



This is based on the temperature of the hot water flowing in and out of the heat exchanger to the surrounding air which increases from “cold” to “hot”. The heat transfer coefficient from water to air via a simple copper heat exchanger is given as 13.1W/m<sup>2</sup>K [96].

Using the same example of 1000 nodes generating 100W power (section 5.1), with the output temperature being 70°C and the returning coolant temperature being 30°C, assuming ambient air (room) temperature of 19°C, which does not rise as a result of the heat exchange, gives a  $\delta T$  of 31°C. Rearranging formula 5.2.2 gives a required area of 0.25m<sup>2</sup> for the simple heat exchanger, which is approximately 490 x 490mm.

This can be compared to a domestic radiator [97] using a standard calculation for the effective heat emission, from the in flow and return flow water temperatures and ambient air temperature. A “type 21” (double panel, single convector) radiator sized 600 x 400mm or 0.24m<sup>2</sup> is quoted as having a power output of 768W (2619 BTU) at 60°C [98] which can be expressed as 3200W/m<sup>2</sup>.

Manufacturers assume a 60°C difference between average water temperature and the ambient air temperature when they quote the performance of a radiator, so the actual (or required) value has to be calculated, so again using a coolant temperature of 70°C and the returning coolant temperature of 30°C and ambient air (room) temperature 19°C we get:

Formula 5.2.4 actual temperature drop: 
$$\delta T = \frac{(T_{InFlow} + T_{Return})}{2} - T_{Ambient}$$

So in this case  $\delta T = 31^\circ\text{C}$ .

Formula 5.2.5 effective heat emission ( $E_2$ ): 
$$E_2 = E_1 \times \left(\frac{\delta T}{60}\right)^{1.3}$$

Where: E = heat emission in W/m<sup>2</sup>

So using a “type 21” radiator, the effective heat emission in this case is 1356W/m<sup>2</sup>. Therefore to dissipate 100W requires a surface area of 0.07m<sup>2</sup> or approximately 270 x 270mm. This suggests that modern domestic radiators are considerably more efficient heat exchangers than the generic heat exchanger used in the first formula, however the use of fins will increase the effective surface area of a standard radiator.

Considering the Aquasar approach [92],[93] where waste heat is transferred to a building's environmental system, a heat exchanger is using water on both sides (with a heat transfer coefficient of 340-355W/m<sup>2</sup>K) to transfer heat from the computer cooling system to the building's under-floor heating system. By running the coolant at the higher temperature the heat transfer is more effective and the large surface area of the underfloor heating pipes provides a very effective transfer of heat from the secondary system into the surrounding air.

If cooling can be achieved this way, then the only power requirement for the cooling system is around 30-40 Watts for a pump to circulate coolant. Now the total system power becomes 140W instead of the 400W for using air conditioning, so even if wire-free power transmission turns out to be only 30% efficient in the end, the overall system would still use less power in total thanks to the water cooling. This could prove to be a major selling point for the Ball Computer concept as the green computing agenda continues to develop.

## **5.5 Summary**

Power management and heat dissipation is a major concern for all large computer installations and there are many ideas which could be incorporated into the Ball Computer to address this. From intelligent power management at the chip level to environmentally friendly water cooling, the system power efficiency could potentially rival current HPC systems.

Clearly there are challenges from a manufacturing perspective, but there are viable solutions for packaging at both the prototyping and final production stages using proven techniques.

## Chapter 6

### 6 Wireless Data Communication Technologies

To support full communication between nearest neighbours in a 3-D array of processing nodes [1], assuming that each node is spherical and the array naturally forms into hexagonal close packing [99], initially it would appear that twelve channels of high bandwidth communication would be required per node. This is based on the fact that each spherical node would be surrounded by twelve immediate neighbours i.e. in contact.

The approach being taken in associated research at the University [100] is to arrange the nodes into zones of four nodes (as a pyramid formation) which is shown in figure 1.1 on page 8 (the four solid spheres). The central node, along with its neighbours, will then form eight zones (pyramids) which must use different channels to communicate in order to avoid overlap and interference.

Each pair of nodes within a zone is allocated two channels, which they share with nodes in the other zones that they are members of. It has been shown that to avoid overlap with other zones, a total of 88 channels is required. However, this figure does not change as the network (number of nodes) grows, which will guarantee a minimum throughput within a large array.

There are two media which will be considered for short range, wireless communications: microwave and optical (infra-red).

#### 6.1 Wireless Communication with Microwaves

A number of researchers have demonstrated the feasibility of using short-range microwave on-chip communication between multiple processor cores. This technique is comparable to conventional hard wired techniques in terms of energy requirements, it is suggested [101] that with microwave transmission power of -10dBm, the transmitter uses only 4.5 pJ/bit of energy.

One factor that needs to be considered here is the total power required by the transmitter, since it will consume more power than it is actually able to transmit. The paper discussed here does not make it clear how the figure 4.5 pJ/bit is calculated, but it does state that by using a very simple design for the transceivers an on-chip wireless network should only consume 1% to 2% of the total power drawn by the chip.

Assuming a bandwidth of 10 Gbps, the power required per channel will be to the order of 4.5mW with up to 16 channels available in the 100GHz to 500GHz range. Obviously this is nowhere near the requirements for the proposed zoning solution.

By way of comparison, the energy required to transmit data via a wired transmission line is quoted as 1.25pJ/bit [102] for small swing, current mode signalling operating at 500 Mbps and low power CMOS gates have a switching energy to the order of 0.1pJ/bit [103].

To get a better idea of the relative performance, take a conventional 32-bit data bus operating at 500Mbps (giving an effective bandwidth of 2GB/s). The total energy used per bus transfer therefore should be 10 pJ/byte, giving a power requirement of 2mW for a sustained transfer.

On-chip wireless data transmission has also been demonstrated to provide lower propagation delays [101] than conventional hard wired, bus architecture. This is a growing problem for large scale multi-core CPUs [104] as the increased bus lengths result in an increasing latency for inter-core communication. Another advantage is that cross-talk is reduced – the longer transmission lines become the greater the cross-talk (due to stray capacitance).

Current research [101] employs a 324 GHz carrier (for 90 nm CMOS) with this expected to increase to 600GHz and possibly 1THz for 16nm CMOS technology. Using carriers in the 100GHz – 500GHz range should provide a data transfer bandwidth of 10GBps and because of the very short wavelength, the antennae become very small, around 1mm for a dipole. If the proposed frequency for RF power transmission is around 900MHz then this could also provide sufficient separation.

Of course, current Wi-Fi technology should not be over-looked, the 802.11ac draft standard, due for approval in November 2013 [105], will offer “very high” throughput working in the 6GHz range, with 1.3Gbps likely to become the standard [106]. The solution for the Ball Computer is going to fall somewhere between these two technologies – high bandwidth, short range.

An alternative to Wi-Fi is Ultra-wideband (UWB) [107] which is being adopted for the IEEE 802.15 standard for wireless personal-area networks (WPAN), designed for use with devices designed to be worn or located near to the body of the user. This offers short range, high bandwidth communication in the 3.1 GHz to 10.6 GHz band with channels being at least 500 MHz wide. With 7500 MHz of spectrum available to UWB, this could provide 15 channels which is close to the number of channels available with Wi-Fi, however, each channel is capable of carrying over 200 Mbps, giving a total capacity of more than 3GBps.

While 15 channels is insufficient for the proposed zoning solution, the IEEE 802.15 standard is designed for multiple access, with the possibility of using frequency-division multiplexing (along with the more commonly used time-division multiplexing and code-division multiplexing).

One potential barrier to the use of microwave data communications is the proposal to use RF for power transmission as well. Provided much lower frequency is used (in the 900MHz range) and the WiFi antennae are designed to give a good rejection at the lower frequencies this should not present a problem. The maximum power for Wi-Fi transmission is 100mW (20dBm) which is a limit set by the EU for spectral regulation, although in the confines of a Ball Computer system this would be irrelevant.

## **6.2 Wireless Communication with Infra-red**

While peripheral devices using IrDA protocols seem to have fallen out of favour in recent years, infra-red is still widely used for wireless data communications and should be considered alongside microwave as a potential solution for the Ball Computer.

For example, the “Free-Space Optical Communications on High Altitude Platforms” project [108] proposes using Laser-based optical communications links (from high altitude platforms) with ranges of up to 600Km. It is stated that the advantages of using such Lasers include high data-rates (Gbps) and low power-consumption. The “High Altitude Platforms for Communication and Other Services” project is a European initiative to develop these long distance communication (and remote sensing) platforms. A High Altitude Platform (HAP) is defined as “airships or planes, operating in the stratosphere, at altitudes of typically 17 - 22Km (around 75,000 ft).” [109]

In terms of technology available on the scale of proposed Ball Computer nodes, Laser LEDs currently offers transmission speeds of up to 2.5 Gbps [110] and could potentially double as a means of power delivery too (see section 3.4).

The Infra-red Data Association or IrDA [111] announced in 2011 that a working group was being formed to develop specifications for 5Gbps and 10Gbps optical communication using “eye-safe” laser technology. This would provide a very viable alternative to microwave communications, with no problems of interference from the proposed 900MHz RF power transmission system.

More detailed information on the new standards (as of October 2011) from the 10GigaIR working group [112] suggest that the 10Giga-IR will be more than sixteen times faster than 802.11n WiFi. It will operate in the range 1cm to 10m using less than 5mW of power at the low end and will use wavelengths of 830nm to 1550nm (at the higher end of the infra red spectrum). It appears there will be two options in terms of the angle for line-of-sight operation, a narrow angle which is specified as less than five degrees and a wide angle, an upper limit is not given for the wide angle option.

One of the proposed applications is for device docking, which will work at distances of around 1cm, this would be perfect for the Ball Computer application.

The final specification was due to be released in the summer of 2012, following “ICTON 2012 : 14th International Conference on Transparent Optical Networks” (2<sup>nd</sup> – 5<sup>th</sup> July 2012) but at the time of writing (March 2013) no further information has been made available.

## Chapter 7

### 7 Modelling Wire-Free Power Transmission

To evaluate the feasibility of node to node power transmission as a method for delivering power (wire-free) in a large processor array a software model was constructed to investigate the result of losses and the maximum array size that could be powered in this way. But first it was essential to obtain some results for the efficiency of wire-free power transmission using the inductive method.

#### 7.1 Physical Measurement of Inductive Transmission

To investigate the efficiency of inductive power transfer at close range a simple test rig was constructed (Figure 7.1) composed of two CMOS ICs with a ceramic resonator to provide a square wave at 250KHz, 500KHz, 1MHz or 2MHz (via a DIP switch) and a MOSFET transistor (shown with heat sink) to buffer this signal to drive the transmission coil.

Figure 7.1: Inductive Power Transfer Test Rig





Figure 7.2: Tx/Rx Coil



The transmitter and receiver coils were separated by 1mm of plastic to simulate the body of the nodes. The coils consisted of a ferrite core/former and 24 turns of 24 swg enamelled copper wire (Figure 7.2).

Initially the output from the receiver coil was half wave rectified and smoothed with a capacitor. The coil properties were calculated approximately, based on the wire gauge, the diameter of the ferrite core and the number of turns used.

Formula 7.1 [113] for the calculation of inductance for an air cored induction coil is given as:

$$L(\mu H) = \frac{0.8(r^2 n^2)}{6r + 9l + 10d}$$

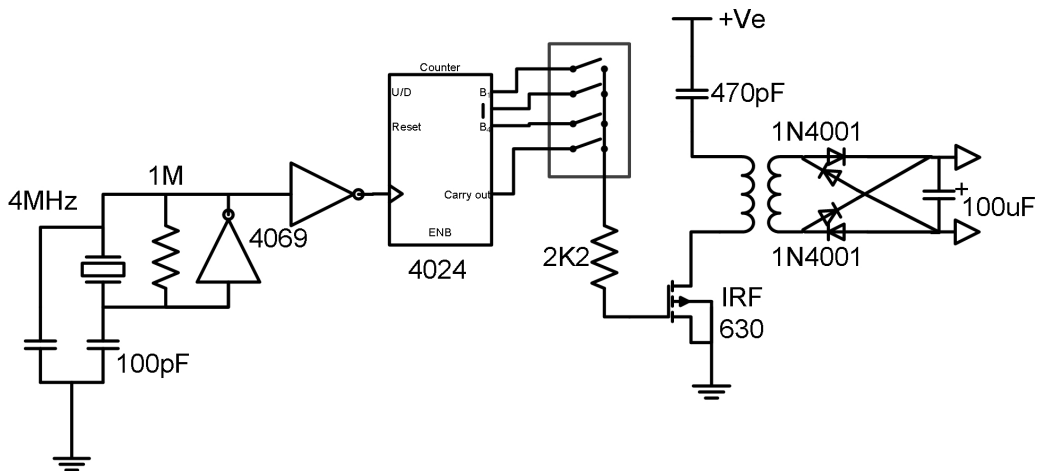
Where:  $r$  = mean radius from the centre to the middle layer of windings  
 $d$  = depth of windings,  $l$  = length of the coil,  $n$  = number of turns

*Note: All dimension are in inches for this formula*

The ohmic resistance was calculated and measured as  $0.04\Omega$  while the inductance was calculated as  $17\mu H$  and measured as  $15\mu H$  (the presence of a ferrite core should have increased the induction of the coil). Operating at 2MHz, the inductive impedance was calculated as  $188\Omega$  ( $Z_L = 2\pi fL$ ) giving a theoretical Q-Factor (quality) for the coil of 4700. For efficient power transfer the Q-Factor should be at least 100, which suggests that this configuration should work well.

Figure 7.3 shows the full schematic of the test rig, which was run at 6V, with the current limited to 250mA.

Figure 7.3: Schematic of Improved Test Rig





As the initial results were poor the transmitter circuit was tuned with the addition of a 470pF capacitor in series with the coil and to improve the output two extra diodes were added to provide full-wave rectification. However, the rig still failed to produce a measurable level of power at the output. The problem appears to be with the transmitter circuit delivering insufficient current to drive the MOSFET. The addition of a low power RF transistor in common emitter mode to act as a buffer between the counter (4024) and the gate of the MOSFET provided a good clean square wave when measured with an oscilloscope, however it was still unable to drive the transmitter coil sufficiently. This is a disappointing result and time did not permit further investigation.

The next solution to be investigated was far field RF power transmission. Due to the risks involved with working with high frequency RF, no physical test rig was constructed, instead a software model was developed.

## **7.2 Software Modelling of a 2-D Array**

To investigate the principle of hybrid RF and optical power transmission a simple software model was initially created to simulate a single layer of nodes in a 4x4, 8x8, 12x12 or 16x16 configuration. Each node is able to hold a local power reserve as charge in a 0.33F supercapacitor. The value of the supercapacitor was chosen based on the physical size – this being the physically largest device that could realistically be accommodated within a 10mm diameter sphere.

The amount of RF power transmitted into the array can be adjusted from 10 to 100 Watts, along with the antenna gain for the receiving antennae, which is -7 dB by default, this figure is quoted as the typical gain for a mobile phone or WiFi antenna. The transmitter antenna is set to unity gain (0dB), the typical figure for a half-wave dipole.

In the model all nodes received RF power from the transmitter antenna, which is located in the centre of the array. The nodes can also request power from their (six) immediate neighbours, using short range optical power transmission. It is assumed this will use a 50mW laser LED with 50% of the power being received, the remainder being dissipated as heat into the surrounding coolant.

To calculate the amount of power received (Rx) at each node, the model employs the Friis equation [114] as suggested by Prof. Ben Allen. It uses the power of the transmitter (Tx), the gain of each antenna, the separation distance (R) and the wavelength of the RF (calculated from the speed of light, c and the frequency of the signal, f) as Formula 7.2.1:

$$Power_{Rx} = \frac{Power_{Tx} Gain_{Tx} Gain_{Rx} c^2}{(4 \pi R f)^2}$$

When running a simulation it takes a finite amount of time for all of the nodes to build a full charge and they will not start to share power until they have reached 85% of full capacity. Nodes will stop sharing their power when the local reserve drops below 75% of full capacity. Nodes will only request power from their neighbours when their own local power reserve is less than 75%, this avoids the situation of nodes passing bursts of power back and forth. Only the six immediate neighbours can request power, in a 3-D array that would be twelve neighbours.

The model assumes that CPU load will follow a normal distribution centred at 50%, the CPU only becomes active once local power reaches 85% of full capacity and it shuts down when local power drops below 25%. The power drain is proportional to the instantaneous CPU load, plus a fixed rate due to support circuitry and static power losses. To improve stability, the model limits the maximum CPU activity (proportional to power reserve) when the local power reserve is less than 60% of full capacity to simulate a power aware CPU scheduling algorithm, which would prevent nodes with low power reserves accepting very intensive processes, avoiding them shutting down too often.

The nodes are represented by an array of **NodeType** objects (as defined in code sample 1 on the following page). It should be noted that although the term “Power” is used, the model actually uses the energy stored in the supercapacitor (in Joules) and as the calculations are done every 10mS, the conversion from Watts to Joules is a simple division by 100 (as Watts are a measure of Joules per Second).

Code Sample 1: Definition of Node objects (in *Visual BASIC*)

Type NodeType

Power As Single 'Energy held in node (Joules)

Availability As Boolean 'Node available to supply power

Active As Boolean 'Node available for processing

Utilisation As Single 'CPU utilisation (%)

Average As Single 'CPU averaged over time

Throttle As Boolean 'Enable CPU throttling

Supplier As String 'Nodes that supplied power (for monitoring)

End Type

The instantaneous CPU utilisation is used to calculate the power consumption while the average is used for the display and to determine when to stop supplying power to neighbours (this happens when utilisation averages 65% or over).

The supplier attribute keeps track of which nodes supply power via the optical link. This is only used for monitoring purposes, but could be used to implement a more intelligent power sharing strategy, to avoid nodes (wastefully) passing power back and forth.

Code Sample 2: Node Behaviour Rules (in *Visual BASIC*)

Rem now check if current node able to run processes and supply power

If Node(X, Y).Power < Capacity \* 0.25 Then 'Stop working at 25%

Node(X, Y).Active = False

ElseIf Node(X, Y).Power < Capacity \* 0.6 Then 'If node has less than 60%

Node(X, Y).Throttle = True 'power reserve then throttle CPU usage

ElseIf Node(X, Y).Power < Capacity \* 0.75 Then 'Stop providing power at 75%

Node(X, Y).Availability = False

Node(X, Y).Throttle = False

ElseIf Node(X, Y).Power > Capacity \* 0.85 Then 'Don't start working until 85%

Node(X, Y).Active = True

If Node(X, Y).Average < 0.65 Then 'Don't start supplying power until 85%

Node(X, Y).Availability = True 'Unless CPU has high utilisation

```

Else
    Node(X, Y).Availability = False 'In which case don't share power
End If
End If

```

The maximum energy capacity of the node (U) is calculated from the value of the supercapacitor (C) and the potential (V, which is assumed to be 5 Volts) using formula 7.2.2 :

$$U = \frac{1}{2} C V^2$$

Therefore a 0.33F supercapacitor working at 5 Volts can store 4.125 Joules of energy, allowing a 100mW node to run for about 40 seconds on a full charge.

As the nodes are updated every 10mS, the energy stored at each node (referred to as “Power” in the object model) is reduced by one hundredth of the power rating for the node. The power consumed by each node is calculated from CPU utilisation, static losses, data communications and power sharing (via laser LED). Since this ends up as heat, the total system power is calculated for each refresh cycle and added to the heat energy held by the coolant. The coolant loses heat energy at a variable rate which depends on the temperature gradient between the coolant and the surrounding air, known as Newtonian Cooling.

Additional forced cooling is automatically determined, based on the transmitter power. The model assumes the cooling system will draw the same power as the RF transmitter (which is the input power to the system) and that it is 50% efficient i.e. a 10W cooling system is able to remove 5W of heat i.e. 5 Joules per second. Forced cooling will only be employed if the coolant temperature reaches 50°C.

Using a standard formula, the total energy held in the coolant is converted into a temperature, based on the volume of coolant:

#### Code Sample 3: Calculating Coolant Temperature (in *Visual BASIC*)

```
CoolantTemp = (TotalHeat / 4000) / (Coolantvol / 1000)
```

```
'Coolant vol is in cc, specific heat capacity of water is 4000 J/Kg/C
```

The volume of coolant is calculated as 24% of the volume of the nodes in the array, since the density of water is 1 Kg/Litre, the coolant volume (calculated in cubic centimetres) is simply divided by 1000 to give the volume in litres and therefore the mass in kilogrammes.

### 7.3 Initial Results of 2-D Software Modelling

Following a number of simulation runs the parameters of the model were tuned to produce a stable array of 144 nodes (12x12). Although, at this size some nodes towards the outer edges were found to power down after around five minutes of operation. However, with 64 nodes (8x8) all nodes remained fully powered.

Figure 7.4: 8x8 Array in the Process of Charging





Figure 7.4 shows a screen-shot of an 8x8 array during the first two minutes of a simulation run as the nodes gradually accumulate power. The amount of power stored in each node's supercapacitor is represented by the colour, from deep red (0%) to bright green (100%) and the average CPU activity is represented by the number of concentric rings (0-20% per ring). It shows how the nodes immediately surrounding the transmitter antenna (the antenna is not actually shown but would be located in the middle of the 4 nodes at the centre of the array) charge most rapidly, however, as the nodes at the outside of the array are also charged by optical energy it is the four nodes located in columns 2 and 7, rows 2 and 7, which are the last to receive a full charge, in just under four minutes.

The following results are based on a run of simulations using a 12x12 array. Figure 7.5 on the following page shows the overall performance of this array, which appears to be the largest configuration that can be effectively powered via single RF transmitter and close proximity optical links. The overall power consumption rapidly falls as the inner nodes charge, from 36 Watts to a fairly steady 12-13 Watts. This is based on a maximum power consumption of 100mW per node (40mW required for an ARM A5 CPU core).

The total system performance plateaus in about three minutes with a total of around 10 GFLOPS available. There is a small variation over time as nodes are throttled (when their local power reserves fall) and as some nodes power down. However, over a period of an hour the performance remains quite stable. This is based on a peak performance of 180 MFLOPS per node, giving a theoretical maximum performance of 25.9 GFLOPS. Combining the average power consumption and the average performance, the array appears to be working at an overall efficiency of 1.3mW/MFLOPS, which is a very promising result.

Figure 7.5: Performance of a 12x12 Array



The temperature of the coolant gradually rises, stabilising at around 44°C after fifteen minutes. Additional (powered) cooling was not required as the model is able to lose the heat naturally, assuming the ambient temperature is 20°C. Since the modelling of the cooling system is greatly simplified an additional cooling mechanism may be required.

Figure 7.6 on the following page shows how a node towards the centre of the 12x12 array maintains almost 100% local power while running at full processing capacity (averaging 50% utilisation).

The nodes are numbered from the top left corner, node (1,1) as seen in figure 7.4, to the bottom right, node (8,8). The nodes which are logged to file for analysis are located along the diagonal from top left to the centre, node (6,6) in the case of the 12x12 array. They are herein referred to by a single coordinate i.e. node 1 to node 6.

Figure 7.6: Performance of an Inner Node


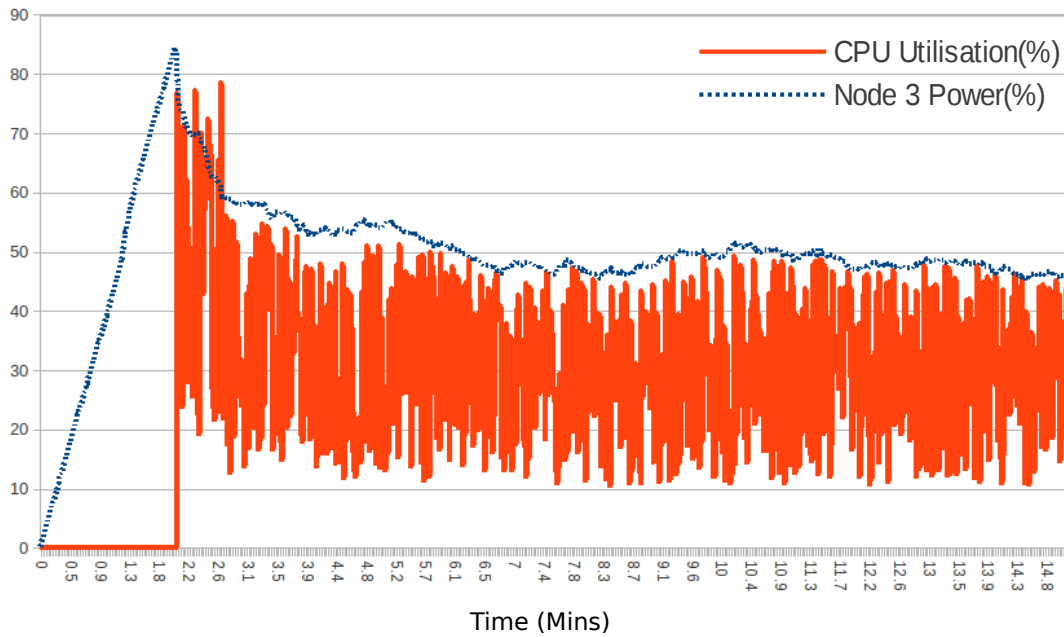


Figure 7.7 on the following page shows how a node further out (but not on the edge) of the array is still able to keep running, but with reduced CPU utilisation, averaging 30%. This demonstrates how a power aware CPU scheduling algorithm would result in lighter workloads for nodes which are receiving less power, avoiding node hibernation.

By introducing additional transmitter antennae to the model, the performance of larger arrays can be greatly improved. To avoid the problem of the transmitted waves interfering (destructively), the antennae would have to be placed at distances equal to integer multiples of the wavelength or they could be switched cyclically. One frequency band normally used for RF power transfer is 900 MHz, which has a wavelength of 333mm, but with 10mm nodes the antennae would need to be placed 40-50mm apart (a fraction of the wavelength) as nodes any further away from an antenna would not receive sufficient power.



Figure 7.7: Performance of an Outer Node



The power transfer model was further developed to use four antennae in arrays larger than 8x8 nodes. The signal to each antenna is switched at 10mS intervals so that each one delivers a burst of power. This appears to require 30% less power overall to achieve the same level of performance as a single antenna in a 12x12 array and allows arrays as large as 20x20 nodes to power up all of the nodes.

The level of average power reserve and CPU utilisation over a thirty minute period for a single antenna simulation and a multi-antenna simulation is shown in the charts on the following pages (figure 7.8 and 7.9) to demonstrate the distribution of power across the array from the edge to the centre.

Figure 7.8 shows a cross-section of nodes from node 1 (top left) to node 6, in terms of average power (stored locally) and average CPU utilisation as a percentage of full capacity. The amount of power available falls rapidly moving away from the centre of the array (and the transmitting antenna). Note that the nodes at the outside edge also receive a small amount of additional power from the ambient light via the PV cells.

Figure 7.8: Power and CPU Utilisation (%) with Single Antenna




Figure 7.9 on the following page shows the same cross-section of nodes, however there is now an antenna located between node 3 and node 4. While the available power still falls off rapidly, the nodes on the extremity of the cluster are half the distance from the transmitting antenna and so receive four times the power of the single antenna configuration due to the inverse square law nature of RF propagation i.e. power reduces proportionally to the square of the distance.

The other three antennae are located between nodes (9,3) and (10,4), between nodes (3,9) and (4,10) and between nodes (9,9) and (10,10). Potentially the inner most nodes will receive small amounts of power from all four of the antennae.

Figure 7.9: Power and CPU Utilisation (%) with Multiple Antennae





Table 7.1 on the following page shows the comparative performance of the two models. While both models set the maximum power available as 24 Watts, the multiple antennae model was able to utilise more of this power (this is displayed as average power i.e. the total power actually received by all the nodes) and delivers much better performance in the total amount of computational power available. This was due to the dramatic improvement (30% better) in the average number of nodes available for processing and the average CPU utilisation. This results in an overall improvement for system efficiency of about 5% in terms of the mW/MFLOPS efficiency measurement.

Table 7.1: Comparison of 12x12 Array with Single and Multiple Antennae

Antenna	Avg. Power	Avg. GFLOPS	mW/MFLOP	Avg. Nodes	Avg. CPU
Single	9.97W	6.84	1.46mW	70.8%	29.9%
Multiple	14.72W	10.56	1.39mW	99.4%	42.3%

The other problem with the single antenna model is that even when it is running at 24 Watts maximum power, the number of available nodes is unstable, as shown in figure 7.10 below:

Figure 7.10: Comparison of System Stability in 12x12 Arrays



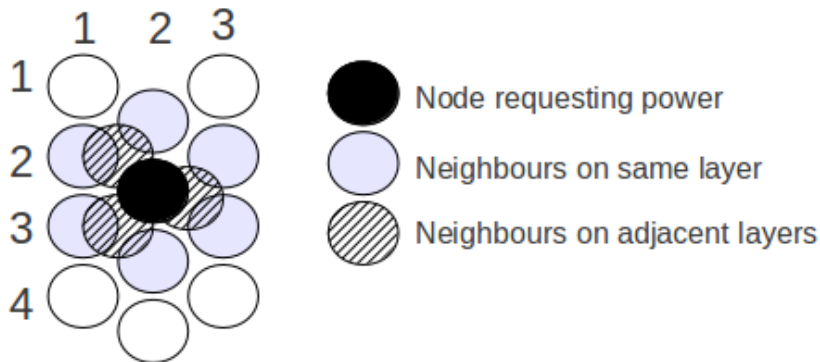
This instability can be seen in every array size, once the maximum transmitter power drops below a threshold. The effect of this instability is that the available processing power will be reduced and so the overall efficiency of the array will be reduced.

## 7.4 Software Modelling of a 3-D Array

With the 2-D model giving good results, it was further developed to simulate a full 3-D array of nodes. Each layer is modelled as a 2-D array for RF power transmission and the total system power is then simply calculated as the total of the power required for each layer.

However, the node to node power sharing now has an additional six nodes to take into account, due to close hexagonal packing. There are three immediate neighbours in the layer above and three in the layer below. In this model the nodes are referred to by three coordinates, with the layer number being the third. Layer 1 is taken as being the upper-most in the array, so node(1,1,1) becomes the top, left corner of the array cube.

Figure 7.11: Numbering Nodes in the Array



In figure 7.11 node(2,2,2) is indicated in solid black, surrounded by twelve immediate neighbours, six on the same layer and another six split between layers 1 and 3. If the node is low on local power reserve it will request a burst of power from these neighbours. Their position is calculated using the following rules, where  $(x, y, z)$  is the coordinate of the node requesting power:

If the column number ( $x$ ) is odd then the six neighbours on the same layer will be nodes  $(x \pm 1, y - 1, z)$ ,  $(x \pm 1, y, z)$ ,  $(x, y \pm 1, z)$ .

If the column number ( $x$ ) is even then the six neighbours on the same layer will be nodes  $(x \pm 1, y, z)$ ,  $(x \pm 1, y + 1, z)$ ,  $(x, y \pm 1, z)$ .

If the layer number ( $z$ ) is odd, then the six neighbours above and below will be nodes  $(x, y, z \pm 1)$ ,  $(x+1, y, z \pm 1)$ ,  $(x+1, y+1, z \pm 1)$ .

If the layer number ( $z$ ) is even then the six neighbours above and below will be nodes  $(x, y, z \pm 1)$ ,  $(x-1, y, z \pm 1)$ ,  $(x-1, y+1, z \pm 1)$ .

The simple power sharing algorithm remains the same, a node will request power from all of its neighbours (twelve in the 3-D model) if its local power reserve is less than 75%.

The neighbours will only supply a short burst of power, at 50mW (of which 25mW is actually transferred to the requesting node) if their own power reserve has already reached 85% and not dropped to less than 75%. The simulator allows individual nodes to be monitored, which reports local power reserve, CPU utilisation and the coordinates of all nodes supplying power over a one second interval.

In order to investigate the effects of node to node power transfer (via optical power transmission) the latest version of the model includes an option to enable/disable this feature. Disabling also prevents nodes at the periphery from harvesting power via ambient light. While this will result in less power available, the savings achieved by not passing power node to node may prove more beneficial overall. The big impact, however, is likely to be that the nodes on the periphery suffer from low power levels.

In order to provide a more realistic power management / CPU throttling model, the CPU limiting has been modified so that the maximum CPU workload is now based on the node charge-up time. If a node takes less than half a minute to get an initial full charge then the CPU workload will not be limited. If it takes between a half and three minutes, the CPU will be limited to between 90% and 30% of full workload otherwise it will be limited to 25%. This is based on observation of the original model, where the nodes which charge within the first minute tend to run at up to 100% utilisation (and don't shut down) and is backed up by calculation.

The calculation at the end of Chapter 4 shows that a full charge of the supercapacitor should be able to power the node for over half a minute running at 100% capacity, which is approximately the same amount of time as it takes to deliver a full charge via RF only to the nodes surrounding a transmitting antenna.

The final significant improvement to this model was to implement a model for the cooling system based on the findings in Chapter 5. The simple formula for calculating the effectiveness of a standard domestic radiator is used to calculate the amount of heat that can be dissipated (based on a given area for the heat exchanger) and the power used by a small circulating pump is calculated from the volume of coolant surrounding the nodes.

An estimated size for the heat exchanger is calculated based on the maximum power of the transmitter and a proportion of the ideal heat output of a “type 21” radiator. The power consumption of the coolant pump is now included in the calculation. Since a 35 Watt pump is quoted as circulating 36 litres per minute, the model adds 2 Watts to the total system power consumption for each litre of coolant used to allow for a change of coolant up to twice a minute (as noted before the the resistance to flow caused by the closely packed spheres is ignored). Only when the coolant temperature rises above the ambient temperature of 19°C will any heat loss actually take place. If the coolant temperature exceeds 75°C the transmitter power is reduced in steps of 1% until the temperature is reduced. This allows a safety margin of 10°C below the safe maximum working temperature of a typical processor.

It is anticipated that these improvements will yield results which are not simply a reflection of the single layer model, with factors just scaled up by the number of layers. However, some other changes have been made to reflect further research that has been carried out since the first (2-D) model was built. It has been established that the ARM FPU does not give the 180 MFLOPS originally estimated (via a simple conversion from the DMIPS score) and the true figure is actually around half that. The other factor that has been changed is the power consumption of the nodes, only 30% of the power consumption is now determined by the CPU utilisation.

The assumption is 100mW per node, with 40mW for the processor core of which 10mW is assumed to be static power. Since the original model demonstrated an average 50% efficiency in RF power transfer, the displayed result for power consumption now shows as double the total power received by all nodes, plus the power to drive the pump. This figure is now used for calculating overall system efficiency (mW/MFLOPS).

In the original model, the power transmission calculations were carried out one hundred times per second i.e. every 10mS, however, the 3-D model couldn't perform in real-time (running on a good specification laptop under Windows XP) with such a rapid refresh cycle, so it has been reduced to just ten times per second with the relevant calculations adjusted to match. The calculation of CPU utilisation was only done on a second by second basis and this remains unchanged.

## **7.5 Initial Results of 3-D Software Modelling**


The model of the cooling system delivers the level of performance expected, in a 12x12x12 array consuming approximately 300 Watts of power, it requires a heat exchanger measuring 870 x 870mm with the circulating pump drawing 29 Watts.

Over the course of an hour of simulation, the coolant temperature stabilises around 60°C. This makes sense as the system power will remain fairly constant and the cooling system will find a point of stability where the output heat is equal to the power input, since the effective heat output is proportional to the difference between the coolant temperature and the ambient air temperature.

The chart in figure 7.12 on the following page shows the temperature of the coolant for four different sized heat exchangers used with a 12x12x12 array running for 100 minutes. Using a heat exchanger 500 x 500mm resulted in the coolant overheating in less than 15 minutes, so the sizes simulated were 750 x 750mm, 870 x 870mm, 1000 x 1000mm and 1250 x 1250mm.



Figure 7.12: Temperature Rise of Coolant



From these results it has been possible to derive a simple formula for predicting the stable running temperature ( $T_s$ ) based on the power drawn by the array and the size of the heat exchanger ( $L \times L$  where  $L$  is in metres) used:

Formula 7.5.1: 
$$T_s \approx \frac{\text{Power Drawn}}{8.65 \times L^{1.54}} + T_{\text{Ambient}}$$

The next series of simulations compares the performance of an 8x8x8 array with different maximum power levels ( $T_x$ ) for the transmitter, from 20 Watts per layer to 40 Watts per layer. After an hour of activity the results were averaged out and summarised in table 7.2 on the following page.

The average power is calculated from the RF power received by each node (doubled) with the coolant pump power added to this. The average computational power of the entire array is expressed in terms of GFLOPS, while the efficiency is measured in mW/MFLOPS (which includes the pump power).

The last two columns show the average percentage of active i.e. powered up nodes and the average CPU utilisation.

The first three rows are simulations using the node to node power sharing via an optical link, while the last row\* is a simulation with the optical power sharing (and harvesting) disabled.

Table 7.2: Comparison of 8x8 Array With/Without\* Optical Power Sharing

<b>Tx Power</b>	<b>Avg. Power</b>	<b>Avg. GFLOPS</b>	<b>mW/MFLOP</b>	<b>%Active</b>	<b>Avg. CPU</b>
<b>160W</b>	83.2W	16.4	5.1	98.8%	37.6%
<b>240W</b>	93.8W	20.1	4.7	99.7%	42.2%
<b>320W</b>	99.7W	21.6	4.6	100.0%	46.2%
<b>320W*</b>	100.9W	20.3	5.0	99.6%	41.4%

With the 8x8x8 array (512 nodes) running at 160W total transmitter power, some nodes never manage to power up, but with 320W the entire array remains powered up for the full hour, it also delivers a better performance in terms of efficiency, i.e. only 4.6mW/MFLOPS as opposed to 5.1mW/MFLOPS.

Even though there are only six more nodes on average running in the third simulation, the overall CPU utilisation is almost 10% greater. This will be due to the throttling of under-powered nodes where the available power is less. The theoretical average for CPU utilisation is 50% so 46.2% utilisation is a good result.

Unsurprisingly, when the optical power transfer is disabled the performance drops noticeably, as more power is drawn from the RF transmitter, to compensate for the loss of power harvested from ambient light by the nodes on the periphery. Also, the average CPU utilisation falls, as does the proportion of active nodes.

The next series of simulations compares the performance of a 12x12x12 array with maximum transmitter power levels ranging from 30 Watts per layer to 50 Watts per layer. The results are summarised in table 7.3 on the following page.

Table 7.3: Comparison of 12x12 Array With/Without\* Optical Power Sharing

<b>Tx Power</b>	<b>Avg. Power</b>	<b>Avg. GFLOPS</b>	<b>mW/MFLOP</b>	<b>%Active</b>	<b>Avg. CPU</b>
<b>360W</b>	284.9W	54.9	5.2	98.4%	36.0%
<b>480W</b>	306.3W	67.1	4.6	99.7%	42.0%
<b>600W</b>	315.1W	71.3	4.4	100.0%	44.7%
<b>600W*</b>	312.1W	66.6	4.7	99.6%	42.2%

With the 12x12x12 array running at 360W total transmitter power, some nodes never manage to power up, but with 600W the entire array remains powered up for the full hour, it also delivers a better performance in terms of efficiency, i.e. only 4.4mW/MFLOPS as opposed to 5.2mW/MFLOPS when running at a lower input power. This is a slight improvement on the performance of the 8x8x8 array.

The reason for the improved efficiency of the larger array is that it uses multiple antennae, so the radius from antenna to receiving node is going to be approximately one node diameter less than the 8x8x8 array which has only a single antenna in the centre. In addition, the nodes in the middle of the array will also receive a small amount of power from all four antennae and so will charge quicker than the nodes on the outside edge.

Again, there is a drop in computational performance when the optical power sharing is disabled, although the total power consumption also drops slightly (in the 8x8x8 array slightly more power was required). The question is whether a 7% improvement in system efficiency (for 6% better CPU utilisation) justifies the extra cost and complexity of implementing the optical power transmission.

Working on the basis of a 12p per kWh tariff, a 12x12x12 array running 24x7 is only going to save approximately £34 per year in electricity if it is able to harvest solar energy and use node to node power sharing. Over a five year working life that is just £170, which means the cost of the optical hardware would need to be under 10p per node to make it a viable proposition i.e. 1728 nodes costing an extra 9p each would add approximately £150 to the cost of the system.

The final series of simulations were aimed at pushing the array to extremes and finding the optimal performance: pushing the transmitter power even higher, finding the optimal performance point, using a larger array of 20x20x20 nodes, studying the long term stability e.g. running for 6+ hours, reducing the heat exchanger size to find the minimum required for effective operation i.e. maintaining the coolant at just under 70°C while delivering enough power to keep all nodes running; and finally, running at 100% CPU utilisation.

**Extreme test 1:** running a 12x12x12 array with a maximum transmitter power of 1200 Watts produced a slight improvement in the CPU utilisation, averaging 48.9% with all nodes remaining powered. The slight increase in utilisation is most likely due to all nodes reaching full charge within half a minute so there was no limiting of CPU capacity. The average power drawn was only 383.5W, so the additional available power was simply not used. However, the overall efficiency ratio was 4.9mW/MFLOPS, which is not as good as when the transmitter power was 480W or 600W. As the heat exchanger was automatically sized at 1450 x 1450mm the system maintained a stable temperature under 40°C. This shows that the formula for calculating an approximation for the heat exchanger size is on the pessimistic side, as the coolant could be running at 60°C.

Some nodes had local power reserves of 107% at times, this is because the model delivers a fixed quantity of power every 100mS, switching between each of the four antennae. The Friis equation is used to calculate the power delivered to each node in each burst. If the node has less than a full charge it receives the full burst irrespective of how much more charge it actually needs, which can cause some nodes to over-charge. This would be easily addressed in the software model with a quick test of the node power level so the excess power is not accepted. However, this is a minor matter as it will not affect the performance of those nodes, since they won't have their CPU availability limited.

**Extreme test 2:** the optimal transmitter power for a 12x12x12 array initially appeared to be around 600 Watts, the next set of simulation runs used small changes in the transmitter power level to identify the best performance in terms of mW/MFLOPS. The settings for transmitter power ranged from 576 Watts to 864 Watts in increments of 12 Watts (or 1 Watt per layer) in order to find the most efficient configuration. From 600 Watts to 864 Watts the efficiency ratio remained steady at 4.4 or 4.5 mW/MFLOPS, but CPU utilisation increased from 45.2% to 48.6%. Above this the efficiency ratio fell off and utilisation plateaued.

At transmitter power levels below 600W (50W per layer) a noticeable proportion of nodes had to shut down due to low power reserves, so this seems to be the threshold for full node availability. Using this observation and the fact that RF power transmission follows an inverse square law, an approximation for this threshold can be described by the formula shown below:

Formula 7.5.2:  $Power T_x = NodePower \times Length^2 \times MaxRadius \times 10$

Where: NodePower = maximum power required by each node (0.1W)

Length = number of nodes along one side of the array

$$MaxRadius = \sqrt{2 \times (Length \div 2)^2}$$

Using this formula a 16x16x16 array and 20x20x20 array were simulated and the results are compared (below) with the performance of a 12x12x12 array:

Table 7.4: Comparison of Arrays for Optimal Power

Size	Tx Power	Avg. GFLOPS	mW/MFLOP	%Active	Avg. CPU
12x	612W	71.6	4.5	99.97%	40.1%
16x	1456W	167.1	4.7	99.9%	44.8%
20x	2820W	323.6	5.0	99.9%	45.4%

The larger arrays have yielded a better performance in terms of CPU utilisation but the transfer of power is, understandably, less efficient as more nodes are further from the transmitter antennae, so the actual efficiency ratio is worse.


**Extreme test 3:** after running a 12x12x12 simulation with a maximum transmitter power of 612 Watts (at 51W/layer) for 7 hours the following results were obtained: average power consumption of 318.7W, 99.97% of nodes active, 40.1% CPU utilisation overall, 71.6 GFLOPS average total processing power giving an efficiency ratio of 4.5mW/MFLOPS. The coolant temperature stabilised at 58.8°C after 2.4 hours, while the derived formula (7.5.1) predicted a slightly more pessimistic 61.6°C for the (automatically calculated) 910 x 910mm heat exchanger. This demonstrates the long term stability of the model.

**Extreme test 4:** after running a 12x12x12 simulation at a maximum transmitter power of 612 Watts (at 51W/layer) with a range of heat exchanger sizes, it was found that using a 780 x 780mm heat exchanger would keep the temperature at a steady 69.2°C. Reducing that by just 10mm per side (to 770 x 770mm) resulted in the system overheating within the first hour. This demonstrates that correct sizing of the heat exchanger is critical, so where users may be increasing processing capacity i.e. adding more nodes, the heat exchanger should be specified for the maximum number of nodes which the containment vessel (and power supply) can cope with.

**Extreme test 5:** finally, running a 12x12x12 simulation at a maximum transmitter power of 612 Watts (at 51W/layer) with the CPU utilisation at 100% for every node (except where it is automatically limited due to slow charging) as shown in figure 7.13 on the following page. Looking at this (top) layer, 47% of the nodes were running at 100% utilisation, while the remainder were limited, with the corner nodes (worst case) limited to between 76% and 83% processing capacity. Pushing the array to its maximum performance in this simulation yielded 141GFLOPS average total processing power, giving an efficiency ratio of 2.6mW/MFLOPS. Due to the extra CPU utilisation the heat exchanger had to be up rated to 840 x 840mm in order to keep the temperature under 70°C, which is still smaller than the automatically calculated size for the given transmitter power, so this formula seems to hold even at the extreme.

The overall proportion of active nodes was 99.8% as the nodes at the far edges did power down more frequently and the overall power consumption increased to 365.4 Watts as required to support the increase in processing workload.

Figure 7.13: 12x12 Array Running at 100% CPU Utilisation



Note that node 1 is indicated by a cross in the centre of the single circle, in the simulator this means the node that is being actively monitored and data such as CPU activity (and limit), local power reserve and which nodes, if any, supplied a burst of power via the optical link.

## 7.6 Summary

The software simulations have demonstrated that a large array of nodes powered by a combination RF and optical power transfer can, in theory, deliver a performance to match the current generation of supercomputers. The software model has been based on appropriate formulae and perhaps most importantly, in terms of the amount of power transfer that could theoretically be achieved in this scenario by wave propagation “[the] figures seem reasonable” [115] according to an expert in RF power harvesting. This suggests that the model is producing realistic results, although there are some short-comings which potentially lead to overly optimistic results.

There appears to be an optimal range for operation giving a good performance in terms of power efficiency and CPU utilisation. If transmitter power is below this range then CPU utilisation falls, although the system efficiency remains the same. Above this range the CPU utilisation plateaus and the system efficiency drops. This seems to be intuitively correct behaviour for a Ball Computer array.

In terms of the actual software model which has been developed, there are two minor short-comings which have already been identified.

The first has already been mentioned, some nodes close to the antenna can end up with more than 100% power level in the supercapacitor. This will be addressed with the addition of a simple check following the addition of a burst of charge to each node. Should the burst take the node over full capacity then the local power level will be capped (at 100%) and the unused power from that burst will be deducted from the running total of power used. This may have the effect of improving power efficiency very slightly.

The second problem is that these same nodes are probably charging faster than is physically possible. A maximum charging rate could be calculated and used to limit the power received by the nodes, as the Friis equation does not take this into account. However, the simplest approach would be to assume that the rate of charging is proportional to the rate of discharging i.e. the maximum node power consumption; so in the current model it could be limited to say 200mW per node (still allowing nodes to charge faster than they discharge).



While this should not affect overall performance dramatically, the nodes closest to the antennae will charge more slowly than before and therefore could potentially have their CPU utilisation limited. However, the threshold values could then be modified to counter this.

One other aspect that needs further investigation is the supercapacitor, the model assumes that as long as there is at least a 25% charge then the node will be able to run. However, the formula for calculating the running time of a supercapacitor [78] suggests that the discharge time is under 20 seconds rather than the 30+ seconds calculated previously.

A quick calculation would seem to suggest that the power-down threshold needs to be increased to at least 30% of full charge. This does require further work, ideally involving some experiments with real hardware to measure the actual charge/discharge time.

Following further discussion of the model, two more problems came to light. In modelling the effects of heat, the current model fails to take into account the variation in CPU power consumption as the temperature increases. If the array is to take full advantage of water cooling and run the nodes at higher temperatures i.e. 70°C plus, then this effect cannot be ignored.

Finally, it has been noted that the Friis equation is based on free space propagation of RF, however, as the Ball Computer array will be surrounded by water rather than air there will be significant attenuation of the RF. As mentioned in Chapter 3, this power will end up additional waste heat.

Further research will be undertaken to evaluate the degree of power losses, which may require a simple adjustment of the Friis equation (below) to give a more rapid attenuation of power, by increasing the value of the index in the divisor term. The difference between transmitter power and the total power received by the nodes is automatically converted to heat in the current model.

$$Power_{Rx} = \frac{Power_{Tx} Gain_{Tx} Gain_{Rx} c^2}{(4 \pi R f)^2}$$

## Chapter 8

### 8 Conclusions and Future Work

#### 8.1 Conclusions

The work covered in this thesis has shown that the Ball Computer, as proposed by Prof. Jim Austin, is indeed technically feasible, with the main factor being the continuous improvement in microprocessor design and power efficiency. This is because supplying power to “energy hungry” CPUs without power rails would prove too inefficient and they would generate a great deal more waste heat.

With the very low power consumption and high performance of widely used CPU cores, such as the ARM Cortex range, the idea of wire-free power transmission has become a reality. In just ten years the power required to deliver one MFLOPS of computational performance has reduced more than one hundred fold. To be competitive with current supercomputers a Ball Computer array would have to be delivering a performance of 3mW/MFLOPS, which the software model suggests could be achievable even with the inefficiencies of wire-free power transmission.

In terms of the mechanisms available to deliver power without the use of conventional wiring (PCB, backplane, wiring loom etc.) there are a wide range of solutions available which are used for similar applications as well as power harvesting or scavenging. Some of the more exotic solutions are even being employed by industry leaders, such as IBM (redox flow battery, combined with cooling) and Boeing (optical power and data transmission for fuel sensors).

The most obvious solution (as discussed in Chapter 3) inductive power transmission, has proven to be unworkable in the context of the Ball Computer, simply due to the scale of the individual nodes and the very poor efficiency of inductive coupling between physically small windings. However, the proposed solution of a hybrid using both RF and Optical power transmission does appear to have a great deal of potential.

The software model developed as part of this work shows that, in theory, a large 3-D array with over 1700 nodes could operate effectively using this hybrid technique for power delivery and potentially harvest additional (optical) energy from the environment.

A unit designed to house (and power) 1000 nodes could easily be built to stand on a desk, it being approximately 100mm along each dimension. The pattern of power distribution does suggest that a circular or hexagonal form would be best, avoiding that handful of nodes towards the outer corners which receive very little power. However, work on the data communications shows that the most efficient configuration for maximising the use of available bandwidth (using the zoned approach) is in fact square. This is a matter for further investigation.

For local power storage, there are some interesting technologies being developed, but the current supercapacitors offer a very viable solution for the Ball Computer. Off the shelf components could provide ample storage capacity at a physical size that could be accommodated in nodes of 10mm diameter. It has been shown that these supercapacitors could hold enough energy to run a 100mW node for over half a minute, although a custom device is going to be required which is both compact and able to handle sufficient discharge current. However, if silicon nanowire batteries do become commercially available they may offer even better power density than supercapacitors as well as having the benefit of a rapid charge/discharge cycle.

Cooling is still a major issue for all large computer installations as well as being a major consumer of power. The proposed technique of using water instead of air seems highly efficient when modelled, but of course this is backed up by work being carried out by IBM with their micro-channel architecture and the Aquasar project in Switzerland. The idea of recycling the waste heat is one which could easily be adopted for the Ball Computer and would further improve the overall system efficiency. With the combination of low power processing units and the use of water as the coolant, the Ball Computer potentially offers significant power savings in the cooling system which could more than offset any inefficiencies due to wire free power delivery.

The final question of data transmission medium is still very much open ended. Optical data transmission offers high bandwidth, good efficiency at short distances, immunity to interference from the RF power source and the possibility of combined data/power transmission between nodes. However, microwave data transmission is widely used both for large scale networks and increasingly for small scale, on-chip communications, with well defined and proven transmission protocols already in existence. At the distances involved here it could operate at very low power and also provide suitably high bandwidth.

## **8.2 Future Work**

There are eight key areas which require further investigation, development and some collaboration with specialists within the fields of optical and RF power, and data transmission in particular.

### ***8.2.1 RF Power Transmission***

Transmitter design could be based on the existing transmitters for RFID readers, adding in the ability to be able to vary transmitter power depending on the load demand. The modelling has shown that initially the power demand is high as the nodes take up their first full charge, but then this reduces to about one third as the nodes achieve capacity and start sharing power between neighbours. RFID solutions which have been investigated during this research use very low power transmitters [116] to comply with safety regulations, in Europe this is limited to 2 Watts. The proposed design for the Ball Computer is going to require transmitters at least 2 or 3 orders of magnitude more powerful. The efficiency of the transmitter will also be crucial to overall system power efficiency, but RF transmitters are notoriously inefficient.

Antenna design is critical to the efficiency of power transmission via RF/microwaves as shown in the Friis equation. While the transmitting antennae could easily be implemented as half wave dipoles, as has already been mentioned, that approach will not work for the nodes. One possible solution is a loop antenna, which is very well suited to the proposed spherical node.

A quick initial survey of information available on-line suggests that for a frequency of 900MHz, a basic loop aerial [117] would have to be 106mm in diameter, which is ten times the diameter of the proposed nodes. This is because the diameter of the loop should be the wavelength (333mm) divided by  $\pi$ . A loop 10mm in diameter would be better suited to a frequency of 9GHz. Further simulations would need to be run with this transmitter frequency, but a quick look the Friis equation [114] suggests that a ten fold increase in transmitter frequency will result in a hundred fold reduction in power received by the nodes. The alternative may be to use a multi-turn loop antenna instead, however, these introduce further complexities which go beyond the author's current understanding.

Another approach is to use similar designs to the modern generation of mobile telephones, which typically offer a gain of -7dB, as used in the software model. These are referred to as printed loops (as in printed circuit board). Again a very quick review of on-line material suggests a size of around 10mm x 25mm for a printed loop antenna working at 900MHz (including a ground plane). The gain of a small loop antenna is quoted [116] as being -20db to -15dB, considerably less than the gain of -7dB used for the simulations.

The final link in the chain is the receiver and power conversion unit. Again, the design for the Ball Computer nodes can look to conventional RFID technology here. Efficiency is the most important consideration in the design of the receiver as every small loss will be multiplied by the (potentially large) number of nodes in an array.

### ***8.2.2 Optical Power Transmission***

It is proposed that node to node power transfer could be done via optical devices. The transmitters will most likely be based on laser LEDs. The frequency of the laser will be selected to match the maximum sensitivity of the photovoltaic modules to ensure the most efficient power transfer. It has already been stated that this is unlikely to be better than 50%. The software model assumes a 50mW LED will be used, that figure is simply based on currently available, high brightness LEDs.

As has been discussed, there have been significant improvements in the design of photovoltaic devices in recent years and there are commercial systems which employ them as a method of transmitting power to small devices. Space is limited by the proposed dimensions of nodes and in order to be able to receive power from the twelve immediate neighbours will require twelve photovoltaic modules. The power output is proportional to the surface area and an optimal solution will have to be found.

The other consideration is the alignment of the nodes to ensure that all of the light from the laser LEDs hits the photovoltaic modules. There are two ways to approach this, the first is to shape the surface of nodes to ensure good alignment, either flat faces or small protrusions and matching indentations. The other approach is to use lenses to capture the light from a wider range of contact angles and focus it onto the photovoltaic modules. If the nodes are packaged in a translucent material, the lenses could be moulded into the outer casing so as to maintain the integrity e.g. against ingress of the liquid coolant. This second solution would seem to be the better option as it doesn't inherently rely on a very specific alignment of all nodes, which no doubt would prove very difficult to achieve, especially in very large arrays with hundreds or even thousands of nodes.

As has been highlighted in Chapter 7, the gains in power efficiency achieved by the use of optical (node to node) power transfer are less than 10%, with around 1% saving in power. A very rough calculation suggested that the additional hardware would have to cost no more than 9p per node to yield any kind of economic benefit in terms of power consumption.

### ***8.2.3 Data Communications***


Data communications is the least explored aspect of this thesis, largely because it is the subject of contemporary Ph.D. research work. This other work has already established a principle for zoning network communication to facilitate high speed, low contention data communication between adjacent nodes using microwaves.

The appeal of microwave communication is that it is a well established technology for high speed data links with widely used transmission protocols (IEEE 802.11 or 802.15), however, due to the nature of the Ball Computer these protocols may be unnecessarily complex, with a lot of overhead. The requirement to provide a communication channel between just three nodes, which are likely to remain fairly static means that a much simpler, light-weight protocol could be used which in turn should improve the overall data transmission rate.

The alternative solution is to use point-to-point optical communication, which should offer at least the same bandwidth as microwave and potentially far simpler protocols, as the zoning would no longer be required. While this would rely on accurate alignment of nodes, this is already a requirement of using optical power transmission between nodes, making this idea even more appealing.

As has been mentioned, the constraint at present seems to be the switching speed of optical devices. Power transfer would only require a standard, high intensity LED while data communication would require specialist, high speed devices. However, the work on the 10Giga-IR specification suggests that this has been overcome, but at the time of writing there is no information available on the final release of the new standard. This is something that will need to be followed up.

Figure 8.1: Test Rig for Optical Data Communication



Prof. Jim Austin has successfully carried out some initial experimentation using an optical data link between two ARM M4 based micro-controllers. This has demonstrated that alignment is not as critical as initially assumed and works in water. The test rig (figure 8.1) uses truncated cubes rather than spheres (for ease of handling), the dummy cubes are weighted with coins to prevent them floating.

The other consideration is communication with the outside world. Using a standard WiFi protocol could make nodes on the edge of the array accessible via standard networking devices, although the shielding required to support RF power transmission would prevent that and an antenna would have to be placed within the shielded containment vessel. If optical communication were used, provided the containment vessel is made of transparent material and the shielding is perforated, then external communication is easy enough. However, in this case custom interface devices would be required.

Clearly there are many factors and trade-offs to be considered before settling on a solution here.

#### ***8.2.4 Cooling System***

The initial calculations have indicated that water cooling using simple a heat exchanger could provide ample system cooling for a fraction of the power requirement of a forced air refrigeration unit. Further investigation and modelling is required since only simplistic calculations have been carried out in this work.

It has been assumed that the transfer of heat from the nodes to the surrounding coolant will be effective, however no consideration has been given to the effect of circulating water through a closely packed block of spheres. As the coolant will most likely be circulated by a pump, there will inevitably be some degree of turbulence and this is likely to affect the transfer of heat. Approximate calculations have been carried out in terms of the pump required, along with the simple model of temperature rise against flow rate, which indicates that a standard domestic central heating pump is more than adequate (and of course designed for constant running with water in the relevant temperature range).



The Aquasar project referred to throughout Chapter 5 uses a heat exchanger to separate the primary cooling system from the building's heating system. In order to work effectively, the primary coolant is run at a higher temperature than normal water cooled systems. The use of an intermediate heat exchanger has not been investigated within this work but is likely to be well documented by other disciplines, in particular building services design.

It would be an interesting exercise to extend the model of the simple cooling system to a full integrated system re-using the waste heat in underfloor heating. A number of assumptions and simplifications have been made for the benefit of this initial investigation, but through being able to adjust the various parameters in a full model, an optimal solution would probably be identified. The cooling system has been investigated for different power outputs from the array (and different sizes of array), but further investigation can be done in terms of different coolant temperature ranges (and temperature drops), a range of coolant flow rates (which will affect the system power consumption), different heat exchangers and a range of ambient air temperatures. It may even prove more efficient to run the primary coolant system at pressure in order to improve the rate of heat transfer, as in a conventional heat pump.

### ***8.2.5 Node Internals***

The modelling work presented here has been based around the technical specifications of an ARM Cortex-A5, which currently offers a good power consumption / computational performance ratio. The Cortex-A5 is also available as an IP core, which will allow customisation to suit the Ball Computer. As has been suggested in Chapter 2, the implementation of a soft IP core could form the basis of a student project, although the actual specification of the core / system on a chip is probably outside the scope of such a project and has not been addressed within the scope of this work either.

An alternative ARM core has also been suggested, the ARM M4 is a low power micro-controller, with a floating point unit and a number of UART devices which would be required for data communications. Like the ARM A5 it requires very low power, but is not as efficient in terms of mW/MFLOPS.

Further investigation is required to determine whether an ARM A5 based system on chip could provide sufficient I/O at low power levels to match the capabilities of an ARM M4 based micro-controller chip.

No discussion has been presented, for example, as to the merits of a CISC or RISC Instruction Set Architecture (ISA), although the CPU cores investigated have all been RISC.

It is widely acknowledged that RISC processors use less power than CISC, simply due to the reduced complexity of the hardware. However [118], this is an oversimplification of the argument since CISC can offer greater power efficiency for a given computational task. In the case of the Ball Computer, low power consumption at each node will be critical, so the basic RISC/CISC argument is probably adequate to justify the choice of a RISC processor core. It should be noted that the ARM Cortex processors are all RISC architectures.

Being self-contained processing elements, each node will include CPU, memory and I/O. Stacked memory has been referred to already, this offers a very neat solution to local memory without requiring a large and complex single die, instead the memory could be stacked on the CPU. No consideration has been given so far to the optimal amount of memory for individual nodes.

The idea of specialised CPU hardware has been mentioned in Chapter 1 and it was suggested that the Ball Computer could be made up of both general purpose and specialised nodes. The flexibility of using a custom IP core and stacked memory would allow a range of node types to be produced, after all this is a very successful solution in the biological world, where cells all have an underlying level of functionality but they all specialise for a particular function.

During initial research, Field Programmable Gate Arrays (FPGA) were investigated as a potential basis for the CPU, allowing the architecture to be modified.

However, it was soon found that the power consumption of typical FPGAs is at least two orders of magnitude greater than the low power cores being investigated. It is likely, though, that FPGAs would be used as part of the design process, for the implementation of a soft IP core.

The idea of an array consisting of a majority of general purpose nodes and a selection of specialist nodes offers some interesting possibilities, but adds a further layer of complexity to the interaction of nodes. Each node would not only need to be aware of the power reserves and availability of processor cycles of its twelve neighbours, it would potentially need to be aware of any specialised functions and be able to share this information so that appropriate tasks can be routed to the specialised nodes.

This will require a statistical analysis of typical workloads running on large parallel processor arrays to identify the proportion of tasks which would run on general purpose or specialised ISAs. Consideration would also need to be given to the overhead in routing tasks to specialised nodes rather than the nearest available node. Using zoned, node to node communication, as has been proposed for the Ball Computer, large volumes of data could end up being passed along a chain of nodes, tying up communication bandwidth for only a modest improvement in performance. Clearly there are many questions still to be answered here.

### ***8.2.6 Node Packaging***

Throughout this thesis and in the original patent [1] it has been assumed that the nodes will be spherical. As has already been highlighted, spheres allow nodes to be efficiently packed while allowing full circulation of coolant. However, spherical nodes will prove difficult to align, which could be an issue if line of sight optical links are used for node to node power and/or data transmission.

Slight modifications have been considered, such as bumps/indentations or flattened faces. While these solutions may provide a way to ensure correct alignment of nodes, they will inevitably introduce a further degree of complexity, and therefore cost, into the manufacturing process.

There has been no mention so far of alternative shapes for the node packaging, other platonic solids such as the cube, octahedron, dodecahedron have been considered during discussions, along with a truncated cube and a rhombicubocthedron [119] which would pack in the same way as cubes be with space to allow for coolant circulation.

This would need to be considered within the modelling of the cooling system and the nature of the liquid flow around smooth and angular surfaces.

Consideration has been given to packaging materials and the conclusion drawn at this point is that for prototyping, at least, a polyester epoxy resin could be used to encapsulate the nodes while the long term solution could be ceramic packaging.

This is another area where specialist knowledge is vital to the success of the project.

Working in partnership with a chip fabricator is inevitable, which would bring with it the kind of specialist knowledge required here.

### ***8.2.7 Containment Vessel***

The design of the Ball Computer containment vessel will be complex due to a number of requirements for power transmission, data communication, and cooling.

The proposed solution for power transmission is to use a 900MHz RF system and for the power levels required a Faraday cage will be mandatory. Not only will this have to protect the user from the RF energy, it will also have to be tuned to minimise losses – which would inevitably end up as excess heat. An obvious starting point for investigation is the design of Microwave ovens, which rely on safely containing high powered RF transmissions and maximising the energy transfer. Wave guides for channelling high power microwave transmissions are well understood and are to be found in applications such as RADAR. These are impedance matched for minimal losses and presumably similar approaches could be employed for the design of the containment vessel.

As has been highlighted in Chapter 5, the design will also have to allow for thermal expansion of the nodes as well as being able to support the loosely formed array in hexagonal close packing, this may require a shaped internal surface, with spherical indentations at the height of alternate layers to encourage the correct alignment. The base of the interior could also follow a similar form (a little like an egg box) to align the bottom layer of nodes correctly.

Consideration must also be given to the cooling system. Water needs to circulate freely, cold water being pumped in at the bottom and hot water drawn off from the top, to utilise the natural convection current. This also means of course that the containment vessel must be water tight and not contain any components which would corrode in water i.e. no exposed iron-based parts.

The overall shape also needs to be considered, in the initial work on data communication it was shown that the most effective shape for an entire array is actually a cube – this is based on the number of nodes on the outside edge which will have fewer neighbours to communicate with. However, a more natural shape would be hexagonal as this mirrors the shape that the array naturally forms and in terms of RF power transmission, it provides the best coverage. The software models clearly show how the level of power transfer around each antenna forms concentric hexagons, with a square array there are inevitably a few nodes which receive insufficient power. Needless to say there is a trade-off here which needs further investigation.

One other point that should not be overlooked is the aesthetic quality of the final design, one should not overlook the significance of the revolutionary design of the Apple iMac which broke the personal computer away from the established beige box look. The Ball Computer is a unique design which will look so different to conventional machines and ought not be concealed in an anonymous looking box. To enhance the aesthetics, the containment vessel should be transparent so that the nodes can be seen. This means that the Faraday cage cannot be a solid foil layer, instead it will have to be perforated, as in the door of a microwave oven for example.

In considering the overall shape, the aesthetics should also be considered – while a cube appears to offer optimal data communication performance, a hexagonal or cylindrical form may be more visually appealing.

If the vessel is transparent, then it will probably be necessary to include an opaque section, possibly in the base, in which to house the power transmission hardware, data communications interfacing and the coolant pump(s). As it is proposed to recycle the heat, any heat exchangers could be located away from the unit, connected by a length of hose.

As it is envisaged that a system can be expanded by adding more nodes, there should be easy access into the top of the containment vessel. However, there must also be a safety system which will immediately stop the RF power transmission in the event that the vessel is opened. Again, this is familiar territory for designers of microwave ovens which use two or three daisy chained micro-switches to sense the status of the door.

### **8.2.8 Software**

This is the final area to consider, although the selection and design of CPU core will be influenced by the software and typical application of a Ball Computer. Obvious applications for the Ball Computer include neural networks and real-time data processing, which are fairly standard applications for massively parallel machines.

If the nodes are to work as independent, self organising units then presumably each will host a light-weight, distributed operating system. Each node will be aware of its own power reserve and recharge rate, which will determine the amount of computing power it can offer. It will also be aware of its twelve immediate neighbours and their capabilities. With this information the nodes should distribute and organise work loads as well as routing data through the array.

As has already been mentioned, a power management protocol will be required if node to node (optical) power transfer is to be implemented. This will form part of the “power awareness” function.

In the software models developed, power is automatically shared when a node has reached the 85% threshold for local power reserves, but in reality a node that is low on power will have to negotiate with its neighbours to receive top-up power from them. This raises some questions that haven't been addressed so far with regards to power sharing. Since the power sharing protocol will use the same communication channel as the data, simply polling all twelve neighbours while ever local power is below the 75% threshold is potentially going to tie up valuable bandwidth.

Therefore nodes will also need to have a degree of awareness of their neighbours' status so power requests will only be made to neighbours with a good power reserve, which are not processing at high intensity. If several neighbours could supply power should the node request power from them all or work on a round-robin principle? (This should be answered with further software modelling.)

The other key function of the Ball Computer Operating System is data communication (and routing), which may or may not end up being based on standard protocols. For example, TCP/IP is a fully routable protocol and is used over 802.11 (Wi-Fi) for wireless LAN applications. TCP/IPv4 is set to phase out over the coming years, so that any investigation into the potential of this protocol should focus on version 6. Since much of the protocol suite would not actually be required e.g. name resolution (DNS) then a light weight version could be used.

An obvious choice for the Operating System would be a UNIX derivative, one of the popular Linux distributions perhaps. Being open source the Linux Operating System can be heavily customised and stripped down to suit the Ball Computer but at the same time will include standard elements such as a TCP/IP stack, and by its very nature works well in a multiprocessor environment.

### **8.2.9 Summary**

There is a lot more work to do in order to even get to the prototyping stage, with many key decision to be made. The overall conclusion has to be that the Ball Computer is now a feasible proposition and this work needs to be moved forward, appropriately, with a number of expert groups working in collaboration.

# Appendix

## A. CPU Statistics

Table 9.1: Comparison of CPU power/performance

Year	CPU	Power	MFLOPS	DMIPS	mW/MFLOP
2000	Pentium III 1GHz	35W			
2000	Celeron 1GHz	28W		3444	95.81
2000	Athlon 1 GHz	55W		4108	157.78
2001	Pentium 4 2 GHz	91W	459	5402	198.3
2001	Duron 1 GHz	46W			
2002	Athlon XP 2 GHz	70W		7527	109.59
2003	Athlon 64 3 GHz	70W		8325	99.09
2003	Sempron 64 2.7 GHz	59W	1190		49.59
2003	DEC Alpha 21364 1 GHz	125W			
2004	DEC Alpha 1.3 GHz	125W			
2005	Pentium D 3 GHz	154W			
2007	Core 2 Duo 3GHz	105W	2658		39.51
2007	Core 2 Quad 3 GHz	167W	7462		22.38
2007	PhenomX4 2.2 GHz	70W	7525		9.3
2008	Phenom II X3 3 GHz	73W	3907		18.69
2009	Core i5-750 3 GHz	95W	7851		12.1
2009	Core i7-970 3.2 GHz	130W			
2009	Turion II 2.7 GHz	35W			
2009	Cortex A5 1GHz	80mW		1600	0.59
2009	MIPS32 24K 1.4 GHz	150mW		2114	0.84
2010	Core i3 -540 3.06 GHz	170W			
2010	Cortex A9 830 MHz	400mW		2075	2.27
2011	Tilera Gx-100	55W			

Sources used (see references 10-13 for further details):

<http://www.cpu-world.com/sspec/index.html>

<http://www.maxxpi.net/pages/result-browser/top10---flops.php>

<http://www.intel.com/products>

<http://www.tomshardware.com/charts/cpu-charts-2004/Sandra-CPU-Dhrystone,449.html>



## B. Example of Raw Array Simulation Data

The following page is an example of raw data from just over an hour of a simulation run with an 8x8x8 array.

The column headings are:

Time(S) – Time in seconds

Power(W) – Total power consumed by the array

Temp(C) – coolant temperature

Perf(MFLOP) – total computational performance in MFLOPS

Active(%) - proportion of nodes powered up

Nn Pwr(%) - Node n power reserve as a proportion of full power

CPU (%) - Node n CPU utilisation (averaged)

This data was initially written to a CSV file which was imported into a spreadsheet for analysis. Four nodes have been sampled, along the diagonal of the cube from node (1,1,1) to node (4,4,4).

Note that for the first minute data is logged every second, after that is every ten seconds. The sample data therefore starts one minute into the simulation.

Time(S)	Power(W)	Temp(C)	Perf(MFLOS)	Active(%)	N1 Pwr(%)	CPU(%)	N2 Pwr(%)	CPU(%)	N3 Pwr(%)	CPU(%)	N4 Pwr(%)	CPU(%)
60	95	20.5	17346.76	98.4	75.1	0	84	51.7	100.1	48.8	103	55
70	91.1	20.7	20281.25	99.7	81.5	23.7	83.9	44.5	100.1	48.7	102.1	44.1
80	91.2	20.9	20660.77	99.9	82.6	27.7	83.1	39.1	100.1	49.5	102	53.9
90	94.4	21.1	20529.2	100	79.5	30.7	82.8	34.8	100.1	50.4	102	48.8
100	94.1	21.3	21115.38	100	76.2	41.4	82.5	45.5	100.1	48.4	102.2	51
110	91.1	21.5	20782.46	100	72.2	38.7	81.5	44.6	100.1	41.9	102.1	53.2
120	91.3	21.6	21040.9	100	68.4	39.1	77.4	33.5	100.1	57.8	102.1	55.1
130	94.3	21.8	20698.19	100	64.3	40.5	74.6	44.4	100.1	48.5	102.1	39.3
140	94.1	22	20744.39	100	60.7	39.1	74	49.5	100.1	55.4	102.1	54.9
150	92.9	22.2	20907.74	100	57.1	28.7	73.6	38.1	100.1	41.4	102.2	60.3
160	93.3	22.4	20993.16	100	53.9	28.4	73.8	44.6	100.1	54	102.1	56.7
170	92.3	22.6	20622.04	99.9	50.3	39.2	73.1	41	100.1	59.4	102	52.1
180	93.8	22.7	20213.07	99.7	46.5	47.4	73	39.3	100.1	43.7	102.1	42.6
190	94.1	22.9	20183.12	99.7	42.4	36.6	73.1	28.7	100.1	47	102.2	55.8
200	92.6	23.1	19337.67	99.7	38.6	39	73.1	41.5	100.1	41.6	102.1	37.9
210	92	23.3	20256	99.7	34.8	40.7	72.7	41	100.1	48.8	102.2	49.7
220	94.4	23.4	20104.41	99.7	31.1	41	72.4	43.5	100.1	48	102.1	57.5
230	93.5	23.6	20051.02	99.7	27.5	28.4	71.8	34.4	100.1	49.1	102	46.9
240	92.9	23.8	19935.87	99.5	28.4	8.2	71.6	41.4	100.1	47.7	102.2	44.2
250	95.1	24	19829.54	99.5	40.4	0.9	71.4	40.1	100.1	47.2	102.2	47.3
260	94.9	24.1	19367.43	99.2	53	0.1	71.4	38.6	100.1	42.1	102.2	39.1
270	91.4	24.3	20035.33	99.2	65.6	0	71.4	45.9	100.1	57.7	102.1	50.6
280	90.6	24.5	19742.19	99.4	78.3	0	70.7	37.6	100.1	50.4	102.1	40.3
290	94.4	24.7	19713.89	99.5	78.2	30	70.6	36.4	100.1	55.4	102.1	48
300	94.7	24.8	20727.95	99.7	71.7	35.2	70.2	44.1	100.1	52.1	102.1	46.6
310	92.1	25	19814.61	99.7	67.1	43.8	70.1	40	100.1	56.6	102.2	51
320	93.1	25.2	20140.16	99.7	63.1	39.6	70.1	44.2	100.1	47.1	102.2	47.6
330	94.1	25.3	19831.53	99.8	59.7	35.9	70.4	35.5	100	55.8	102.1	43.6
340	94.9	25.5	20142.75	99.8	56.2	28.6	70.1	53.1	100	53	102.1	43.2
350	91.6	25.7	20244.09	99.8	53.1	37	69.9	47.6	100.1	41.7	102.1	42.6
360	93.7	25.8	20318.55	99.8	49.3	33.1	70.1	41.3	100.1	58.6	102.1	38.6
370	93.9	26	19848.69	99.8	46.1	33.4	70	44.1	100.1	53.6	102.1	57.1
380	93.5	26.2	20626.47	99.8	42.7	37	70	45.8	100.1	51	102.1	48
390	94.2	26.3	20193.17	99.8	39.1	35.7	70	40.1	100.1	46.1	102.1	47.5
400	93.3	26.5	19863.39	99.8	35.6	32.2	70.1	46.8	100.1	47.6	102.1	44
410	91.5	26.7	19927.86	99.8	32.6	34	70.1	31.5	100.1	49.5	102.1	41.3
420	92.8	26.8	20904.33	99.8	29.4	36.7	70.1	38.6	100.1	55.8	102.1	39
430	94.8	27	20537.67	99.7	28.3	32.6	70	47.2	100.1	48.5	102.1	51
440	95.2	27.1	19181.65	99.6	26.5	12.7	70.1	33.3	100.1	38.7	102.2	48.3
450	94.3	27.3	19372.49	99.5	36.4	1.4	70	49.2	100.1	55.5	102.1	55.3
460	92.1	27.5	20242.77	99.4	49	0.1	70	39.9	100.1	44	102.1	58.4
470	94.1	27.6	20061.22	99.4	61.6	0	69.9	42.4	100.1	52.2	102	54.6
480	92.9	27.8	19711.46	99.5	74.2	0	69.9	48.4	100.1	51.9	102.2	44.5
490	94.4	27.9	20158.02	99.5	80.2	26.5	70	39.3	100.1	47.4	102.2	57.8
500	94.1	28.1	20051.61	99.6	72.9	38.8	69.9	48.8	100.1	56.1	102.1	52.4
510	92.3	28.2	19477.7	99.6	69.7	44.6	70	48.4	100.1	48.8	102.1	45.3
520	91.8	28.4	19556.17	99.6	62.6	38.5	69.9	51.7	100	54	102.1	46.1
530	91.4	28.5	20407.29	99.6	53	34.3	70.1	36	100.1	48.2	102	55.2
540	94.5	28.7	20813.87	99.7	49.1	39.5	70	35.6	100.1	47.8	102.1	37.2
550	94.3	28.8	19539.16	99.7	45.5	39.7	70	46	100.1	43.9	102.2	48.8
560	94.9	29	20482.87	99.8	41.9	44.5	70	37.4	100.1	49.4	102.1	53.7
570	93	29.1	20360.84	99.8	40.5	43	69.9	42.2	100.1	47	102	55.8
580	94.4	29.3	20741.37	99.8	37.4	27.1	70	39.9	100.1	49	102	46.9
590	95.8	29.4	19572.28	99.8	34.4	32	70	42.6	100.1	52.2	102.2	61.6
600	93.4	29.6	20194.97	99.8	30.9	35.9	70.1	30.8	100.1	41.3	102.1	50.2
610	91.9	29.7	20174.38	99.8	27.4	39.4	70.2	43.9	100.1	49.9	102.1	52.8
620	93.9	29.9	20408.09	99.7	29.1	6.5	69.9	45.6	100.1	48.6	102.2	51.3
630	93.5	30	19651.16	99.6	41.4	0.7	70	42.3	100.1	54.7	102.1	46.4
640	91.6	30.2	20537.36	99.5	54	0.1	70.1	38.6	100.1	56.9	102.1	46.8
650	94.6	30.3	19745.64	99.5	66.7	0	70.2	38	100.1	42.1	102	46.8

## References

- [1] J. Austin (Cybula Ltd), "Computing Devices" International Patent Application No. PCT/GB02/04104, September 10, 2002
- [2] University of York, Dept. of Computer Science. (2013) *Advanced Computer Architectures Group* [Online] Available: [http://www.cs.york.ac.uk/arch/computer-architectures/index\\_html](http://www.cs.york.ac.uk/arch/computer-architectures/index_html)
- [3] H. Gietl and H. Meuer. (2009, May, 20) *The Top Trends in High Performance Computing* [Online] Available: [http://www.top500.org/blog/2009/05/20/top\\_trends\\_high\\_performance\\_computing](http://www.top500.org/blog/2009/05/20/top_trends_high_performance_computing)
- [4] R.M. Hord, *Understanding Parallel Supercomputing*. IEEE: New York, 1999
- [5] nVIDIA. (2013) *CUDA* [Online] Available: [http://www.nvidia.com/object/cuda\\_home\\_new.html](http://www.nvidia.com/object/cuda_home_new.html)
- [6] Scientific Computing. (2012, Mar. 7) *UK's Emerald Supercomputer Enters Service* [Online] Available: <http://www.scientificcomputing.com/news-HPC-UK-Emerald-Supercomputer-Enters-Service-070312.aspx>
- [7] nVIDIA. (2011, Jun.) *TESLA M2090 DUAL-SLOT COMPUTING PROCESSOR MODULE BD-05766-001\_v02 | June 2011 Board Specification* [Online] Available: <http://www.nvidia.com/docs/io/43395/tesla-m2090-board-specification.pdf>
- [8] Scientific Computing. (2012, Jun. 20) *Japan Reclaims Top Ranking on Latest TOP500 List of World's Supercomputers* [Online] Available: <http://www.scientificcomputing.com/news-HPC-Japan-Reclaims-Top-Ranking-on-Latest-TOP500-List-of-Worlds-Supercomputers-062011.aspx>
- [9] Top500. (2012, Jun.) *June 2012* [Online] Available: <http://www.top500.org/lists/2012/06/>
- [10] MaxxPI2. (2012, Jul.) *Top 10 – FLOPS* [Online] Available: <http://www.maxxpi.net/pages/result-browser/top10---flops.php>
- [11] ARM. (2012, Jul.) *Processors* [Online] Available: <http://www.arm.com/products/processors>
- [12] CPU World. (2010) *Specifications numbers of Intel processors* [Online] Available: <http://www.cpu-world.com/sspec/index.html>
- [13] Tom's Hardware. (2004) *Sandra – CPU Dhrystone* [Online] Available: <http://www.tomshardware.com/charts/cpu-charts-2004/Sandra-CPU-Dhrystone,449.html>
- [14] E. Grochowski and M. Annavaram. *Energy per Instruction Trends in Intel® Microprocessors* [Online] Available: <http://support.intel.co.jp/pressroom/kits/core2duo/pdf/epi-trends-final2.pdf>
- [15] Energy.Gov. (2012, Jul. 1) *Room Air Conditioners* [Online] Available: <http://energy.gov/energysaver/articles/room-air-conditioners>
- [16] Engineering Toolbox. *Air Conditioner Efficiency* [Online] Available: [http://www.engineeringtoolbox.com/air-conditioner-efficiency-d\\_442.html](http://www.engineeringtoolbox.com/air-conditioner-efficiency-d_442.html)
- [17] S. Kamil et al, "Power Efficiency in High Performance Computing" presented at the High-Performance, Power-Aware Computing, Miami, Fl, April 14, 2008

- [18] Power Electronics Technologies. (2008, May, 13) *Switch-Mode Power Supplies for Beginners: An Efficiency Primer Part 1* [Online] Available: [http://powerelectronics.com/spotlight/power\\_primer/switch-mode-power-supplies-part1-pp0513/](http://powerelectronics.com/spotlight/power_primer/switch-mode-power-supplies-part1-pp0513/)
- [19] D. Lautzenheiser. (2008, Feb.) *Good Embedded Communications is the Key to Multicore Hardware Design Success* [Online] Available: [http://www.low-powerdesign.com/article\\_silistix\\_lautzenheiser.htm](http://www.low-powerdesign.com/article_silistix_lautzenheiser.htm)
- [20] E. Salman and Qi Qi, "Path Specific Register Design to Reduce Standby Power Consumption" *Journal of Low Power Electronics and Applications*, Vol. 1, pp. 131-149, 2011
- [21] H. Iwai. (2009, Mar. 25) *Roadmap for 22 nm and beyond* [Online] Available: [www.iwailab.ep.titech.ac.jp/pdf/iwaironbun/0906infos.pdf](http://www.iwailab.ep.titech.ac.jp/pdf/iwaironbun/0906infos.pdf)
- [22] M. Dubois et al, *Parallel Computing Organization and Design*. Cambridge University Press: Cambridge, 2012
- [23] H. Esmaeilzadeh et al, "Dark Silicon and the end of Multi-core Scaling" Presented at the 38th International Symposium on Computer Architecture, San Jose, CA, June 4-8, 2011
- [24] 7-Zip. *7-Zip LZMA Benchmark* [Online] Available: <http://www.7-cpu.com/>
- [25] Fierce PC. (2012) *Intel Quad Core i7 2600K 3.4GHz Socket LGA1155 CPU Processor* [Online] Available: <http://www.fiercepc.co.uk/cpus/processors/intel-quad-core-i7-2600k-3.4ghz-socket-lga1155-cpu-processor/>
- [26] Freescale. (2012) *i.MX515: Applications Processor* [Online] Available: [http://www.freescale.com/webapp/sps/site/prod\\_summary.jsp?code=i.MX515](http://www.freescale.com/webapp/sps/site/prod_summary.jsp?code=i.MX515)
- [27] S. K. Moore. (2008, Nov.) *Multicore Is Bad News For Supercomputers Adding cores slows data-intensive applications* [Online] Available: <http://spectrum.ieee.org/computing/hardware/multicore-is-bad-news-for-supercomputers>
- [28] K. Bergman et al, "Let There Be Light! The Future of Memory Systems is Photonics and 3D Stacking" Sandia National Laboratories, New Mexico, June 5, 2011
- [29] D. Flynn. (2007, Sep. 18) *AMD announces three-core desktop CPU* [Online] Available: [http://apcmag.com/amd\\_announces\\_threecore\\_desktop\\_cpu.htm](http://apcmag.com/amd_announces_threecore_desktop_cpu.htm)
- [30] Wikipedia. (2012, May, 10) *Multi-core processor* [Online] Available: [http://en.wikipedia.org/wiki/Multi-core\\_processor](http://en.wikipedia.org/wiki/Multi-core_processor)
- [31] Scientific Computing. (2012, Jul. 18) *Achieving One Million Times more Computing Efficiency* [Online] Available: <http://www.scientificcomputing.com/news-HPC-Achieving-One-Million-Times-More-Computing-Efficiency-071812.aspx>
- [32] S. Patil et al. *Spintronic Logic Gates for Spintronic Data Using Magnetic Tunnel Junctions* [Online] Available: <http://www.ece.umn.edu/users/pati0036/spingatespindata.pdfShare>
- [33] ARM. (2010, Jul. 28) *Dhrystone and MIPS performance of ARM processors* [Online] Available: <http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.faqs/ka3885.html>

- [34] Tiler. (2012) *The TILE-Gx processor family* [Online] Available: <http://www.tiler.com/products/processors/TILE-Gx-3000>
- [35] ARM. (2010) *Cortex-A5 Technical Reference Manual* (Revision: r0p1) [Online] Available: <http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.ddi0433b/index.html>
- [36] MIPS Technologies. (2002, Jun.) *Choosing an Intellectual Property Core* [Online] Available: <http://www.mips.com/products/product-materials/whitepapers/>
- [37] MIPS Technologies. (2012) *Open Virtual Platforms Simulator Models for MIPS Family Cores* [Online] Available: <http://www.mips.com/products/system-software/simulator/>
- [38] ARM. (2013) *ARM Artisan Physical IP Solutions* [Online] Available: <http://www.arm.com/products/physical-ip/index.php>
- [39] Verilog Dot Com. (2012) *Verilog Resources* [Online] Available: <http://www.verilog.com/>
- [40] A. Milenkovic and D. Fatzer, "Teaching IP Core Development: An Example" Proceedings of IEEE International Conference on Microelectronic Systems Education, Austin, Texas, USA, 1-2 June 2003
- [41] E. Waffenschmidt. *Proximity Power is Efficient and Safe* [Online] Available: <http://www.wirelesspowerconsortium.com/technology/>
- [42] X. Liu and S. Y. Hui, "Optimal Design of a Hybrid Winding Structure for Planar Contactless Battery Charging Platform" IEEE Transactions on Power Electronics, Issue 1, January 2008
- [43] J. Gao, "Inductive Power Transmission for Untethered Micro-Robots", IEEE, Page 2011, 2005
- [44] Texas Instruments. (2011, Aug.) *Qi Compliant Wireless Power Transmitter Manager (bq500210)* [Online] Available: <http://www.ti.com/lit/ug/slvu467/slvu467.pdf>
- [45] Texas Instruments. (2011, Aug.) *Wireless Power Transmitter Manager EVM (Users Guide)* [Online] Available: [www.ti.com/lit/ds/slusao2/slusao2.pdf](http://www.ti.com/lit/ds/slusao2/slusao2.pdf)
- [46] Texas Instruments. (2011, Aug.) *Integrated Wireless Power Supply Receiver (bq51010, bq51011, bq51013)* [Online] Available: <http://www.ti.com/general/docs/lit/getliterature.tsp?genericPartNumber=bq51011&fileType=pdf>
- [47] D. Wageningen and E. Waffenschmidt. *Transfer Efficiency* [Online] Available: <http://www.wirelesspowerconsortium.com/technology/transfer-efficiency.html>
- [48] J. D. Kraus, *Electromagnetism*. 5th Edition. Boston : WCB/McGraw-Hill, 1999.
- [49] W. J. Duffin, *Electricity and Magnetism*. 4th Edition. London : McGraw-Hill, 1990.
- [50] Murata. *Capacitive coupling wireless power transmission technology* [Online] Available: [http://www.murata.com/products/wireless\\_power/tech\\_intro/index.html#A002](http://www.murata.com/products/wireless_power/tech_intro/index.html#A002)
- [51] P. Camurati and H. Bondar, "Device for transporting energy by partial influence through a dielectric medium" US Patent Application No. US 2009/0206675 A1, August 20, 2009

- [52] RFID Journal. (2005, Jan. 16) *The basics of RFID Technology* [Online] Available: <http://www.rfidjournal.com/article/articleview/1337/1/129/>
- [53] B. Jiang et al, "Energy Scavenging for Inductively Coupled Passive RFID Systems", presented at Instrumentation and Measurement Technology Conference , Ottawa, 17-19 May, 2005
- [54] B. Allen. (2011, Sep.) *University of Bedfordshire Knowledge Network presentation slides – Green Batteries* [Online] Available: <http://www.beds.ac.uk/howtoapply/departments/computing/staff/benallen>
- [55] Powercast. (2012) *P1110 Datasheet* [Online] Available: <http://www.powercastco.com/products/powerharvester-receivers/>
- [56] T. Suzuki et al, "MOS-FET DC-to-RF Inverter for Power Transmission via Insulation layer", IEEE, Page 1479, 1995
- [57] R. F. Cleveland, Jr. et al, "Evaluating Compliance with FCC Guidelines for Human Exposure to Radiofrequency Electromagnetic Fields." Appendix A, OET Bulletin 65 , Edition 97-01 , August 1997
- [58] H. Lehpame, *RFID Design Principles*. 2<sup>nd</sup> Edition, London: Artech House, 2012.
- [59] R. Rosner, *MacMillan Encyclopedia of Physics*, vol. 4, London: Macmillan Library Reference, 1996. Cited by The Green Tank (2009, Dec.) [Online] Available: <http://thegreentank.blogspot.com/2009/12/solar-heat-number.html>
- [60] J. M. Crow, "Sun Strokes", *New scientist*, 26 November 2011, pp. 38-40
- [61] Daily Tech. (2012, Jul. 12) *Lockheed Martin Stalker UAV Powered by Laser Light for 48 Hours* [Online] Available: <http://www.dailytech.com/article.aspx?newsid=25156>
- [62] W. Ferguson and D. Hambling, "A small buzz in the air", *New scientist*, 21 July 2012, pp. 19.
- [63] A. Basanskaya, "Electricity Over Glass", *IEEE Spectrum*, October 2005, pp.18
- [64] RP Photonics *RP Photonics Encyclopedia: Laser Diodes* [Online] Available: [http://www.rp-photonics.com/laser\\_diodes.html](http://www.rp-photonics.com/laser_diodes.html)
- [65] Fibre Optic FX Ltd. *Side Emitting Cable* [Online] Available: <http://www.fibreopticfx.co.uk/Side-Emitting-Cable.html>
- [66] D. Worthington (2012, Jul. 31) *Scientists invent paper LEDs* [Online] Available: <http://www.smartplanet.com/blog/intelligent-energy/scientists-invent-paper-leds/18222>
- [67] D. N. Sinha. (2005, Aug. 12) *Power Generation in Pipeline* [Online] Available: <http://preview.tinyurl.com/7y2p78q>
- [68] P. R. Tripathi and A. Khan "Wireless Electricity Transmission (using acoustic resonance through piezoelectric crystals)" presented at International Conference on Advances in Electrical and Electronics Engineering (ICAEE'2011), Pattaya, Thailand Oct 7-8, 2011
- [69] Sunon. (2010) *Mighty Mini Fan & Blower* [Online] Available: [http://www.sunon.com/pro2\\_page.php?pkid=4](http://www.sunon.com/pro2_page.php?pkid=4)
- [70] D. Graham-Rowe, "Liquid Power for Chips", *New Scientist*, 19 November 2011, pp. 25

- [71] M. Skyllas-Kazacos et al, “The vanadium redox battery - an energy reservoir for stand-alone ITS applications along motor and expressways “ Intelligent Transportation Systems, 2005. Proceedings. 2005 IEEE Issue Date: 13-15 Sept. 2005, pp. 391 – 395
- [72] Z. Xu et al. *Self-Charged Graphene Battery Harvests Electricity from Thermal Energy of the Environment* [Online] Available: <http://arxiv.org/pdf/1203.0161v1.pdf>
- [73] British Computer Society, “Green Impact of Hardware (Batteries of the Future)”, IT Now, May 2011, pp. 15-17
- [74] C.K. Chan et al, “High-performance Lithium Battery Anodes Using Silicon Nanowires”, *Nature Nanotechnology*, vol. 3, January 2008, pp. 31-35 [Online] Available: <http://www.nanoarchive.org/417/1/opr001OE.pdf>
- [75] IBM. *Battery 500 Project: Lithium-air battery* [Online] Available: [http://www.ibm.com/smarterplanet/us/en/smart\\_grid/article/battery500.html](http://www.ibm.com/smarterplanet/us/en/smart_grid/article/battery500.html)
- [76] Battery University. (2013) *Lithium-based Batteries* [Online] Available: [http://batteryuniversity.com/learn/article/lithium\\_based\\_batteries](http://batteryuniversity.com/learn/article/lithium_based_batteries)
- [77] J. Miller and P. Simon, “Supercapacitors: Fundamentals of Electrochemical Capacitor Design and Operation”, *The Electrochemical Society Interface*, Spring 2008.
- [78] Farnell. (2012) *CORNELL DUBILIER - EDLSG105V5R5C - CAPACITOR, 1F, 5.5V, STAKED COIN* [Online] Available: <http://uk.farnell.com/cornell-dubilier/edlsg105v5r5c/capacitor-1f-5-5v-staked-coin/dp/2113011>
- [79] S. Zurek. *Supercapacitors Chart* [Online] Available: [http://en.wikipedia.org/wiki/File: Supercapacitors\\_chart.svg](http://en.wikipedia.org/wiki/File:Supercapacitors_chart.svg)
- [80] M. Pant, “Intel. Microprocessor Power Impacts.” SIGDA–NSF–SRC–DAC Design Automation Summer School, San Diego, CA, June 5–6, 2011
- [81] W. L. Bircher and L. K. John, “Analysis of Dynamic Power Management on Multi-Core Processors”, Presented at ISCA’09, Austin, Texas, USA, June 20–24, 2009 [Online] Available: [http://www.eecs.harvard.edu/~dbrooks/isca09\\_rangan.pdf](http://www.eecs.harvard.edu/~dbrooks/isca09_rangan.pdf)
- [82] Intel. *Intel® Centrino® Duo Mobile Technology* [Online] Available: [http://www.intel.com/technology/itj/2006/volume10issue02/art03\\_Power\\_and\\_Thermal\\_Management/p01\\_abstract.htm](http://www.intel.com/technology/itj/2006/volume10issue02/art03_Power_and_Thermal_Management/p01_abstract.htm)
- [83] AMD. (2013) *Cool n' Quiet and CoolCore technology* [Online] Available: <http://www.amd.com/us/products/technologies/cool-n-quiet/Pages/cool-n-quiet.aspx>
- [84] ARM. (2010) *Cortex -A5 MP Core Technical Reference Manual* [Online] Available: <http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.ddi0434b/CHDIGBJG.html>
- [85] Green Droid. (2010) *Using Dark Silicon to Improve Smartphone Processors* [Online] Available: <http://greendroid.ucsd.edu/>
- [86] Technology Review. (2011, Apr. 28) *App-Specific Processors to Fight Dark Silicon* [Online] Available: <http://www.technologyreview.com/computing/37478/?a=f>
- [87] UCSD. (2011) *The Greendroid Mobile Application Processor: An Architecture for Silicon’s Dark Future* [Online] Available: <http://greenroid.ucsd.edu>

- [88] STATS ChipPAC. (2013, Mar.) *Data Sheet for Fine Pitch Ball Grid Array* [Online] Available:  
<http://www.statschippac.com/services/packagingservices/~~/media/Files/Package%20Datasheets/FBGA.ashx>
- [89] Comair Rotron. (2012) *Airflow Calculation* [Online] Available:  
[http://www.comairrotron.com/airflow\\_note.shtml](http://www.comairrotron.com/airflow_note.shtml)
- [90] Plumbnation. *GRUNDFOS Selectric/Super selectric* (Data Sheet) [Online] Available: <http://www.plumbnation.co.uk/site/grundfos-ups-15-50-selectric-130-bare-pump/grundfos-ups-selectric-pumps-brochure.pdf>
- [91] E. G. Colgan et al, "A Practical Implementation of Silicon Microchannel Coolers for High Power Chips" presented at the 21st Annual IEEE Semiconductor Thermal Measurement and Management Symposium. Piscataway, NJ, 15-17 March, 2005
- [92] Scientific Computing. (2009, Jun. 7) *Water-cooling System Enables Supercomputers to Heat Buildings* [Online] Available:  
<http://www.scientificcomputing.com/news-hpc-Water-cooling-system-enables-supercomputers-to-heat-buildings-070609.aspx>
- [93] Scientific Computing. (2012, Jul. 17) *First-of-a-kind Hot water-Cooled Supercomputer Goes Live* [Online] Available:  
<http://www.scientificcomputing.com/news-HPC-First-of-a-kind-Hot-Water-Cooled-Supercomputer-Goes-Live-071712.aspx>
- [94] M. Ashby et al, *Materials Engineering, Science, Processing and Design*. Butterworth-Heinemann: Oxford, 2007.
- [95] D. A. Bartlett, "The Fundamentals of Heat Exchangers ", *The Industrial Physicist*, Vol. 2, Issue 4, 1996, pp. 18-21 [Online] Available:  
[www.aip.org/tip/INPHFA/vol-2/iss-4/p18.pdf](http://www.aip.org/tip/INPHFA/vol-2/iss-4/p18.pdf)
- [96] Engineering Toolbox. *Overall Heat Transfer Coefficients for some common Fluids and Heat Exchanger Surfaces* [Online] Available:  
[http://www.engineeringtoolbox.com/overall-heat-transfer-coefficients-d\\_284.html](http://www.engineeringtoolbox.com/overall-heat-transfer-coefficients-d_284.html)
- [97] F. Hall, *Building Services and Equipment*. vol. 3, Second Edition. Longman: London, 1987.
- [98] Heat and Gas dot Com. (2008) *Technical Specification* [Online] Available:  
[http://www.heatlineradiators.co.uk/menu\\_radiators\\_spec.htm](http://www.heatlineradiators.co.uk/menu_radiators_spec.htm)
- [99] E. W. Weisstein. *Hexagonal Close Packing*, From MathWorld--A Wolfram Web Resource [Online] Available:  
<http://mathworld.wolfram.com/HexagonalClosePacking.html>
- [100] C. Bailey et al. "Evaluating 3D Wireless Interconnect Schemes towards A Wire-Free Processing Array", currently unpublished
- [101] S. B. Lee et al. "A Scalable Micro Wireless Interconnect Structure for CMPs" presented at MobiCom 09, Beijing, 2009
- [102] W. J. Dally and J. W. Poulton. *Digital Systems Engineering*. Cambridge, UK: Cambridge University Press, 2001, pp. 7
- [103] J. Crowe and B. Hayes-Gill. *Introduction to Digital Electronics*. London, UK: Newnes, 2004, pp. 233



- [104] V. Agarwal et al. "Clock Rate versus IPC: The End of the road for Conventional Microarchitectures." presented at the 27th Annual International Symposium on Computer Architecture, Vancouver, British Columbia, June 10-14, 2000
- [105] IEEE. (2012, Aug. 10) *OFFICIAL IEEE 802.11 WORKING GROUP PROJECT TIMELINES – 2012-08-10* [Online] Available: [http://www.ieee802.org/11/Reports/802.11\\_Timelines.htm](http://www.ieee802.org/11/Reports/802.11_Timelines.htm)
- [106] C. Mathias. (2012, Aug. 28) *Wireless LAN's New Standard, 802.11ac: Prep Time* [Online] Available: <http://www.informationweek.com/mobility/802dot11x/wireless-lans-new-standard-80211ac-prep/240006274>
- [107] G.R. Aiello and G.D. Rogerson, "Ultra-wideband wireless systems" IEEE Microwave Magazine, Vol. 4, Issue 2, June 2003, pp. 36-47
- [108] HAPCOS. *Free-Space Optical Communications on High Altitude Platforms* [Online] Available: [http://www.kn.dlr.de/freespaceoptics/index.htm?group\\_BenGurion.htm](http://www.kn.dlr.de/freespaceoptics/index.htm?group_BenGurion.htm)
- [109] HAPCOS. *What are High Altitude Platforms (HAPs)?* [Online] Available: [http://www.hapcos.org/what\\_are\\_HAPS.php](http://www.hapcos.org/what_are_HAPS.php)
- [110] Laseroptronics. (2012) *The Next Generation of Wireless IR Links* [Online] Available: <http://www.laseroptronics.com/>
- [111] Infra Red Data Association. (2011) *Infra Red Data Association* [Online] Available: <http://www.irda.org/>
- [112] R. J. Green (2011, Oct. 18) *10GigaIR Working Group - Update* [Online] Available: [www.see.ed.ac.uk/~hxx/owc2011/green.ppt](http://www.see.ed.ac.uk/~hxx/owc2011/green.ppt)
- [113] R. Myers (Editor) *The Radio Amateur's Handbook*. 52nd Edition. Newington : American Radio Relay League, 1975, pp25-29.
- [114] Antenna-Theory.com (2011) *The Friis Equation* [Online] Available: <http://www.antenna-theory.com/basics/friis.php>
- [115] B. Allen (2013, Jan. 7) Power Harvesting [Online]. Available e-mail: [rjh540@york.ac.uk](mailto:rjh540@york.ac.uk) Message: re:Power Harvesting
- [116] N. Roy et al. (2006) *Designing an FPGA-Based RFID Reader* [Online] Available: [www.linear.com/docs/40013](http://www.linear.com/docs/40013)
- [117] J. Nie. (2007, Aug. 5) *Antennas for RFIC Transmitter and Receiver Part 1: Loop Antennas* [Online] Available: [www.rfm.com/products/apnotes/antenna\\_appnote.pdf](http://www.rfm.com/products/apnotes/antenna_appnote.pdf)
- [118] C. Edwards. (2011, Mar. 23) *Back to the future as CISC vs RISC argument reopens* [Online] Available: <http://www.electronicweekly.com/blogs/low-power-design/2011/03/cisc-vs-risc-power-argument-reopens.html>
- [119] D. Sutton, *Platonic and Achimedean Solids*. Wooden Books.com: Glastonbury, 2005.