

Rochester Institute of Technology

## RIT Digital Institutional Repository

---

Theses

---

2011

### Towards FPGA hardware in the loop for QCA simulation

Alan Olson

Follow this and additional works at: <https://repository.rit.edu/theses>

---

#### Recommended Citation

Olson, Alan, "Towards FPGA hardware in the loop for QCA simulation" (2011). Thesis. Rochester Institute of Technology. Accessed from

This Thesis is brought to you for free and open access by the RIT Libraries. For more information, please contact [repository@rit.edu](mailto:repository@rit.edu).

# Towards FPGA Hardware in the Loop for QCA Simulation

by

**Alan Olson**

A Thesis Submitted in Partial Fulfillment of the Requirements for the  
Degree of Master of Science  
in Electrical Engineering

Supervised by

Associate Professor Dr. Dorin Patru  
Department of Electrical and Microelectronic Engineering  
Kate Gleason College of Engineering  
Rochester Institute of Technology  
Rochester, New York

May 2011

Approved by:

---

Dr. Dorin Patru, Associate Professor  
*Thesis Advisor, Department of Electrical and Microelectronic Engineering*

---

Dr. Marcin Lukowiak, Assistant Professor  
*Committee Member, Department of Computer Engineering*

---

Dr. Dhireesha Kudithipudi, Assistant Professor  
*Committee Member, Department of Computer Engineering*

# Thesis Release Permission Form

Rochester Institute of Technology  
Kate Gleason College of Engineering

Title:

Towards FPGA Hardware in the Loop for QCA Simulation

I, Alan Olson, hereby grant permission to the Wallace Memorial Library to reproduce my thesis in whole or part.

---

Alan Olson

---

Date

# Acknowledgements

Dr. Lukowiak and Dr. Kudithipudi for their insightful review

Matt Hibbard for his assistance with the Wildcard 4 system

Konrad Walus and Gabriel Schulhof for their correspondence  
regarding QCADesigner

My parents Grieg and Betty Olson for their undying love and support

Dr. Eric Peskin for his assistance and guidance during the early parts  
of this work

and

God for his faithfulness

# Abstract

## Towards FPGA Hardware in the Loop for QCA Simulation

Alan Olson

**Supervising Professor: Dr. Dorin Patru**

As transistors begin to hit raw physical limits and performance barriers, other technologies are being researched to potentially replace conventional integrated circuit technology. Quantum-dot Cellular Automata (QCA) is one such technology which executes computations using coulomb interactions and quantum-mechanical effects. Part of this research is pursuant to the design of circuits which exploit QCA technology and take advantage of what it has to offer. These circuits must be simulated to ensure their functionality and help prove the viability of QCA. These simulations, like many scientific computing applications, can take a long time to complete; hours or days, depending on their size and complexity. Many scientific applications have benefitted from research into Field Programmable Gate Array (FPGA) application development, which has been used to accelerate the speed at which such simulations execute.

This thesis investigates the possibility of using FPGAs to accelerate the simulation of QCA circuits. The hardware developed is a streaming type architecture using floating point arithmetic and hardware/software techniques. Hardware implementation shows the system to run slower than the existing software code, but demonstrates the ability to simulate a small QCA circuit. Analysis of the design reveals good potential for achieving speedup, and an alternate design is proposed to improve the execution time. In the course of this work, improvements to the existing software are also developed and contributed to the community.

# Contents

<b>Acknowledgements</b> . . . . .	<b>iii</b>
<b>Abstract</b> . . . . .	<b>iv</b>
<b>1 Introduction</b> . . . . .	<b>1</b>
1.1 Motivation . . . . .	1
1.1.1 QCA Advancement . . . . .	2
1.1.2 FPGA Advancement . . . . .	4
1.2 Previous Work . . . . .	5
1.3 Contributions . . . . .	6
1.4 Organization . . . . .	7
<b>2 QCA Background</b> . . . . .	<b>8</b>
2.1 Basic QCA Operation . . . . .	8
2.1.1 QCA Cell . . . . .	8
2.1.2 QCA Wires . . . . .	9
2.1.3 QCA Logic . . . . .	10
2.1.4 The Clock . . . . .	11
2.2 Implementations . . . . .	14
2.2.1 Metal-Dot . . . . .	14
2.2.2 Molecular . . . . .	15
2.2.3 Magnetic . . . . .	15
2.3 QCA Designer . . . . .	16
2.3.1 Bistable Simulation . . . . .	16

2.3.2	Coherence Vector Simulation . . . . .	17
2.3.3	Simulation Speed . . . . .	18
<b>3</b>	<b>Using FPGAs for Application Speedup . . . . .</b>	<b>20</b>
3.1	Basic Architecture and Functionality . . . . .	20
3.2	Use In Application Speedup . . . . .	21
3.3	Number Systems . . . . .	23
3.4	Hardware/Software Design . . . . .	25
3.5	Nature of QCADesigner Computation . . . . .	26
3.6	The Wildcard Hardware . . . . .	27
<b>4</b>	<b>Proposed Simulation Architecture . . . . .</b>	<b>29</b>
4.1	Software Analysis . . . . .	29
4.1.1	Code Analysis . . . . .	29
4.1.2	Numerical Analysis . . . . .	34
4.2	Hardware Architecture . . . . .	38
4.2.1	Proof of Concept . . . . .	38
4.2.2	DRAM Interfacing . . . . .	40
4.2.3	Other Interfacing Issues . . . . .	44
4.2.4	The Arithmetic Core . . . . .	45
4.2.5	Operation and Synthesis . . . . .	48
<b>5</b>	<b>Results and Discussion . . . . .</b>	<b>51</b>
5.1	Cause of Slowdown . . . . .	51
5.2	Hardware Reliability . . . . .	54
5.3	QCADesigner Enhancements . . . . .	55
5.4	Design Alternatives . . . . .	56
5.4.1	New Architecture . . . . .	56
5.4.2	Other Platforms . . . . .	57

<b>6 Conclusion and Future Work . . . . .</b>	<b>59</b>
<b>Bibliography . . . . .</b>	<b>60</b>
<b>A run_bistable_simulation C Code Listing . . . . .</b>	<b>69</b>
<b>B run_bistable_simulation_hardware C Code Listing . . . . .</b>	<b>80</b>
<b>C qcad VHDL Code Listing . . . . .</b>	<b>96</b>
<b>D p_math VHDL Code Listing . . . . .</b>	<b>118</b>



# List of Tables

2.1	Truth table for a majority gate . . . . .	12
-----	---	----

# List of Figures

2.1	A QCA cell (reproduced from [45]) . . . . .	9
2.2	A normal QCA wire (reproduced from [45]) . . . . .	10
2.3	An inversion chain (reproduced from [45]) . . . . .	10
2.4	A wire crossing in QCA using an inversion chain (reproduced from [45]) . . . . .	11
2.5	Majority gate in QCA . . . . .	12
2.6	An inverter gate in QCA (reproduced from [45]) . . . . .	13
2.7	Waveform showing how QCA clock zones operate (reproduced from [52]) . . . . .	13
2.8	Clock zones as they should appear in a wire . . . . .	14
3.1	Basic mathematical structure of 2D convolution algorithm (reproduced from [37]) . . . . .	22
3.2	Binary reduction tree structure used in [32] . . . . .	23
3.3	A 32 bit floating point number . . . . .	24
3.4	Block diagram of Wildcard 4 hardware . . . . .	28
4.1	QCADesigner simulation using single precision floating point . . . . .	36
4.2	QCADesigner simulation using double precision floating point . . . . .	37
4.3	Hardware results from minisim test . . . . .	40
4.4	Software results from minisim test . . . . .	41
4.5	Block diagram of the polarization math pipeline . . . . .	46

4.6	Block diagram of high level system architecture . . . . .	49
4.7	Flowchart for the operation of the combined hardware/- software system . . . . .	50
5.1	Result for software run of a majority gate . . . . .	52
5.2	Result for a hardware run of a majority gate . . . . .	53

# Chapter 1

## Introduction

There is constant demand for computing resources to become more powerful, more energy efficient, and less expensive. For decades, the semiconductor industry has been able to keep up with this demand, doubling processor performance every 18 months according to Moore's Law [31]. Recently, however, semiconductor-based transistors have begun to approach absolute physical barriers that will prevent further progress at this pace. The size of state of the art transistors is not far from the atomic scale, and their behaviour is increasingly determined by quantum effects. As such, it is becoming ever more difficult to continue with the current scaling trend, and alternate paths are being researched, which specifically leverage the quantum effects at this level. QCA is one such technology, with research ongoing to demonstrate its feasibility and utility. Accelerating the simulations that are performed in the course of this research can help speed up the overall progress, just as with other research fields. This can be carried out by using FPGAs to implement specialized hardware that exploits parallelism for a more efficient architecture.

### 1.1 Motivation

There are two primary motives to perform this work. First is to advance the state of QCA research as a means to overcome the performance barriers currently faced by the semiconductor industry. Second is to advance the state of FPGA utilization in support of simulation speedup.

### 1.1.1 QCA Advancement

Currently, the primary manufacturing technology for integrated circuits is based on complementary metal-oxide-semiconductor (CMOS). This technology is presently facing absolute theoretical scaling limits as current leakage, power consumption, oxide thickness, mobility degradation, threshold voltage variation and many other characteristics become increasingly difficult to control [1]. This means that traditional scaling techniques are no longer sufficient to maintain the pace of semiconductor development.

The International Technology Roadmap for Semiconductors (ITRS) [1] outlines myriad directions industry research may take in addition to traditional scaling to keep advancement proceeding apace. These include equivalent scaling, design equivalent scaling, Emerging Research Materials (ERM), and Emerging Research Devices (ERD). Equivalent scaling and design equivalent scaling focus on new forms of architectures, such as parallel or three dimensional computing structures to improve performance through design rather than implementation technology. ERM and ERD focus on developing alternate technologies to augment or replace CMOS. The field of ERM includes the development of new channel materials to facilitate the advancement of CMOS technology as well as materials to enable further development of ERD topics such as carbon nanotubes, silicon nanowires and spin materials.

QCA is one of the technologies mentioned in the ITRS. Materials research regarding macromolecules and self-directed assembly benefits QCA, and the concept of QCA is under scrutiny as a possible replacement for CMOS. It has many potential advantages over CMOS with regard to device density [51], operating speed [28], and power consumption [9], and it may even serve as the basis for a reversible computing architecture [26, 25]. These are all desirable characteristics for a technology that supports data processing to have, and indeed are sought-after goals stated in the ITRS. However, a direct comparison to CMOS is difficult because transistors and QCA cells are not directly analogous (see Section 2.1). Furthermore, none of these aspects of QCA have been

fully realized in practice yet, nor have QCA circuits been manufactured at any significant scale. QCA research is ongoing, focusing on two distinct aspects: implementation and architecture.

One direction of QCA research is that of physical implementation. Although the technology shows promise, the current state of QCA is highly experimental, and it must be made more practical if it is to compete with and supercede, CMOS. QCA must be cheaply and reliably manufacturable, operational at room temperature, and able to deliver the high speed and low energy consumption demanded of it.

There are three basic types of QCA proposed: metal-dot, molecular, and magnetic. Magnetic QCA (see Section 2.2.3) has been shown to operate at room temperature, but is not able to be clocked much faster than 100 MHz [7]. Metal-dot QCA (see Section 2.2.1) promises to operate at speeds up to 10 GHz, but has not been shown to work at room temperature [43]. Molecular QCA (see Section 2.2.2) purports to be able to achieve room temperature operation [42] at up to THz speeds [28], but no actual circuits have been manufactured at any scale. These operating speeds have been determined using simulations of the adiabatic switching properties of the materials involved, but have yet to be verified in a laboratory. This research is currently the domain of physicists, material scientists, and microelectronic engineers.

This thesis relates more closely to the other important aspect of QCA research, which is nanoarchitecture research [45]. This research models the physics of QCA and ignores the manufacturing concerns so that the engineer may focus on more familiar elements like cells, wires, and gates. It is important to demonstrate that novel circuits can be created using QCA, thus making the aforementioned physical implementation research worthwhile. It also provides optimized circuits “in the pipe,” ready to manufacture as soon as possible. This research can proceed in parallel with the other, without the immediate need for physical circuits to test on, but is still held back by the large amount of time needed to simulate each design. If accurate simulations could be sped up, rapid progress in this field will ensue, verifying and correcting the functionality of large QCA circuits in hours rather than days or

weeks.

### 1.1.2 FPGA Advancement

Field Programmable Gate Arrays (FPGAs) are hardware devices that can be programmed to implement any arbitrary logic function. There are many ways they can be designed, but usually they consist of a large array of multiplexer-based Look-Up-Tables (LUTs) that can be individually programmed and interconnected according to a programming bit stream [16]. FPGAs are used in a variety of contexts today, from Application Specific Integrated Circuit (ASIC) prototyping, to implementation platforms for end products, to a means of scientific research.

FPGAs are powerful and versatile. Unlike ASICs, which can only serve one specialized purpose after manufacture, FPGAs can be easily reconfigured to implement a large variety of hardware tasks. They also have the advantage compared to CPUs because their hardware structure can be modified and optimized for very specific tasks, and can have a very high degree of parallelism. Therefore they have become a way to achieve performance similar to ASICs without the high monetary and labour cost of developing and manufacturing an actual ASIC chip. Furthermore, a single design can be shared between groups that use different models or brands of FPGAs, and that design can be quickly implemented, studied, and used without any need to purchase a new expensive ASIC. Thus, FPGAs are very valuable to the scientific community and electronics industry as they enable the rapid distribution, testing, and implementation of novel designs.

Up to now, FPGAs have been typically used to implement fixed point or limited precision applications; due to constraints on hardware resources, IEEE-compliant floating point cores were not able to be implemented at useful speeds. However, the performance of FPGAs has been increasing even more rapidly than that of CPUs [47], and their floating point performance has improved significantly. In [47], Keith

Underwood shows that floating point performance on FPGAs has indeed advanced to match and exceed that of general purpose CPUs. Given the importance of floating point precision to scientific applications, and the usefulness of FPGAs in general, this is a promising development. Despite this advancement, relatively few applications have been developed which take advantage of it, and none have been found that combine floating point and extensive hardware/software design. Hardware/software design (see Section 3.4) can be exploited to speed up the development process by porting scientific applications from software to hardware. Showing that hardware/software design can be used effectively with floating point applications would be significant.

## 1.2 Previous Work

Up to now, a fair bit of work has been done to develop a number of novel circuits that could be implemented using QCA, including a microprocessor [50], fault tolerant architectures [44, 52], and various programmable logic structures [33, 23, 9, 46, 11]. Specifically, work has been done at the Rochester Institute of Technology (RIT) to develop a Combinational Logic Block (CLB) that would form the basis of a QCA-based FPGA [46]. Thus it can be seen that nonoarchitecture research is progressing, and various circuit designs exist that could be implemented with QCA when the time comes.

There are several simulation tools that have been made to help circuit designers verify the functionality of their designs. These tools include QCADesigner, MAQUINAS, QBART and HDLQ [36]. QCADesigner and MAQUINAS both focus on the cell level (for metal-dot and molecular implementations respectively), and execute numerical simulations that are based on Schrödinger equations. QBART instantiates QCA devices into a grid, and uses *ad hoc* simulation rules. This allows simulations to run faster, but the simulation environment is less flexible and not as accurate. HDLQ is a Verilog library which provides a basis for behavioural simulation of QCA circuits using existing Verilog simulation tools. The difference between HDLQ and QCADesigner is



somewhat akin to the difference between Verilog and SPICE. At RIT, QCADesigner has been used for verification of small-scale circuits, and HDLQ for circuits where QCADesigner’s simulation time proves impractical. For reference, the simulation of one CLB (with 22,558 QCA cells) in [46] takes approximately 50 minutes, while the QCADesigner simulation of four CLBs takes between 8 and 12 hours.

There has been much work done developing applications for FPGAs in many categories to varying degrees of success. Image processing is commonly implemented due to its relatively light numerical demands and significant parallelizability. There are also numerous Molecular Dynamics (MD) simulations implemented on FPGAs typically achieving speedups of 5x over pure software [14]. MD simulations often start as software applications using floating point, but find themselves using some other number system such as logarithmic floating point [6] or semi-floating point [15] (see Section 3.3) that makes the hardware more compact. The precision suffers somewhat, but the size of mathematical cores is reduced to allow for greatly enhanced parallelism. A limited number of applications using fully floating point numbers have also been created to speed up various applications, in some cases up to 35x [32], showing the ability of FPGAs to handle floating point applications.

### 1.3 Contributions

This thesis attempts to make steps toward accelerating the execution of QCA simulations, a task which is not shown to have been attempted in the literature. The design used to this end is novel in concept among other similar applications because it attempts to use IEEE 754 compliant floating point arithmetic in conjunction with hardware/software design in order to achieve high numerical precision simultaneously with shortened development time and ease of use. One possible architecture is designed and tested, but fails to achieve any speedup due to interfacing bottlenecks. The computational core itself, however, demonstrates the operational speed necessary to serve as the basis for faster systems in the future. Another design is proposed that would eliminate the

most significant observed issues and potentially run much faster, while maintaining floating point precision and hardware/software usability. Additionally, some minor improvements are made to the QCADesigner software, including a bug fix in the algorithm and an optimization that significantly speeds up a common function.

## 1.4 Organization

This thesis is organized as follows: Chapter 2 provides the background to QCA necessary to understand what is being simulated, including descriptions and background of existing software simulators. Chapter 3 offers a detailed look at FPGAs, their basic architecture and functionality, and how they are used to accelerate various computing applications. This section also contains a brief examination of different number systems, and how they may be applied in different designs. Chapter 4 details the design process, including the prerequisite software analysis and explication of the final implemented hardware. Chapter 5 reports and analyzes the observed results and Chapter 6 discusses their significance and suggests directions for future work.

# Chapter 2

## QCA Background

Quantum-dot Cellular Automata (QCA) is a technology concept devised by Tougaw and Lent in 1993 [24]. It is a way to create digital electronic devices without using transistors that has potential to be faster, denser, and more energy efficient than CMOS technology. Rather than computing with voltage levels requiring current flow, QCA computes via field polarization [2], in such a way that the ground state of the system corresponds to the solution state of the computation. Much progress has been made since 1993 in terms of design and fabrication techniques. This section explains in detail how this form of computation works and can be implemented. Note that the logical architecture and functionality is the same independent of the method of manufacturing and implementation.

### 2.1 Basic QCA Operation

This section describes how data is stored, transferred, and manipulated in QCA, starting with a description of a basic cell, then building wires and gates out of cells. The clocking structure, which is central to the functionality of QCA, is also described.

#### 2.1.1 QCA Cell

A QCA Cell is the basic building block of any QCA circuit [3]. It consists of two pairs of charge sites, each pair joined by a tunnel junction,

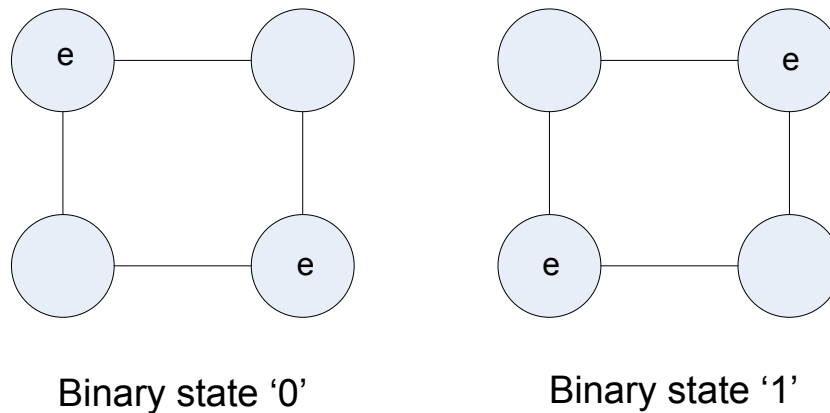


Figure 2.1: A QCA cell (reproduced from [45])

and two distinct charges. The specific configuration of the cell represents a binary value stored in the cell as shown in Figure 2.1. There are only two stable configurations in a QCA cell, corresponding to the electrons occupying opposite corners, hence representing two alternate binary values. The charges can move between charge sites by means of quantum tunneling, a process which is facilitated by the clock (section 2.1.4). The configuration of a cell can be calculated in simulation by using a Schrödinger equation with a quantum-mechanical Hamiltonian, and is influenced by any other QCA cells surrounding it. A cell can also theoretically be arranged with the charge sites rotated by 45 degrees to facilitate wire crossings [9].

### 2.1.2 QCA Wires

When QCA cells are arranged in a line, they form a wire. In contrast to CMOS technology where simple metal lines are used to transfer data, it is the same complex functional devices that process data in QCA which transmit data. A QCA wire functions by local coulombic interactions from one cell to the next, generating a “domino effect” that propagates data from a fixed input to the measured output as shown in Figure 2.2.

Alternately, rotated cells may be arranged to form an inversion chain




Figure 2.2: A normal QCA wire (reproduced from [45])




Figure 2.3: An inversion chain (reproduced from [45])

as shown in Figure 2.3. This is like a normal QCA wire, but the polarization flips between each cell. Therefore an inversion chain must be an odd number of cells in order to preserve the value from input to output. The main use of an inversion chain is to allow coplanar wire crossings [9]. In a wire crossing such as that shown in 2.4, the cell at the crossover point in the primary wire has a neutral coulombic effect on either of the nearby cells in the secondary wire, allowing data to cross over unaltered. Multi-layer QCA circuits have not been shown to be feasible, but coplanar crossings such as the ones described here may be. Wire crossings, however they are implemented, are important to having novel circuits in a compact area.

### 2.1.3 QCA Logic

The logical basis for QCA circuits are majority gates and inverters. A QCA inverter has the same function as the familiar CMOS inverter, *i.e.* inverts a logical value from input to output. A majority gate is a three-input gate whose output matches the most popular input value. The boolean expression for this is  $F = AB + BC + AC$ , and the truth table for a majority gate is shown in Table 2.1. The QCA schematics of an inverter and a majority gate are shown in Figure 2.5 and 2.6.

It can be seen from the schematic that a three-input majority gate is fundamentally the most complex single gate structure that can be

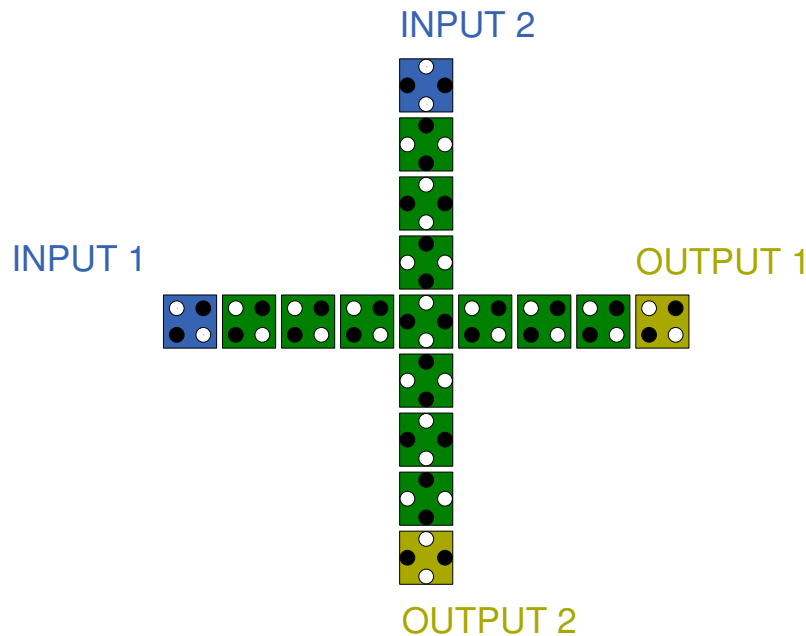


Figure 2.4: A wire crossing in QCA using an inversion chain (reproduced from [45])

implemented in QCA. It functions due to the bistable nature of a QCA cell, which will “snap” into one configuration or another based on which cell polarity has a more dominant influence. A majority gate can be programmed to function as either an AND or an OR gate by holding one of the inputs at a logical zero or one respectively. With AND, OR and INV, a complete set of logical devices can be implemented in QCA, enabling the creation of any digital circuit that can be made currently in CMOS.

#### 2.1.4 The Clock

All QCA functions are driven by an underlying clock structure which is essential to the controlled operation of each cell and governs the flow of data. There are actually four clocks in a QCA circuit, each separated by 90 degrees of phase shift, and each cell is attached to one of the four clock zones. The four phases of each clock period may be labeled as relax, switch, hold, and release, indicating each phase’s effect on a cell [52]. A timing diagram showing their basic progression is shown in

A	B	C	F
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	1
1	0	0	0
1	0	1	1
1	1	0	1
1	1	1	1

Table 2.1: Truth table for a majority gate




Figure 2.5: Majority gate in QCA

## Figure 2.7

During the relax phase, a cell is free to switch polarizations based on the surrounding cells' influence. In the switch phase, the clock level progresses from low to high, and the polarization of the cell becomes locked in place. When the clock reaches the high level, this is the hold phase, and the polarization of the cell is fully locked and able to influence other cells without being altered itself. After the hold phase comes the release phase where the clock progresses from high to low, and the cell gradually loses its fixed value to re-enter the pliable relax state.

These alternating clock phases allow the controlled progression of




Figure 2.6: An inverter gate in QCA (reproduced from [45])

data in one direction through a QCA circuit, from input to output, with minimal error. Figure 2.8 illustrates the prescribed arrangement of clock zones in a wire. Input is on the left, output is on the right, and each colour represents a different zone, progressing in order from Clock 1 to Clock 4 from left to right. However, because every cell is clocked, including those making up the wires, QCA is a deeply pipelined architecture. Every full set of clock zones between input and output adds one clock period of latency to the circuit, and there is a limit to how many consecutive cells may be in one clock zone. Therefore, even wires will add to circuit latency, meaning the circuit designer must take extra care in aligning data arrival times.

The clock structure in QCA is currently implemented as a standard CMOS clock network directly below the QCA circuit whose voltage level directly affects the depth of the energy wells holding the charges. The clock can generally be routed wherever it is needed in the circuit, but with four clocks to route, the actual layout could be quite complex, and minimizing skew can be difficult. Research is ongoing to develop techniques to design QCA circuits with more regular and easily routable clock zones.




Figure 2.7: Waveform showing how QCA clock zones operate (reproduced from [52])






Figure 2.8: Clock zones as they should appear in a wire

Concerns have been raised that the power savings of QCA may be offset by the power consumption of a clock running at the proposed speeds of 10 GHz and more. This is addressed in [4], where Lent *et al.* simulate the power consumption of the type of clock network used in QCA. Their findings show that the clock does not significantly contribute to the overall power consumption of a QCA circuit, and that a 100 GHz QCA clock consumes similar power to a 2 GHz clock driving a conventional CMOS circuit.

## 2.2 Implementations

Three methods of implementation for QCA have been proposed: Metal-dot, molecular, and magnetic. They each have their own advantages and disadvantages, and are all actively being researched.

### 2.2.1 Metal-Dot

The most popular form of QCA implementation being researched is the metal-dot implementation. This uses small dots of conductors such as Aluminum [22], Silicon or Gallium-Arsenic [2] as charge sites. Each 4-dot cell in a metal-dot QCA cell is typically a square of about 20 nm on a side. It has been shown to work in a lab environment, in which some very basic, though important, circuits have been demonstrated at temperatures below 1 K [34, 35]. In theory, based on adiabatic switching simulations of metal-dot cells, the clock speed of this form of QCA circuit can reach into the tens of GHz, but only up to MHz range has been demonstrated [43]. Metal-dot is the form of QCA best-suited to use current CMOS manufacturing lines without too much upgrading or modification being needed; it would use the same basic lithography techniques, and has similar feature sizes as the current state of the art

transistor technology. It is proposed that if the individual cells and dots can be made smaller, higher operating temperatures may be achieved.

### 2.2.2 Molecular

Molecular QCA is a step up from metal-dot in every way, except feasibility. Molecular QCA proposes composing a QCA cell out of just two molecules specially designed to have two primary charge sites and mobile electrons. In theory, because individual molecules are being used, the “dots” are small enough to allow operating temperatures up to 450 K [51], *i.e.* well above room temperature. It is also theorized that molecular QCA supports the highest clock speeds out of the three implementations: up to 1 THz [28]. However, development of molecular QCA has not progressed very far beyond simulation. Several candidate molecules are being studied, and have been experimentally shown to behave as a QCA cell, but no complex circuits have been manufactured to date.

### 2.2.3 Magnetic

Magnetic QCA is slightly different from metal-dot and molecular, because instead of charges jumping between charge sites in a cell, a cell consists of a magnetic dipole that changes its alignment based on its neighbours. The same type of architecture and logic applies, only the mechanics of how it works changes. Magnetic QCA has been shown to work at room temperatures, and possesses a natural hardness against radiation, but simulation of the switching properties of a magnetic QCA cell suggest that it may only work at speeds up to 100 MHz [7]. Despite the low operating speed, magnetic QCA may find applications in areas where low power consumption and high reliability are more important than operating speed, such as in medical implants or space exploration.

## 2.3 QCA Designer

The simulation tool that RIT and many other institutions [44, 18, 41, 46] use to test their QCA circuits is QCADesigner [48, 49]. It is a program that allows a designer to easily create and test a QCA circuit using a graphical user interface (GUI). It offers two different simulation engines (coherence vector and bistable) with a variety of customizable parameters. QCADesigner was initially developed by Konrad Walus *et. al.* in the Microsystems and Nanotechnology Group at the University of British Columbia, but is now an open source project, so obtaining and modifying the source code is very easy, making it an ideal choice for an application to accelerate.

### 2.3.1 Bistable Simulation

The bistable simulation engine is the preferred QCADesigner engine for simulating circuits at RIT, and therefore the engine selected for the work of this thesis. It models a QCA cell as a system having two stable states, and the polarization of a cell is calculated based on the surrounding polarizations and kink energies. The kink energy between two cells refers to the energy cost of the two cells having opposing polarizations. Naturally, this value tends to decrease as the distance between cells increases. The main equation used to determine the kink energy is shown below [48].

$$E_{i,j} = \frac{1}{4\pi\epsilon_0\epsilon_r} \frac{q_i q_j}{|r_i - r_j|} \quad (2.1)$$

This equation is used to calculate the electrostatic energy between two individual dots in cell  $i$  and cell  $j$ .  $\epsilon_0$  is the permittivity of free space, and  $\epsilon_r$  is the relative permittivity of the material system in use, which can be specified by the user in the simulation setup window.  $q_i$  and  $q_j$  are the charges stored on the two dots from cell  $i$  and  $j$  respectively, and  $|r_i - r_j|$  is the distance between the dots. This equation is applied and summed for every combination of dots in each cell, and

the difference between the results for aligned and misaligned polarities is taken as the kink energy. Because there is no net charge on the cell, the interaction is quadrupole-quadrupole based, so it decays very rapidly as distance increases. The set of kink energies for all cells only needs to be calculated once during a simulation, during the initialization phase before any polarizations are calculated.

The kink energies between a cell and its neighbours are then used as terms in the equation below to determine the polarity:

$$P_i = \frac{\frac{1}{2\gamma} \sum_j (E_{i,j}^k P_j)}{\sqrt{1 + \frac{1}{2\gamma} \sum_j (E_{i,j}^k P_j)}} \quad (2.2)$$

where  $\gamma$  is the tunneling potential, which changes with the clock value,  $P_j$  is the polarization of the  $j$ th neighbour of the current cell, and  $E_{i,j}^k$  is the kink energy between the current cell and the  $j$ th neighbour cell. This equation is carried out once for every cell in a QCA circuit, for every iteration the simulation goes through. One “sample” of the simulation may consist of several iterations of the circuit, in which the state of each cell is recalculated until a maximum allowable change threshold is reached. This bistable simulation engine is able to simulate a large number of cells relatively quickly to verify logical correctness and basic structural validity, but it lacks certain considerations, such as temperature and dissipative effects, required to execute dynamic simulations with a high degree of accuracy.

### 2.3.2 Coherence Vector Simulation

The coherence vector simulation engine is a slightly more detailed simulation than the bistable simulation engine. Like the bistable engine, it assumes a cell is basically a bistable system, but the state is calculated in a more sophisticated way. The status of a cell is tracked using a coherence vector where the polarization is represented as the z component of this vector. Each component of the vector is calculated by multiplying the trace of the density matrix by each of the Pauli spin

matrices.

$$\lambda_i = Tr\{\hat{\rho}\hat{\sigma}_i\} \quad i = \{x, y, z\} \quad (2.3)$$

The main driver of the coherence vector simulation is

$$\frac{\partial}{\partial t} \vec{\lambda} = \vec{\Gamma} \times \vec{\lambda} - \frac{1}{\tau}(\vec{\lambda} - \vec{\lambda}_{ss}) \quad (2.4)$$

which models the change in the cell's state over time, including dissipative effects. The Gamma vector is defined as

$$\vec{\Gamma}\{x, y, z\} = \frac{1}{\hbar}[-2\gamma, 0, \sum_{j \in S} E_{i,j}^k P_j] \quad (2.5)$$

where gamma is the tunneling potential (tied to the clock),  $S$  is the effective neighbourhood of cell  $i$ , and  $E_{i,j}^k$  and  $P_j$  are the kink energy and polarization respectively of cell  $j$ .  $\tau$  is the relaxation time of the cell.  $\vec{\lambda}_{ss}$  from equation 2.4 is the steady-state coherence vector, and is defined as

$$\vec{\lambda}_{ss} = -\frac{\vec{\Gamma}}{|\vec{\Gamma}|} \tanh \Delta \quad (2.6)$$

where delta is the temperature ratio:

$$\Delta = \frac{\hbar |\vec{\Gamma}|}{2k_B T} \quad (2.7)$$

where  $T$  is the temperature in kelvin and  $k_B$  is Boltzmann's constant. The simulator evaluates these equations for each cell in the circuit for each timestep until the end time is reached, as specified by the user. It is plain to see why this method may be more accurate and take much longer to run.

### 2.3.3 Simulation Speed

Each simulation of a 22,558 cell CLB takes a little less than an hour to complete using the bistable simulation engine. The bistable simulation

engine is used in our research because it provides basic architecture verification an order of magnitude faster than the coherence vector simulation. For larger circuits, exponentially more time is required. For example, a circuit four times as big with four CLBs requires 10 hours to complete the simulation. The CMOS versions of these circuits would not be extraordinarily complex in terms of the number of gates (*i.e.* devices), 572 gates for one CLB; but because gates *and* wires are made up of complex fundamental devices in QCA, the simulation is much more complex and time consuming.

It is observed that for both of these simulation engines, the new polarization of a cell may be calculated concurrently with the new polarization of any and all other cells without flouting any data dependencies. This means that there is a great amount of parallelism that may be exploited when executing these simulations. If one computational core can be assigned to each cell, then a very large circuit would take the same time to simulate as a small circuit currently does, yielding extraordinary speedup. Such a task is well-suited to FPGAs, as they can easily replicate custom processing hardware into multiple instances on one chip.

## Chapter 3

# Using FPGAs for Application Speedup

An FPGA is a chip containing a reconfigurable hardware structure which can be programmed to implement any desired hardware circuit. It has no fixed functionality, and unlike a CPU, programming it results in an actual reassignment and reconfiguration of hardware resources. FPGAs often have a repeating structure with computing and memory resources spread throughout the chip, allowing them to be used for a large number of applications. Their versatility allows them to serve as specialized application specific hardware platforms without the high development costs associated with ASICs. They can be used for hardware prototyping before sending a design to ASIC production, or may be used as end platforms in themselves.

### 3.1 Basic Architecture and Functionality

Although there is no universal design standard for an FPGA, the most common approach is to use a combination of programmable logic blocks and programmable interconnects. Logic may be programmed using look up tables that consist of SRAM cells attached to multiplexer inputs, and interconnects may be programmed using solid-state switches. Some FPGAs may include embedded multipliers, digital signal processing (DSP) units, block RAM, and other features, but the basic principle remains the same. FPGAs are usually laid out in a repeating grid pattern, which facilitates their use for parallelizing calculations. Computing resources are distributed evenly throughout the chip, allowing multiple computational cores to be implemented at once, cores

that may or may not be identical to each other. Memory resources are also distributed such that computing resources in one region may use memory in that region while other memory is free to be used by other resources, resulting in a very large effective memory bandwidth.

## 3.2 Use In Application Speedup

A large class of applications commonly implemented on FPGAs is image processing. These are ideally suited to FPGAs because they rely on simple fixed point mathematics and their operation is often highly parallelizable. Therefore, there is a long history and large body of work related to accelerating image processing on FPGAs. Of particular interest is the 2D Convolution algorithm [37], often used for image noise reduction, which replaces a pixel value with a weighted average of surrounding pixel values. The mathematical algorithm has some similarities to that used in the bistable simulation engine of QCADesigner, both using a sum of products term based on surrounding values. It calculates the new value of a pixel by computing the sum of products of surrounding pixel values, times the convolution weights (see figure 3.1). There is even an example of how to use a basic 3x3 convolver design to compute arbitrarily large convolution kernels. This might be analogous to using a basic QCA simulation core with a 50 nm radius of effect to enable the simulation of any larger radius of effect without having to use a different core.

Many advanced scientific applications are also executed on FPGAs, using a variety of implementation methods. In [6], an alternative number system, logarithmic floating point (see Section 3.3), is implemented and an application to accelerate Quantum Chromodynamics calculation is constructed using Handel-C. Handel-C is one of many alternatives to VHDL, like JHDL [19] and Carte [40, 21], that abstract the hardware description task to a level more familiar to software developers. This can make development easier by lowering the learning curve and making the hardware description language more similar to the original language of the program being adapted. Alternative number systems are



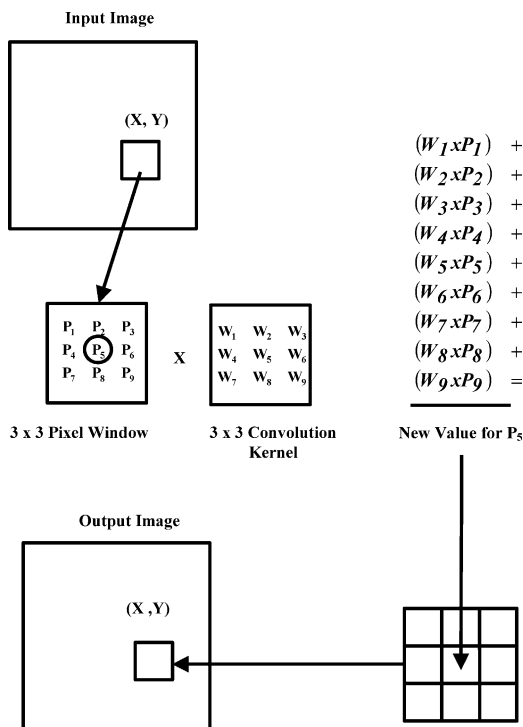


Figure 3.1: Basic mathematical structure of 2D convolution algorithm (reproduced from [37])

also a common practice in FPGA development to reduce the hardware demands of a system where possible [15, 12, 13].

Molecular Dynamics (MD) simulations are also often implemented on FPGAs to good effect, often achieving greater than 5x speedup [21, 15, 14, 13], sometimes much more. A variety of architectures and methodologies are used to obtain these results, including hardware/software implementations [21, 40, 15, 39, 13], discrete event simulation [29], and multigrid computing [14]. In hardware/software implementations, only a portion of the application is ported to hardware, while other parts are left done in software for ease of development and optimum resource usage. To speed up the simulation, discrete event simulations use a variety of simplifications to advance MD simulations by event rather than timestep. The multigrid computing technique breaks the computation up into spatial grids to parallelize computation and memory access, and iteratively uses different sized grids to balance between

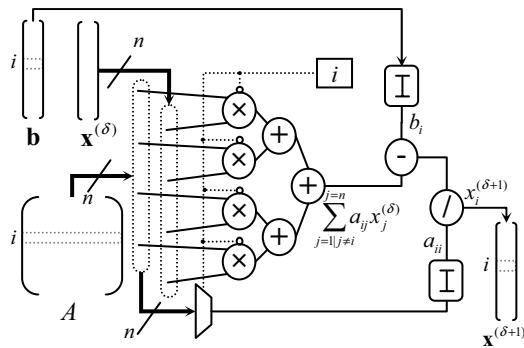


Figure 3.2: Binary reduction tree structure used in [32]

speed and accuracy.

Some floating point applications, including SPICE simulation [20], matrix operations [10, 27, 32], and Fast Fourier Transforms (FFTs) [17] have also been implemented on FPGAs with speedups of up to 18x compared to software implementations. For matrix operations, which involve relatively simple multiply-accumulate mathematics, many floating point cores can be implemented using spatial parallelism to produce a result quickly in a binary reduction tree pipeline. For example, the one used in a Jacobi iterative solver (see Figure 3.2) [32]. However, for the more complex mathematical demands of SPICE and FFT, some spatial parallelism is exchanged for longer pipelines that continuously compute a stream of data.

### 3.3 Number Systems

There are a variety of number systems that can be used for calculation in FPGA application, and which one is selected depends on the requirements of the application. Floating point is typically preferred in scientific applications due to its high dynamic range and good relative accuracy. This means that a very large range of values can be stored while at the same time having a very small percent error, independent of the magnitude. Floating point comes in two basic flavors: single precision and double precision. The IEEE 754-2008 Standard for Floating Point Arithmetic [8] defines single precision to be a 32 bit

0	1	0	0	0	1	1	1	0	0	0	1	0	0	1	1	1	0	0	1	1	0	0	1	1	0	1	1	1	1		= 37235.215
sign	exponent = 142 - 127 = 15								mantissa = 1 .00100010111001100110111 = 1.1363286																						

Figure 3.3: A 32 bit floating point number

number with 1 sign bit, 8 exponent bits, and 23 mantissa bits. Double precision has 1 sign bit, 11 exponent bits, and 52 mantissa bits to make a 64 bit number. Both of these precision levels function in the same way. The exponent is determined by the exponent bits' difference from a bias value to give a positive or negative exponent. The mantissa bits combine with an unrepresented implied leading 1 to give the base value on which the exponent will operate, and the value is read in decimal as  $(-1)^s * m * 2^e$ . Figure 3.3 shows a 32 bit floating point example. This is a very versatile method of representing numbers. However, due to the complexity of operating on such representations, floating point arithmetic cores are large compared to other implementable systems, so they are often eschewed in favor of simpler number systems.

The simplest form of number system is fixed point, often realized as just integer arithmetic. This system is used in applications where very little dynamic range is needed, such as some forms of image processing [37, 38, 30, 5]. Fixed point arithmetic cores are simple, compact, and fast, but only because the computations are simple. They are not very useful in scientific applications where higher dynamic ranges and more accuracy are required.

There are other number systems which are more versatile than fixed point, but easier to implement than floating point. A logarithmic floating point number system converts floating point numbers into a logarithmic floating point domain, allowing logarithm arithmetic to be used [6]. This means that multiplication, division, and square root operations are all made much simpler, but addition and subtraction are more complex than in standard floating point. This is because when multiplying or dividing logarithms, only simple addition is required whereas addition and subtraction operations require large look up tables. This number system is effective when the majority of operations are multiplication, division, or square root, but the advantage is lost if too many

additions or subtractions are present.

A system called semi-floating point [15] may also be used. In this system, the decimal may shift only to certain places. This allows the arithmetic to be carried out using a limited number of fixed point cores or look up tables, simplifying the hardware requirements compared to floating point. However, it requires that a small amount of dynamic range is available and that all numbers in the computational problem can be represented accurately with the limited set of decimal places. Additionally, unlike floating point, neither logarithmic nor semi-floating point arithmetic are standardized.

FPGAs have not historically been well-known as effective platforms for floating point applications due to the large hardware requirements, but this has changed in recent years. Studies such as [47] have shown that FPGA floating point performance has increased to the point of surpassing general purpose CPUs, especially since the inclusion of embedded multiplication units. Still, only a small number of floating point FPGA applications have been developed, so there is much room to explore this feature.

### **3.4 Hardware/Software Design**

Hardware/software design is an application design technique where the hardware and software aspects are designed simultaneously to allow a software program to communicate with and use an FPGA to carry out specialized functions. This is a practice that is possible because of the programmable nature of FPGAs. It differs from using plug-in ASICs as part of an application because the hardware can be more specialized, and the software can potentially use it more effectively [39]. Hardware/software design allows a developer to exploit the strengths of both software and hardware. For example, software is good at handling tasks that require complex control structures, and can easily present setup or results in GUI form for ease of use and interpretation, while hardware is able to parallelize computation for faster execution. Hardware/software design is also very well suited to porting an application

from a software-only implementation to a platform which allows hardware acceleration, because a lesser portion of the application will need to be redeveloped, cutting down on development time. In order to be effective, however, the communication avenue between the hardware and software must not present a significant bottleneck.

### 3.5 Nature of QCADesigner Computation

The actual algorithmic structure of QCADesigner bears some striking similarity to the 2D convolution image processing algorithm [37]. A central cell (analogous to a pixel) is acted upon by a sum of weighted products in the form of surrounding cells' polarizations and kink energies (analogous to the pixel values and convolution weights). There were techniques presented in [37] which could be used to make arbitrarily large and efficient convolution grids (analogous to a variable radius of effect) for large images (analogous to circuits), using repeated instances of a relatively small mathematical core. However the mathematics present in the 2D convolution problem are comparatively simple, using integer arithmetic and a fixed convolution grid with 4-bit weights for each pixel, and only 16-bit pixel values. QCADesigner requires 32-bit floating point values for both cell polarizations and kink energies (pixel values and convolution weights), so the hardware would consume far too much area to allow this type of architecture to be implemented on available resources.

The numerics of QCADesigner are more similar to N-body physics simulations such as molecular dynamics and quantum chromodynamics. If a QCA cell is considered like a “particle,” the two problems have many similarities. Both the cell and the particle change their state based on the state and proximity of other surrounding particles, and both problems require high-precision mathematics. On the other hand, QCA is a simpler problem than N-body simulations because the cells (“particles”) do not move, and are generally limited to two dimensions instead of three, therefore making some forms of grid-based memory structures [14] easier to manage.

The floating point FPGA applications in [20], [17], and [32], also have similar numerical demands to QCADesigner, but their architecture and controls are more straightforward. They are implemented using floating point computational pipelines and data streaming techniques, easily adaptable to other mathematical requirements. Although, these applications' mathematics are limited to multiplication and addition, while QCADesigner demands division and square roots as well. Division and square roots are more complex functions than multiplication and addition, so there will be more hardware consumed in each core, leaving less room for multiple cores.

### 3.6 The Wildcard Hardware

After completion of the background research, and after it was determined that hardware/software design was viable for this application, an appropriate implementation platform is selected. An Annapolis Microsystems Wildstar II is used in [13] and [14] to implement a hardware/software application, which is a PCI-based peripheral for use with desktop computers. It possesses two Xilinx Virtex II FPGA chips as computing elements, as well as on-board memory and PCI communication drivers for interaction with software.

A similar upgraded card is available, the Wildcard 4, which has a Xilinx Virtex 4 FPGA and increased on-board memory. Ready-made libraries for both hardware and software development exist which support high data transfer rates over the PCI interface, facilitating the rapid establishment of effective communication between hardware and software.

In addition to the Virtex 4 chip, the Wildcard 4 also has 128 MB of DRAM, 8 MB of SRAM, two clocks (one fixed on the communication bus, the other programmable up to 240 MHz), and PCI communication. Refer to figure 3.4 for a block diagram of the system. Also included are libraries for interacting with the card via C programs and simulating a software driver in VHDL. The PCI interface allows for several types of communication, including register transfer, direct memory access

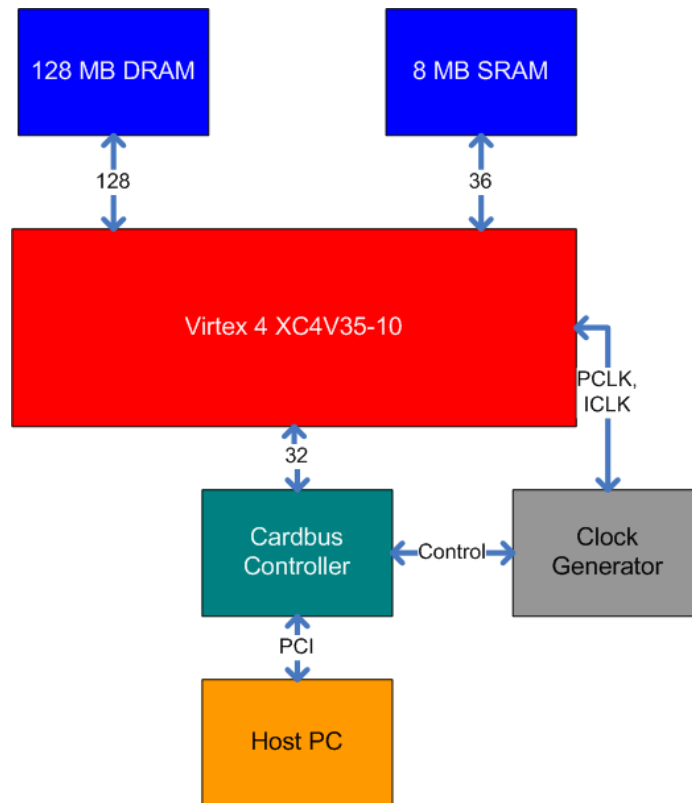


Figure 3.4: Block diagram of Wildcard 4 hardware

(DMA), and interrupt generation.

# Chapter 4

## Proposed Simulation Architecture

The design of a system that could potentially speed up the execution of QCADesigner could take one of many shapes based on a large set of constraints. The functionality of the design is explained in this chapter, as well as some of the related key decisions and software analysis that took place beforehand.

### 4.1 Software Analysis

Prior to committing to any new design, the QCADesigner software is thoroughly analyzed. This analysis includes application profiling using gprof; qualitative analysis of the code by means of examination and debug outputs to determine where most of the computation occurs; further code analysis to determine the memory structures used to store data and assess how best to use the FPGA resources; numerical analysis to track in detail the dynamic range of critical values in the calculation and compare the accuracy of double precision to single precision floating point. The goal here is to identify which section of code should be converted to hardware, and confirm that it has as much parallelizability as expected based on the simulation engine's mathematics.

#### 4.1.1 Code Analysis

The first step in analyzing the QCADesigner source code is done using gprof, which generates an execution profile of the application. This breaks down how much computation time is spent in each function,



and gives other statistics like how many times each function is called, and where it is called from. This is done by adding a `-pg` option to the `gcc` input, running the application, and running the `gprof` tool on the resulting profile data dump. Each such execution profile is generated by opening QCADesigner, loading a circuit, setting up the simulation, running the simulation to completion, then closing QCADesigner. The results of this `gprof` analysis show that for simulations of any appreciable size (*i.e.* more than one second long), the dominant function is `run_bistable_simulation` (see Appendix A for a complete code listing of this function), which takes approximately 90% to 98% of the execution time. This initial analysis shows that there is a high enough concentration of processor time spent executing a single function for QCADesigner to be a promising candidate for hardware acceleration.

Deeper analysis of the code and application execution details reveals greater detail in terms of what specific operations are taking up the most execution time. For a particular run of the CLB simulation, 96.6% of the execution time is spent in the function `run_bistable_simulation`, which includes 141 seconds of “initialization” and 3500 seconds of “simulation.” Initialization takes place in lines 89 through 168 of listing A.1, and most of the time consumed there is actually used by a function called `bistable_refresh_all_Ek`, which calculates all of the kink energies that will be used in the simulation. The simulation takes place in lines 170 through 362 of listing A.1, but the execution time is dominated by a smaller loop of code, lines 268 through 318 in listing A.1 (reproduced below for convenience).

```

1     iteration = 0;
2     stable = FALSE;
3     while (!stable && iteration < max_iterations_per_sample)
4     {
5         iteration++;
6         // — assume that the circuit is stable — //
7         stable = TRUE;
8
9         for (icLayers = 0; icLayers < number_of_cell_layers; icLayers++)
10        {
11 #ifdef REDUCE_DEREF

```

```

12     number_of_cells_in_current_layer = number_of_cells_in_layer [
13         icLayers] ;
14     for (icCellsInLayer = 0 ; icCellsInLayer <
15         number_of_cells_in_current_layer ; icCellsInLayer++)
16 #else
17     for (icCellsInLayer = 0 ; icCellsInLayer <
18         number_of_cells_in_layer [icLayers] ; icCellsInLayer++)
19 #endif
20     {
21         cell = sorted_cells [icLayers] [icCellsInLayer] ;
22
23         if (!(QCAD_CELLINPUT == cell->cell_function) ||
24             (QCAD_CELLFIXED == cell->cell_function))
25         {
26             current_cell_model = ((bistable_model *) cell->cell_model) ;
27             old_polarization = current_cell_model->polarization;
28             polarization_math = 0;
29
30             for (q = 0; q < current_cell_model->number_of_neighbours; q
31                 ++)
32                 polarization_math += (current_cell_model->Ek[q] * ((
33                     bistable_model *) current_cell_model->neighbours [q]->
34                     cell_model)->polarization);
35
36             // math = math / 2 * gamma
37             polarization_math /= (2.0 * sim_data->clock_data [cell->
38                 cell_options.clock].data [j]);
39
40             // — calculate the new cell polarization — //
41             // if math < 0.05 then math/sqrt(1+math^2) ~ math with
42             // error <= 4e-5
43             // if math > 100 then math/sqrt(1+math^2) ~ +-1 with error
44             // <= 5e-5
45             new_polarization =
46                 (polarization_math > 1000.0) ? 1 :
47                 (polarization_math < -1000.0) ? -1 :
48                 (fabs (polarization_math) < 0.001) ?
49                     polarization_math :
50                     polarization_math / sqrt (1 + polarization_math *
51                         polarization_math) ;
52
53             // — set the polarization of this cell — //
54             current_cell_model->polarization = new_polarization;
55
56             // If any cells polarization has changed beyond this
57             // threshold

```

```

46         // then the entire circuit is assumed to have not converged.
47         stable = (fabs (new_polarization - old_polarization) <=
                    tolerance) && stable ;
48     }
49 }
50 }
51 }//WHILE !STABLE

```

Listing 4.1: Main simulation loop

The loop in listing 4.1 iterates through every cell in every layer of the circuit being simulated and calculates a new polarization for that cell, forming the essential basis of the simulation engine. It reiterates until total circuit stability is achieved, or until the maximum number of iterations has elapsed. It executes to its own completion once for every sample in the simulation, thus these 50 lines of code constitute a very large portion of the total execution time of a simulation run by QCADesigner. Thus, listing 4.1 represents a portion of QCADesigner code that would be desirable to be executed in hardware.

Examining the code for data dependencies, which would affect how it could be accelerated, it is clear that each sample of the simulation must be executed sequentially, so there is no possible way to parallelize the computation per sample. It can also be observed that some cells whose polarizations are calculated in later iterations of the loop in listing 4.1 will be affected by cell polarizations that are calculated in preceding iterations. However, this represents a “numerical problem” within the calculation, rather than an actual data dependency, as mentioned in the code’s comments (lines 147-149 of listing A.1). Comparison of the theory of the bistable engine with the actual C code execution indicates that no such data dependency is necessary in the calculation. Its existence in the software is due to the nature of the storage structures and limited resources. Memory resources on the PC are constrained due to operating system overhead and the need for multitasking, preventing the data duplication necessary to avoid the numerical problem altogether. The software attempts to mitigate this numerical problem by randomizing the cell list order (lines 152 to 162 in listing A.1) so that a minimum of influence is bequeathed by the early changes.

The lack of any true data dependencies between cells means the idea presented in Section 2.3 of calculating every cell simultaneously is mathematically plausible. This confirms that great parallelizability exists to be exploited. The limiting factor of the practical implementation would lie with the availability of hardware resources, but given a large enough chip, such an architecture could be realized.

The critical data in the calculation, as can be seen in listing 4.1, are the kink energies, polarizations, and gamma (clock) value associated with each cell. The specifics of how they are stored in memory, and how often they change, influences what the most effective design for retrieving, calculating, and delivering data will be. The communication functions in the wildcard libraries operate much more efficiently with data in contiguous blocks such as arrays, and larger transfers are always more efficient. Kink energies are stored as contiguous arrays associated with each cell, but only contiguous by cell, so high volume transfers are not practical. However, kink energies do not change between iterations or samples in the simulation, so they can be transferred one time at the beginning and kept for the entire execution time. Polarizations are stored by cell, individually, so no bulk transfers can be made without rearranging the data. Additionally, they change for each iteration of the simulation, so they must be updated each time. Gamma values are stored in an array by sample index, but there are only four distinct values that need to be used at any given time. They could be transferred in high volume for all samples, or could use the alternate register transfer method on a per-sample basis.

The total memory this data occupies is also important to determine how it may be used. They are stored in 64 bit double precision floating point variables, and a CLB has about 22,000 cells. Given an effective radius of 500 nm (for reference, RIT's simulations use a radius of 50 nm), the average number of neighbours per cell is about 400, which means 400 kink energies must be stored for each cell. This means that a total of 70 MB of storage is required for kink energies. For polarizations, only one polarization per cell is required, for a total of just 176 kB. For a 300,000 sample simulation as we are running, 9.6 MB would be

required for gamma value storage. These totals would be halved if using single precision floating point. Either way, there is enough memory on-board the FPGA to accommodate all this data locally.

### 4.1.2 Numerical Analysis

An extensive numerical analysis is conducted on QCADesigner to determine what number systems may be acceptable candidates for implementation. The dynamic ranges of the kink energy, polarization, and other intermediate variables are tracked in detail throughout a variety of simulations in order to determine the numerical accuracy demanded by QCADesigner. To begin with, the polarization value of a cell is typically close to 1 or -1, however it can in theory be arbitrarily small. In practice polarizations are not often reported to be any smaller than  $1e-6$ . Kink energies are more extreme. With a large radius of effect, a kink energy of less than  $1e-42$  is not uncommon, and can also come in at  $1e-22$ , representing a minimum of 20 orders of magnitude dynamic range that must be supported by the Kink energy variable. With less extreme radii of effect, kink energies of less than  $1e-39$  become increasingly rare, which is a value readily handled by single-precision floating point.

The first intermediate calculation value that comes up to be analyzed is `polarization_math`, specifically the code shown in listing 4.2. In this code, the `polarization_math` variable is a destination for the sum of all neighbouring cells' polarization  $\times$  kink-energy products.

```

1 for (q = 0; q < current_cell_model->number_of_neighbours; q++)
2   polarization_math += (current_cell_model->Ek[q] * ((bistable_model *)
   current_cell_model->neighbours[q]->cell_model)->polarization);

```

Listing 4.2: Polarization math accumulator

A possible issue with reducing the usable precision of variables here is cumulative rounding errors. Up to 20% of the polarization  $\times$  kink-energy values could be less than  $5e-42$ , which is the floating point half-precision threshold at which point a significant amount of accuracy is

lost. Adding that many rounded values together can create a noticeable rounding error in the final total, depending on what the final total is. This error, once created, could propagate itself through all subsequent iterations and samples of the simulation. However, the nature of this accumulation operation typically involves many higher magnitude values as well. Even in double precision floating point, the least significant digit of a number can have no greater difference from the most significant digit of 16 orders of base ten magnitude. So the net effect of low-magnitude rounding errors will be numerically insignificant if not zero. In fact, analysis shows that the total contribution of low-magnitude values to the total in this accumulation is very rarely more than zero, partly due to having offsetting signs.

```

1 polarization_math /= (2.0 * sim_data->clock_data[cell->cell_options.
   clock].data[j]);
2
3 new_polarization =
4   (polarization_math > 1000.0) ? 1 :
5   (polarization_math < -1000.0) ? -1 :
6   (fabs (polarization_math) < 0.001) ? polarization_math :
7   polarization_math / sqrt(1 + polarization_math * polarization_math);

```

Listing 4.3: Gamma and square root simplification for polarization

Listing 4.3 shows another place in the code where dynamic range consideration is important. Here, the `polarization_math` variable is divided by the clock value, which typically ranges between  $3.8\text{e-}23$  and  $9.8\text{e-}22$ , bringing the magnitude of `polarization_math` back up. Then, to get the final new polarization value,

$$\sqrt{1 + \frac{1}{2\gamma} \sum_j (E_{i,j}^k P_j)} \quad (4.1)$$

is only executed mathematically if the absolute value of  $\frac{1}{2\gamma} \sum_j (E_{i,j}^k P_j)$  is greater than 0.001 and less than 1000. Therefore the squared term has no chance of being less than  $1\text{e-}6$ , and cannot possibly become a bottleneck regarding precision. The simplification in this part of the equation also renders earlier imprecision even more insignificant.

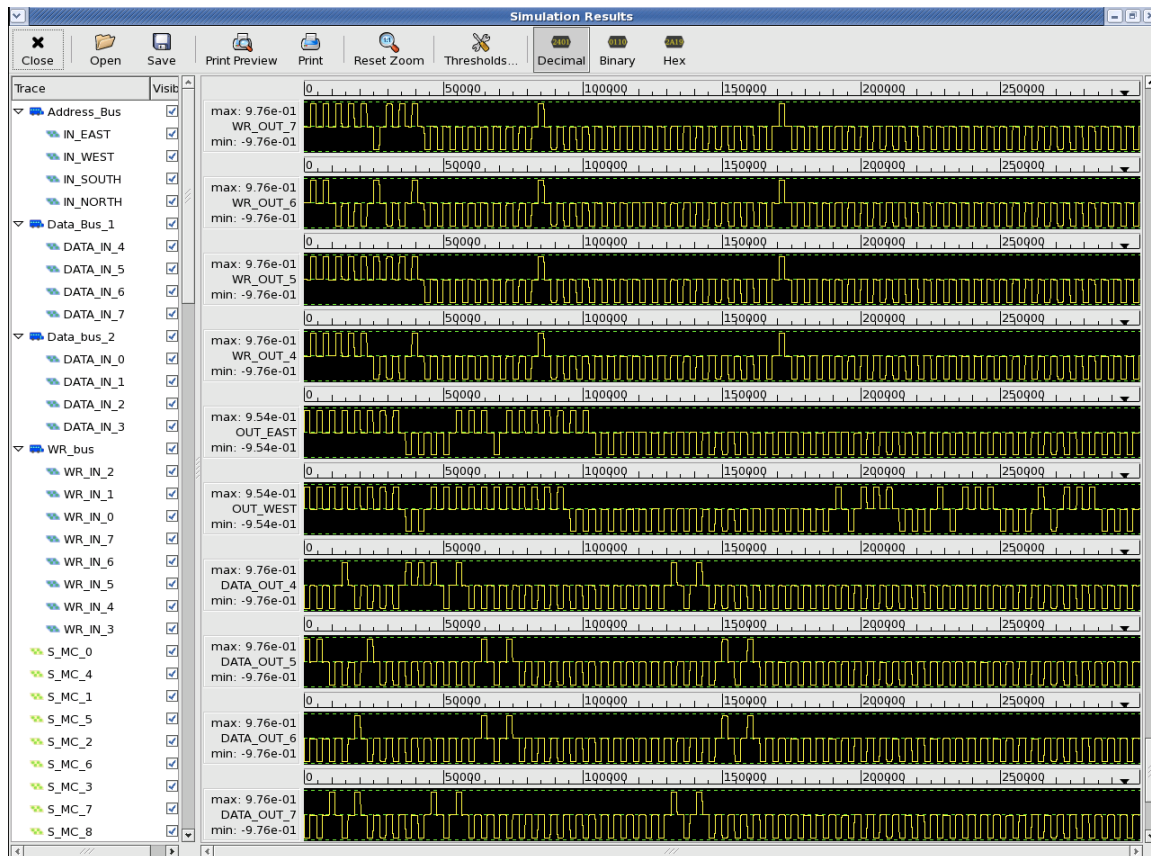


Figure 4.1: QCA Designer simulation using single precision floating point

Considering this analysis, it is clear that floating point is the best number system to use. This application is not suited to use with fixed or semi-floating number systems, because of the large dynamic ranges required, and the fact that a large portion of the operations are additions rules out logarithmic floating point. In addition, floating point arithmetic cores are readily available as IP Cores in the Xilinx IDE, which greatly reduces development time. Considering the data channel for communication with the FPGA is only 32 bits wide, and the fact that double precision floating point cores are amazingly large, single precision floating point is preferable to double precision for this design.

With this in mind, QCA Designer is tested using single precision floating point in software to confirm that simulation accuracy is maintained. By simply changing variable types from `double` to `float` and

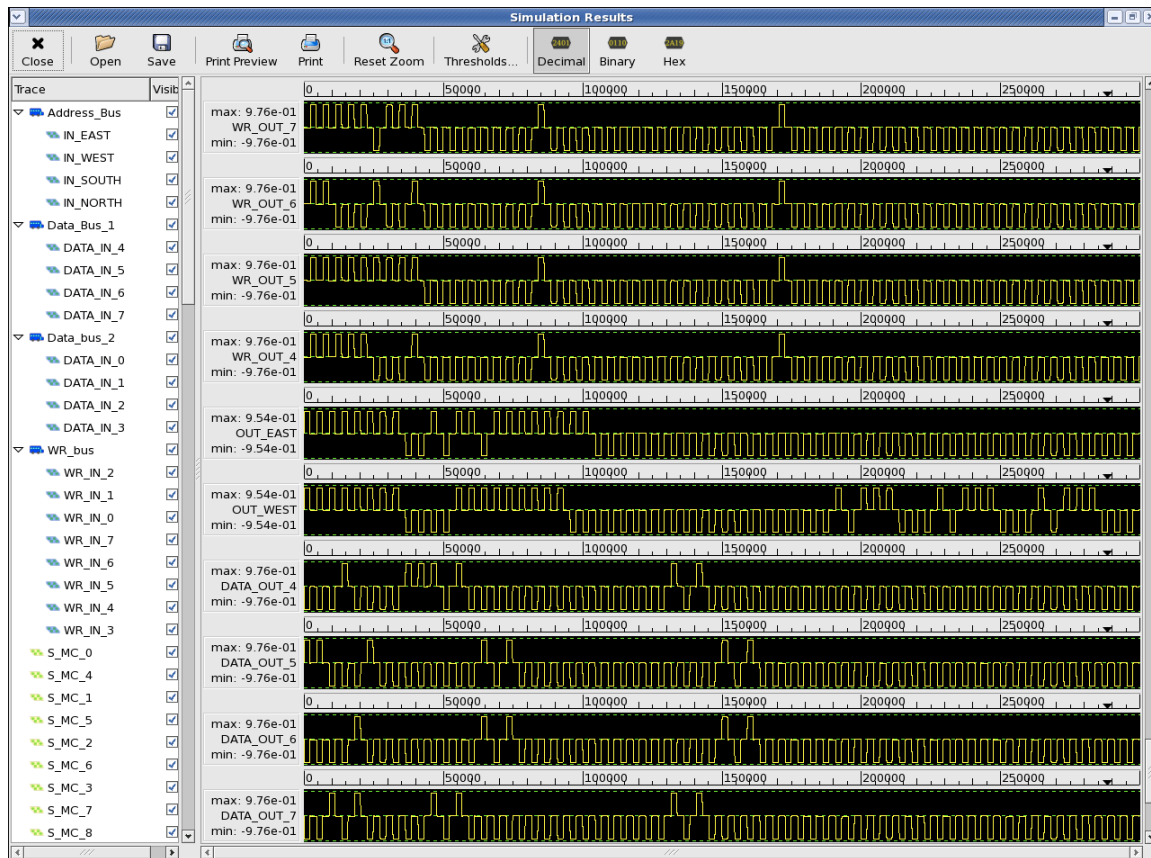


Figure 4.2: QCA Designer simulation using double precision floating point

comparing the simulation results, it is determined that single precision floating point is a safe implementation to use. Figure 4.1 shows the simulation results using single precision, and figure 4.2 shows the simulation results using double precision. It can be seen that the expected results (samples 190000 to 300000 on output OUT\_WEST) are the same and correct for both. There is some minor discrepancy between the implementations for the early portion of the results, but this is inconsequential. Because the latency to the outputs is 16 cycles [45], any values that appear on the outputs before this latency has elapsed are simply noise that has been amplified, so they have no special relevance to the simulation. Closer inspection of the numerical results shows that there is far less than 1% error between double precision and single precision results, even at the end of the simulation where



rounding errors would have a greater chance to accumulate to significant values. Thus single precision floating point is deemed acceptable for this application.

## 4.2 Hardware Architecture

To accelerate the above simulations, several architectures are considered for implementation: including streaming, grid-based parallelism, cellular automata, and instance specific design. A streaming architecture takes a constant stream of data in, processes it, and sends a stream out, achieving application speedup by means of parallelization in a computational pipeline. Grid-based parallelism would break the problem into spatial regions and compute results within each grid, achieving speedup through spatial parallelism. Cellular automata (unrelated to QCA) refers to a very fine-grained form of grid parallelism in which one computational core corresponds to one element that needs to be computed (*i.e.* one core per QCA cell), and these cores are laid out in a large repeating structure. Instance specific design uses a basic starting design and procedurally generates a large design to be implemented on the FPGA, unique to the specific circuit being simulated [16]. Cellular automata would be the ideal implementation to achieve maximum speedup, but the floating point mathematics are too complex to support a large number of cores, so it must be discounted. The Wildcard 4 libraries do not support any means of procedural circuit generation, so instance specific design is not feasible. A streaming architecture, possibly one with multiple pipelines to increase parallelization as much as possible, is decided on. This also allows easy integration with the software.

### 4.2.1 Proof of Concept

Before designing a final general purpose version of the hardware, a small proof of concept version is created, dubbed “Minisim.” It is an instance specific design, simulating a QCA majority gate in one clock

zone using a streaming architecture. It is meant to confirm that a QCA circuit can be simulated accurately using single precision floating point on hardware in a reasonable time while using a reasonable amount of resources, and to gauge how much of the simulation can be put on hardware. Initially, the design attempts to run with a starting point of just cell positions, dynamically calculating the kink energies to be used in the simulation. The additional hardware required to calculate the kink energies proves to be excessive, and because kink energy calculation only contributes a small fraction of execution time, the concept is abandoned.

Minisim's final design involves hard-coded kink energies implying the structure of a majority gate, with the input values and a stream of gamma values coming from software, and a stream of output values going back to the software. This hardware/software execution proceeds to completion in the same perceived amount of time as the software implementation, even for long simulations. The numerical results match software to within 0.001% for steady state values, but deviate from each other more during transitional values. This greater deviation is only in terms of relative error, the absolute error is still quite small, on the order of thousandths. Plots of the hardware and software results of minisim are shown in Figures 4.3 and 4.4 respectively. The synthesis report of Minisim shows that 33% of the board's total available resources are used by the hardware, giving initial indications that three or four computational cores could be included on the FPGA simultaneously. However, these are still rather large cores, so any sort of granular grid computing structure is out of the question.

It was determined earlier in Section 3.6 that the total set of kink energy data for a 22,000 cell circuit with a radius of effect of 500 nm uses about 35 MB of the Wildcard's available 128 MB of DRAM. Using this same radius of effect, circuits with 50,000 or even 80,000 cells can be easily accommodated, or larger if the radius of effect is reduced. And because the kink energies do not change throughout the simulation, they can be loaded once and referenced later. All the kink energies are therefore stored on the Wildcard DRAM rather than the host memory when

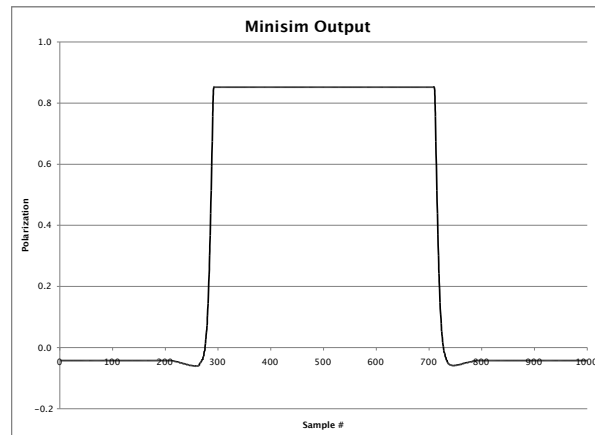


Figure 4.3: Hardware results from minisim test

they are calculated. There is no need for them to be in host memory when all further calculations are done in hardware, so extra transfers are avoided. Furthermore, they are stored in one solid contiguous block with no form of demarkation between cells. Instead, reliance is put upon having a consistent list of cells that does not change its order, and whose neighbour counts are known.

Therefore, cell randomization must be abandoned. This is of no consequence however, because the hardware will calculate an entire iteration without replacing any values in memory, thus completely avoiding the numerical problem the randomization was trying to reduce. Recall from Section 4.1.1 that cell list randomization is done to minimize the effect of the numerical problem by shuffling the order in which cell polarizations are calculated.

### 4.2.2 DRAM Interfacing

Using the DRAM to store kink energies is not completely without issue, however. First of all, the DRAM has a latency of 173 cycles between getting a read request and delivering the requested data, so steps must be taken to ensure this latency is incurred as infrequently as possible. Second, each address refers to a 128 bit word whereas the data being

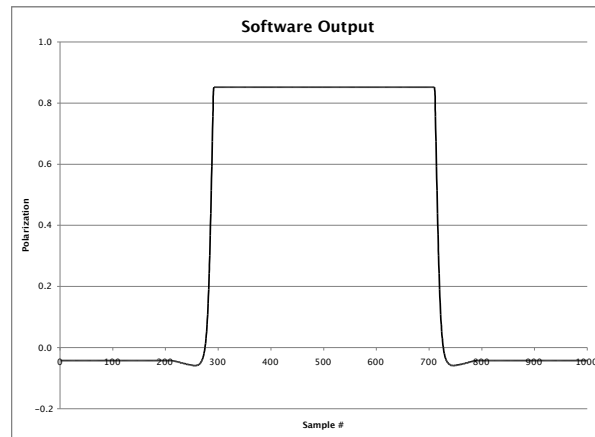


Figure 4.4: Software results from minimim test

stored is only 32 bits wide, so a special interface must be created to write and read data 32 bits at a time.

The DRAM’s “startup latency” of 173 cycles applies when a new set of read operations is started. Additional latency applies when seeking new non-sequential addresses, but after the initial latency has passed, each further sequential address read is delivered on the immediate next cycle. If a constant stream of reads can be sustained, the entire contents of DRAM can be delivered with the only latency cost being the initial 173 cycles. The storage method of the kink energies in the design actually allows for this to take place independent of the radius of effect, because they are stored in the same order as they are processed. However, the mathematical core demands data at a much slower rate than it is delivered from DRAM, so the DRAM output must go through a First-In-First-Out queue (FIFO) which can accumulate data and deliver it when requested. This FIFO and the DRAM are carefully controlled to ensure no over or underflow occur, so that no data is lost or miscalculated. It is 256 words large, and the state machine will only send a read request to the DRAM if less than 50 values are in the FIFO (indicated by `dram_fifo_prog_empty`), allowing the next 173 “cool down” words to be received without overflow. Requesting a new read when the FIFO

data count is back down to 50, allows time for the requested data to get into the FIFO before underflow occurs.

```

1  port map (
2      din => dram_data_in ,
3      rd_clk => i_clock ,
4      rd_en => dram_fifo_rd , --controlled from state machine when data
        is needed
5      rst => dram_fifo_clr , --controlled in the state machine so empty
        flag resets
6      wr_clk => p_clock ,
7      wr_en => dram_data_in_valid ,
8      dout => dram_fifo_out , --goes to DRAM output buffer
9      empty => dram_fifo_empty ,
10     full => dram_fifo_full ,
11     prog_empty => dram_fifo_prog_empty); --set to turn on when less
        than 50 data are present

```

Listing 4.4: DRAM output FIFO

The structures used to manage the data flow into and out of DRAM are shown in their VHDL form in listing 4.5. They refer to `dram_addr(1 downto 0)`, which are two address line bits that are added onto the normal 23 address line bits of the DRAM module. They are used to refer to the 32 bit parts of the 128 bit DRAM word individually. This part of the address line is controlled independently from the other 23 bits so that new DRAM addresses can continue being read at their own pace independent of which value is actually demanded for use on the output. The DRAM output manager is quite simple, with the low address bits essentially controlling a multiplexer between the DRAM's output FIFO and the main math core.

The DRAM data input control is slightly more complex. Since data is written to DRAM in larger portions than it is received, it must be stored in parts to a register whose value is written every fourth address. An override (`dram_super_write`) is built into the logic controlling the write line to allow the final set of data to be written if less than four 32 bit words occur in the last set.

```

1  --DRAM output buffer , goes directly to p_math core
2  Ek_in <= dram_fifo_out(127 downto 96) when (dram_addr(1 downto 0) =
        "11") else

```

```

3     dram_fifo_out(95 downto 64) when (dram_addr(1 downto 0) = "10")
      else
4     dram_fifo_out(63 downto 32) when (dram_addr(1 downto 0) = "01")
      else
5     dram_fifo_out(31 downto 0) when (dram_addr(1 downto 0) = "00")
      else
6     (others => '0');
7
8 --dram input buffer
9 DRAM_input_buffer: process(i_clock , reset)
10 begin
11     if(reset = '1') then
12         dram_data_out <= (others => '0');
13         dram_write <= '0';
14     else
15         if rising_edge(i_clock) then
16             if(dram_buf_write = '1') then
17                 case dram_addr(1 downto 0) is
18                     when "00" =>
19                         dram_data_out(31 downto 0) <= in_fifo_out;
20                     when "01" =>
21                         dram_data_out(63 downto 32) <= in_fifo_out;
22                     when "10" =>
23                         dram_data_out(95 downto 64) <= in_fifo_out;
24                     when "11" =>
25                         dram_data_out(127 downto 96) <= in_fifo_out;
26                     when others =>
27                         dram_data_out <= dram_data_out;
28                     end case;
29
30                     if((dram_addr(1 downto 0) = "11") or (dram_super_write = '1'))
31                         then
32                             dram_write <= '1';
33                         else
34                             dram_write <= '0';
35                         end if;
36                     else
37                         dram_write <= '0';
38                         dram_data_out <= dram_data_out;
39                     end if;
40                 end if;
41 end process;

```

Listing 4.5: DRAM buffer control

### 4.2.3 Other Interfacing Issues

Other interfacing concerns are involved with sending the polarization and gamma data to the FPGA, and reporting back the new polarization results from the FPGA. Methods of efficient transfer must be devised that are able to meet a large range of parameters. Most importantly, the hardware must be able to support any reasonable number of neighbours (effective radius) and all reasonable circuit sizes.

The input into the FPGA is more restrictive than the output from the FPGA, because data comes in faster than it is used, so data overflow is a risk, whereas on the output, data can be written to host memory as soon as it is calculated with less concern of overflow, because DMA writes are more explicitly controlled in hardware. Care must be taken in software and hardware to avoid any read-before-write errors that would corrupt the results of a calculation, while also carrying out necessary data copying on the host in parallel with data computation on the hardware.

Another key communication issue is interfacing between different clock domains. The data transfer bus runs at 40 MHz, while the arithmetic core runs at 140 MHz. Therefore, dual-clock FIFOs are used on both the input and output, which are generated using the IPCore-Gen tool in Xilinx ISE. They are also specifically sized and controlled to prevent over and underflow while meeting the requirements of the application.

The input FIFO is designed to have 1024 locations, which will allow up to 1024 neighbours per cell. This is large enough to accommodate a radius of effect of well over 500 nm, in which no cell in the CLB has more than 650 neighbours. Should it be determined that support for more neighbours is desired, it is easy to create a new, larger FIFO. The DMA buffer on the software side is allocated at a matching fixed size.

For the hardware's output, the FIFO does not need to be as large as there are cells in the circuit, it only needs to be large enough to prevent any over or underflow problems that may arise due to control signals having laggy actuation. 256 locations is deemed an appropriate size for

this, and is tested to confirm functionality. The DMA buffer for the results in software is allocated to be able to receive as many cells as are calculated, which is *not* equal to the total number of cells in the circuit; it only needs to accommodate the cells that are not fixed or inputs, as seen in lines 93 and 94 of listing B.1.

With the architectures in place, the communication algorithm must be created. In order to send kink energy or polarization data to the FPGA without garbage data existing at the end of any block transfers, the FPGA must know how many words to request in the DMA transfer. This is accomplished by sending the number of neighbours to the FPGA using a register transfer, which is slower than DMA for large transactions, but better suited to small ones. For transferring polarizations, the gamma value for the cell to be calculated is also part of the register transaction. At the end of a full phase (*i.e.* all kink energies have been transferred, or all cell data has been transferred for an iteration), a special totem value is sent to the hardware instead of the normal data, so the hardware knows to proceed with the next phase of operation. The control hardware for this is simpler than including cell number counters. The hardware sends interrupts back to the software to indicate when all data has been read or calculated to ensure that no read-before-write errors occur.

Due to the way polarization data is stored (not in solid blocks), each neighbour polarization value must be copied individually to the DMA transfer buffer, as in line 382 of listing B.1. This operation takes the place of where the software would calculate the accumulation of `polarization_math`. The software waits for the interrupt signifying the completion of a cell's calculation *after* this memory transfer is complete to parallelize the workload between host and hardware.

#### 4.2.4 The Arithmetic Core

The main mathematical core is where all of the computation happens. It consists of a pipeline of floating point math cores as shown in Figure 4.5. The floating point cores are all generated using the IPCoreGen tool






Figure 4.5: Block diagram of the polarization math pipeline

in the Xilinx ISE. The VHDL code for this core is shown in Appendix D.

The control signals on each floating point core are for the most part connected in such a way that the next core in the line is automatically started when the data on the previous core is ready. Controlling the multiplier and adder together to form the accumulator in the first stage is slightly more complex, requiring precise timing and counting of inputs. The delay line shown is simply a shift register in place to sync up data propagation through multiple branches of one pipeline. Gamma being multiplied by a constant value of two is accomplished by left shifting the exponent of the floating point value to save on hardware cost. It works without issue because the value in gamma is never large enough where a 1 would be shifted out of the MSB. The circuit's stability is not calculated in this core, or in fact anywhere else on the hardware. Because all polarizations are delivered back to the host platform for each iteration, this comparatively simple operation can be done in software with little extra cost.

The floating point cores are all tuned to have minimum latency, while guaranteeing a certain operating clock speed. Of special importance is the adder's latency in the accumulator, because this is the critical factor of how fast a complete cell can be calculated, and how efficiently the full pipeline may be used. Lower latency for a floating point unit means that more levels of combinational logic must be waited on between stages in the computational pipeline, so more delay is inherent per pipeline stage. Increased latency means the clock may run faster, but more hardware is required to register more stages of the pipeline. So, the least amount of latency is desired while still maintaining a minimum clock rate.

There are two clocks on the Wildcard, one is fixed at 40 MHz and governs data bus transactions. The other runs at a user-specified speed up to 240 MHz, and in this case must be set to at least 130 MHz for the DRAM to function properly. Calculating the total time required to calculate the polarization of one cell with ten neighbours shows that there is little difference between latency configurations, less than 50 ns difference out of a total around 750 ns. A latency configuration is selected which enables an optimum balance between speed and stability at 140 MHz, as indicated by the Xilinx place and route tool.

At this point, the duration of a complete QCA circuit simulation using this core can be calculated. It is based on the simulation of a CLB (22,558 cells) with a radius of effect of 50 nm (average number of neighbours per cell is 5), 300,000 samples, and a convergence tolerance of 0.001 (average reiteration rate per sample of 1.7). Keeping in mind, this calculation assumes maximum pipeline usage and no overhead once the calculation is started, it represents a theoretical minimum time. The combined latency of the accumulator stage is the only real per-cell calculation cost, because a new cell can be accumulated as the previous cell flows down the pipeline. Similarly, the only real per-neighbour cost in the accumulator is in the adder, as the multiplier only contributes its own additional latency once before any addition can start. The accumulator latency is equivalent to the multiplier latency, plus the adder latency for each neighbour, plus one cycle to account for informing the core that a cell is finished with input. This value gets multiplied by the number of cells to get the number of cycles for one iteration, added with the number of cycles for the last part of the pipeline to complete on the last cell, then multiplied by the number of iterations per sample, then multiplied by the number of samples. This total number of cycles for a simulation is divided by the clock speed to get the elapsed time.

$$\frac{((6 + 6 * 5 + 1) * 22558 + 14 + 4 + 5 + 14 + 14) * 1.7 * 300000}{140 \text{ MHz}} = 3040.68 \text{ s} \quad (4.2)$$

This result is on par with the current execution time in software. Although it assumes perfect pipeline usage and no overhead, several cores can fit on the FPGA simultaneously. The execution time could be reduced by a linear factor equivalent to the number of cores present, thus compensating for non-idealities and potentially achieving the desired speedup.

#### 4.2.5 Operation and Synthesis

Hardware operation is selected via a checkbox in the simulation setup window. The block diagram of the system is shown in Figure 4.6, illustrating the basic connections between major components of the design. The combined hardware/software system operates as shown in the flowchart in Figure 4.7.

There is some overhead involved in the frequent DMA transfers, but any such overhead's contributing factor is reduced as the number of neighbours in a circuit increases (*i.e.* as the size of each transfer increases). This hardware/software system is shown to work in simulation using the simplified host software simulator prescribed by the Annapolis tools. The Xilinx place and route report shows that the main arithmetic core uses 1834 slices (3037 LUTs, or 11% of the Virtex 4) of the available hardware resources, and the rest of the control and interface hardware consumes 1809 slices (1608 LUTs, or 12% of the Virtex 4), indicating that up to 6 arithmetic cores may be used simultaneously.

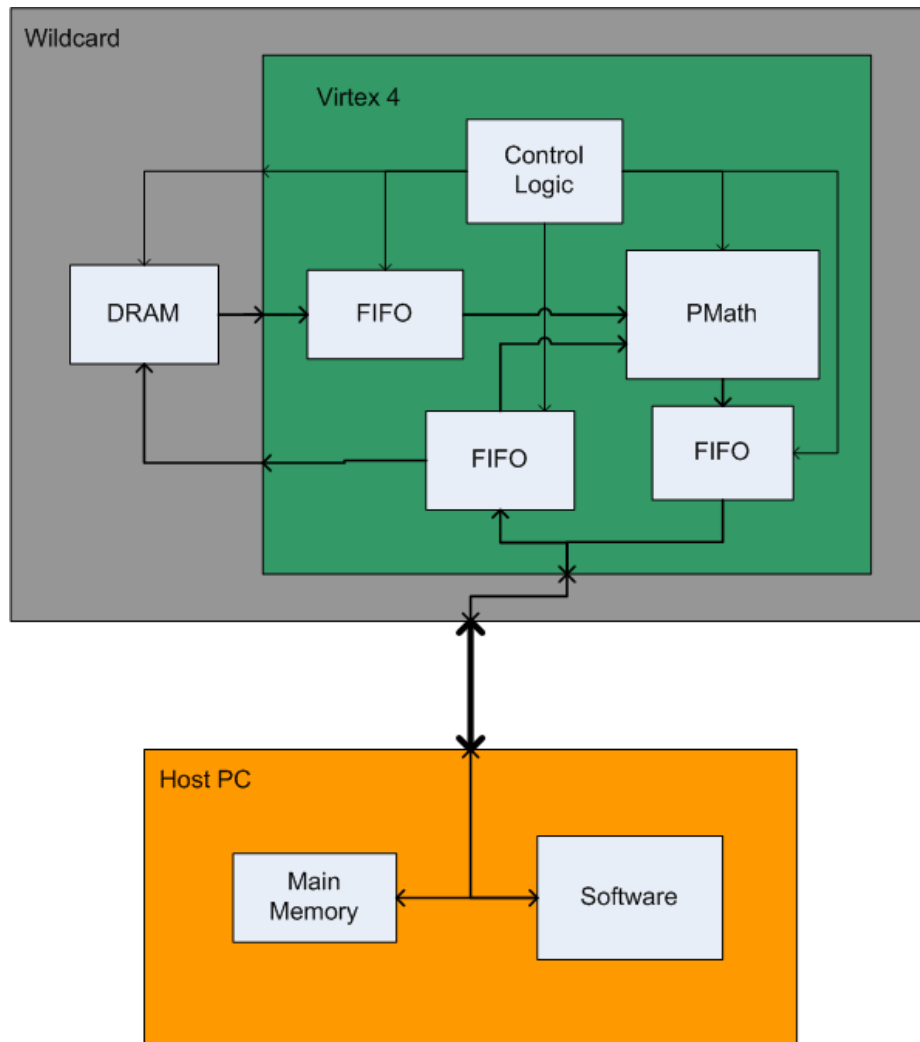


Figure 4.6: Block diagram of high level system architecture

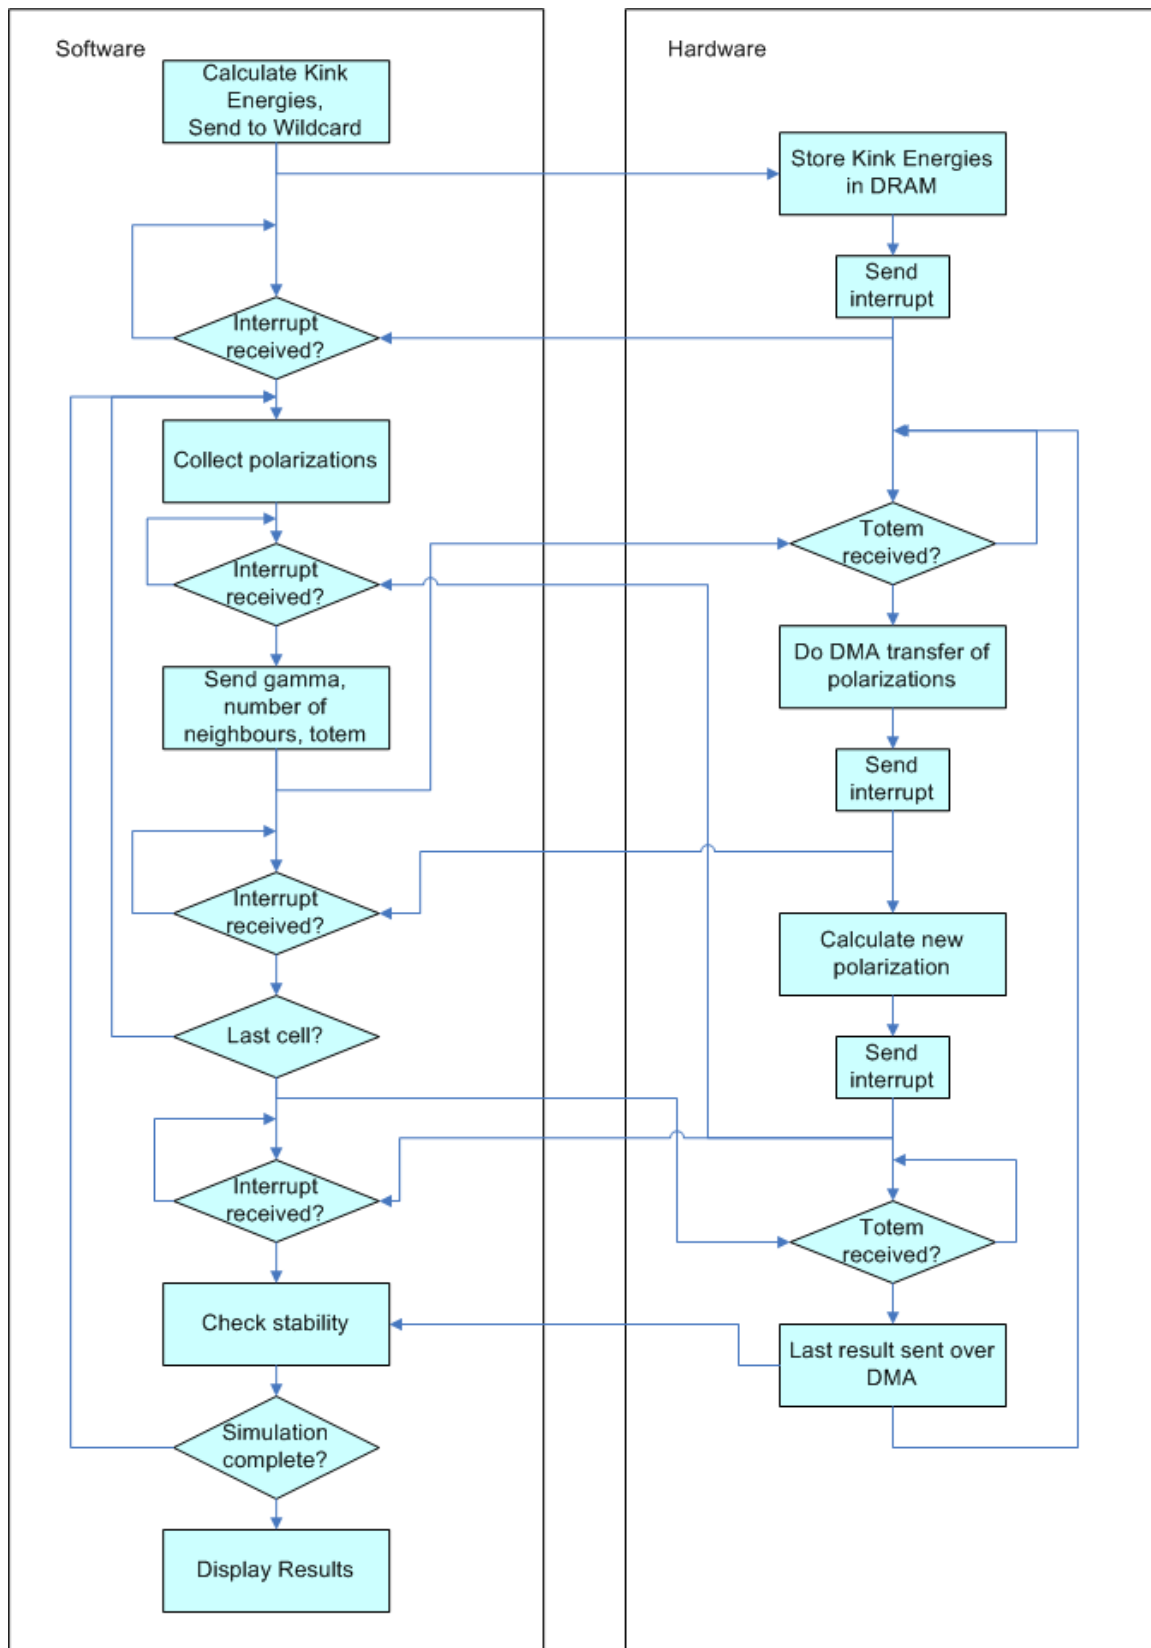


Figure 4.7: Flowchart for the operation of the combined hardware/software system

# Chapter 5

## Results and Discussion

As mentioned in Section 4.2.4, the predicted execution speed of the hardware is expected to be on the same order of magnitude as the software execution speed. However, upon actually using and testing the hardware, the observed execution speed is several orders of magnitude slower. An 80,000 sample simulation of a majority gate with an 8 cell wire on the output takes a reported 1 second to run to completion in software, but 157 seconds in hardware; a very surprising result. Despite this, the hardware is still capable of delivering an accurate result for small circuits. Figure 5.1 shows a simple majority gate simulation running in software and Figure 5.2 shows the same simulation done in hardware, clearly illustrating the same result being obtained in both cases, with some minor numerical discrepancy due to switching the floating point precision level. This result verifies the basic functionality of the floating point pipeline developed. This and the analysis showing its potential speed are encouraging results in terms of a proof of concept design on which to base other designs. Additionally, some notable improvements have been made to the QCADesigner software during the course of this work.

### 5.1 Cause of Slowdown

The first effort to try to determine why the hardware simulation takes so long is to apply `gprof` and see if any of the new functions being applied in the main simulation loop such as `IntWait` are taking an excessive amount of time, which would indicate the hardware rather

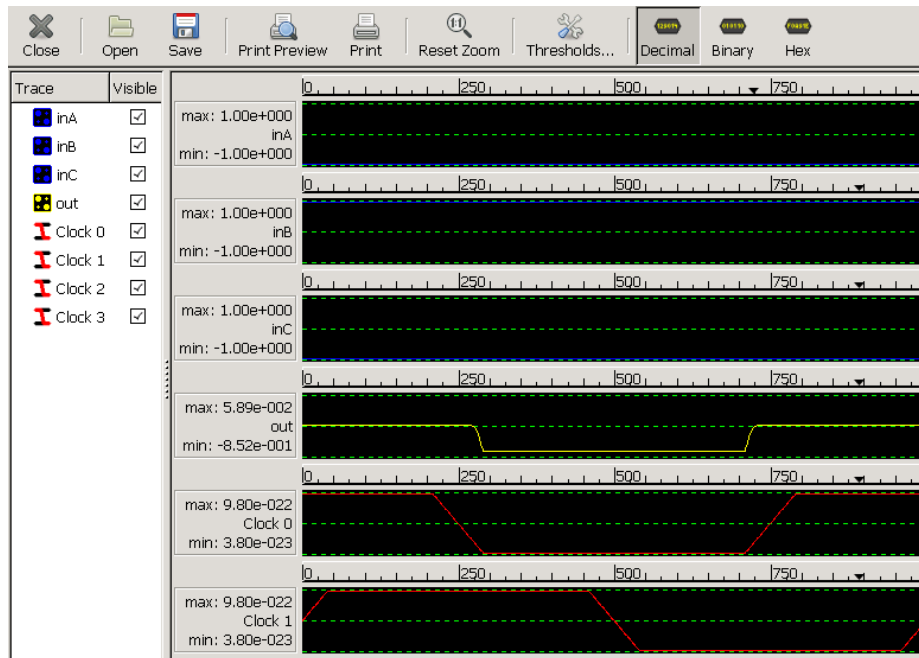


Figure 5.1: Result for software run of a majority gate

than the software is taking too long to finish the computation. This doesn't yield any useful insight. The next step is to manually instrument the code with the functions `QueryPerformanceFrequency` and `QueryPerformanceCounter`. `QueryPerformanceFrequency` gives the number of processor ticks that occur every second (*i.e.* the clock speed), and `QueryPerformanceCounter` gives the number of processor ticks that have occurred since startup. Using these functions, it can be determined precisely how much time is consumed executing certain sections of the code.

To begin with, it is determined that one cell in software takes roughly 600 ns to compute, with some variance depending on the specific run. In hardware mode, the calculation time for a single cell is between 18 and 45  $\mu$ s. Upon further analysis, this rather alarming increase in execution time is determined to be caused by the interrupt reset procedures, of which there are two, taking 10 to 20  $\mu$ s each (usually about 14  $\mu$ s). The repeated interrupt reset, shown in listing 5.1, is a required procedure given the architecture in use, because every time an interrupt is used, it must be reset before being used again. The only way to avoid using

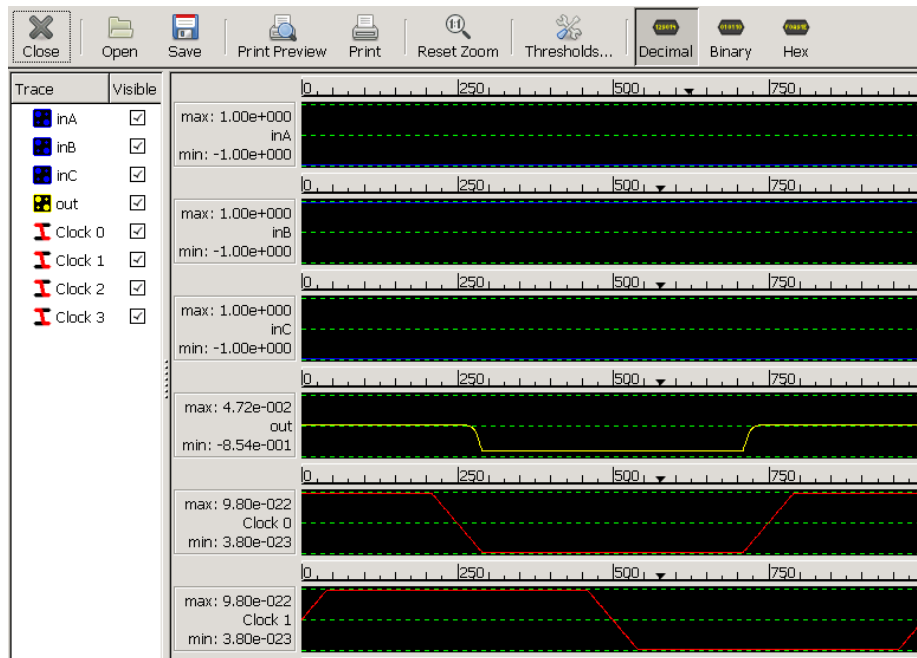


Figure 5.2: Result for a hardware run of a majority gate

interrupts with this architecture would be to wait a fixed amount of time for hardware operations to complete, but there is no good way to estimate the completion time dynamically or wait for such brief precise periods. There also is no way for the software to detect the logical value assigned to the interrupt bit by the hardware, because the value is latched at 1 when triggered until reset by the software.

```

1 rc=WC.IntEnable( TestInfo . DeviceNum , FALSE ); //reset interrupt
2 rc=WC.IntReset( TestInfo . DeviceNum );
3 rc=WC.IntEnable( TestInfo . DeviceNum , TRUE );

```

Listing 5.1: Interrupt reset procedure

Due to the large variance in these measurements, it is impossible to say exactly how long the calculation would take without these interrupt resets. An estimate based on the difference between total execution time and total interrupt reset time might come between 1 and 5  $\mu$ s, low enough to achieve speedup with six pipelines. To be clear, it is just the interrupt reset that takes that long; the interrupt wait function itself takes a much shorter amount of time, indicating that the hardware calculation is proceeding apace. This is consistent with the



determination in Section 4.2.4 that the main arithmetic core should be able to calculate the results at the same rate as the software. If better interfacing and control hardware can be developed, speedup may be achieved when using multiple cores in parallel.

A small test program is created to test in simpler terms what difference the interrupt reset has on execution time. Two separate data transfers are initiated in sequence, with an interrupt between them, and another interrupt at the end, requiring a reset in between. With the reset in place, the program takes 40  $\mu\text{s}$  to execute, but simply waiting the short transfer out without the reset results in a 1.5  $\mu\text{s}$ . There was no indication in any manual or simulation that interrupt resets would take so long.

## 5.2 Hardware Reliability

In addition to the extra long execution time, there are significant reliability problems with the control hardware. Any circuit much larger than 6 cells encounters significant difficulty in proceeding to the end of simulation. One problem appears to be related to bad data being reported back to the host, which could be due to problems on either the input or output sides of the hardware. This causes the simulation to constantly reiterate until it hits the reiteration limit, and results in a flatlined simulation result. Another problem is related to the hardware and software getting out of sync, causing interrupts to timeout, again resulting in flatlined results. The inconsistent nature of these issues suggests a problem related to timing, possibly a clock running too fast. However, reconfiguring the hardware to run on the 40 MHz clock rather than 140 MHz does not inspire any changes.

Further debugging is quite difficult. Because the programming interface is on the PCI bus rather than a JTAG interface, it is impossible to use the Xilinx ChipScope hardware analyzer to examine signal timelines. The closed form factor of the card also prevents examining any signals using an oscilloscope or logic analyzer. The simulation works, so that is not a debugging avenue to pursue. The only way to debug the

hardware is to request debug outputs from the hardware. But the times these data can be requested without interfering with the standard operation of the hardware are very limited, making it nigh impossible to get comprehensive debugging information. Because of the great difficulty in debugging and the timing analysis showing a very low performance ceiling for the given architecture, further development is abandoned.

### 5.3 QCADesigner Enhancements

The results of this work are not a complete write-off, however. Some notable improvements to the QCADesigner software have been made. First, an error was found in the stability checking method. The version in current distribution has the following for a stability check.

```

1 // If any cells polarization has changed beyond this threshold
2 // then the entire circuit is assumed to have not converged.
3 stable = (fabs (new_polarization - old_polarization) <= tolerance) ;

```

Listing 5.2: Original stability check

The comment suggests that the sample should reiterate if *any* cell's polarization changes goes beyond the threshold, but the way the code is written, it only reiterates if the *last* cell's polarization changes beyond the threshold limit. The fix is very simple, changing line three of listing 5.2 to listing 5.3, thus forcing any unstable cell to cause all future stability checks to fail. The fix causes simulations to take roughly four times as long, as expected, because samples reiterate more often. The fix is also confirmed as needed by Konrad Walus, the QCADesigner lead.

```

1 stable = (fabs (new_polarization - old_polarization) <= tolerance) &&
  stable ;

```

Listing 5.3: Fixed stability check

The other improvement made to QCADesigner is an optimization made to the `select_cells_in_radius` function. In the current version of QCADesigner, a cell's neighbours are found by using a complex conditional statement shown in listing 5.4

```

1  if (sqrt ((QCAD_DESIGN_OBJECT(sorted_cells[i][j])->x -
           QCAD_DESIGN_OBJECT(cell)->x) *
2  (QCAD_DESIGN_OBJECT(sorted_cells[i][j])->x - QCAD_DESIGN_OBJECT(cell)
           ->x) +
3  (QCAD_DESIGN_OBJECT(sorted_cells[i][j])->y - QCAD_DESIGN_OBJECT(cell)
           ->y) *
4  (QCAD_DESIGN_OBJECT(sorted_cells[i][j])->y - QCAD_DESIGN_OBJECT(cell)
           ->y) +
5  ((double) ABS(the_cells_layer-i)*layer_separation) *
6  ((double) ABS(the_cells_layer-i)*layer_separation)) < world_radius)
7  number_of_selected_cells++;

```

Listing 5.4: select\_cells\_in\_radius original

This condition contains within it several arithmetic operations, including a square root function, which take time to execute. These operations are computing a distance between cells, which takes time and does not need to be precisely calculated each time to achieve the desired result. A precondition (shown in listing 5.5) can be added to test just the horizontal and vertical components of the distance, and if either is greater than the radius of effect, the detailed distance calculation can be avoided.

```

1  if( ((QCAD_DESIGN_OBJECT(sorted_cells[i][j])->x - QCAD_DESIGN_OBJECT(
           cell)->x) < world_radius) && ((QCAD_DESIGN_OBJECT(sorted_cells[i][j]
           )->y - QCAD_DESIGN_OBJECT(cell)->y) < world_radius) )

```

Listing 5.5: select\_cells\_in\_radius optimization

This modification reduces initialization time for large circuits by 37%.

## 5.4 Design Alternatives

Several options exist for pursuing alternate solutions to this design problem. A variety of other platforms could be used which are larger and better equipped, or different architectures could be explored.

### 5.4.1 New Architecture

One possible design architecture that could be implemented for future research would store all cells' polarizations on the Wildcard's SRAM

and compute the entire length of the simulation in hardware, rather than constantly exchanging polarization data and using repeated interrupts. This would eliminate almost all communication overhead, including most importantly the interrupt resets.

The storage structure of this new architecture would store all cell polarizations in the form of a spatial grid in SRAM, where every 36 bit word in SRAM will store the polarization and clock zone (34 bits of data) for the cell at a particular location. The cell's location would directly correspond to the address in SRAM, and zeros would be in SRAM for locations where there is no cell. The location encoding would work by each sequential address representing the cell with its location at the next 10 nm increment. QCADesigner can only place cells at 10 nm increments (10 nm is half of a cell's length), so no more location precision than this is needed. A new row is simply defined by counting off the number of cells in the previous row. Finding and retrieving neighbour polarization data would be accomplished by simple rectangular address arithmetic; a given radius of effect would be converted to an integer number of half-cell distances, and cells within a square rather than a circle would be counted to simplify the mathematics a little cost to accuracy. There is enough SRAM on the Wildcard 4 to hold a circuit  $4500 \text{ nm} \times 4500 \text{ nm}$ , which is enough for 4 CLBs. Most of the data in SRAM would be zeros, so larger circuits could be stored if some form of compression can be developed.

Kink energies would be stored in much the same way, but the delivery of gamma values would have to be modified to be delivered in bulk rather than per cell, and results would still be sent back to the host as they are calculated using DMA. Extra software development is required to sort the cells into the form which they will be stored in, and to change the calculation of kink energies to consider the new square of effect.

### 5.4.2 Other Platforms

Other hardware platforms may also be used to great effect. For example, a top of the line Virtex 6 card has over 85,000 logic slices, each

containing four 6-input LUTs [53]. The Virtex 4 used in this thesis only has two 4-input LUTs for each of just 15,000 slices. A Virtex 6 also has PCI Express integration, which would facilitate communication with a host PC, as well as a faster clock speed. With the much larger amount of resources, a design based on a Virtex 6 may be able to support up to 100 floating point pipelines in parallel, moving one step closer to being able to implement a cellular automata style computational architecture.

Many hardware platforms also offer on-board processors, or have the ability to generate soft-core processors on the FPGA. This type of hardware may allow the creation of a hardware/software design with much more efficient communication between the hardware and software kernels.

## Chapter 6

# Conclusion and Future Work

The work done in this thesis has shown the potential for a hardware coprocessor to accelerate the simulations needed in QCA research. Although the system designed here did not succeed in speeding up the execution of QCA simulations, this has more to do with interrupt resets taking much longer than expected than any other aspect of the design. It is possible to achieve the desired speedup using the arithmetic core presented here if data can be fed to it fast enough. Avoiding the long interrupt resets, either by avoiding interrupts altogether or moving to a different hardware which resets them more quickly, could help provide the necessary throughput. An implementation of this application that achieves speedup would be a significant achievement for FPGAs, showing that floating point applications may be sped up on easily available reconfigurable hardware using hardware/software design. This would be attractive to the realm of scientific research because it combines floating point accuracy with the relative ease of use and development of hardware/software design.

The implementation of this alternate architecture is beyond the scope of this thesis, but is a prime candidate for future research. Alternate hardware platforms may also be sought that have better interrupt routines or larger FPGAs, and the QCADesigner coherence vector simulation engine may also be accelerated. The vein of this research may also be applied to broader applications such as the physical side of QCA research, advancing more than just QCA architecture development.

# Bibliography

- [1] Wolfgang Arden, Patrick Coge, Mart Graef, Reinhard Mahnkopf, Hidemi Ishiuchi, Toshihiko Osada, JooTae Moon, JaeSung Roh, Carlos H Diaz, Burn Lin, Pushkar Apte, Bob Doering, and Paolo Gargini. 2009 international technology roadmap for semiconductors. Technical report, IEEE Computer Society, 2009.
- [2] Gary H. Bernstein. Quantum-dot cellular automata: computing by field polarization. In *DAC '03: Proceedings of the 40th annual Design Automation Conference*, pages 268–273. ACM, 2003.
- [3] E. P. Blair and C. S. Lent. Quantum-dot cellular automata: an architecture for molecular computing. In *Simulation of Semiconductor Processes and Devices, 2003. SISPAD 2003. International Conference on*, pages 14–18, 2003.
- [4] Enrique Blair, Eric Yost, and Craig Lent. Power dissipation in clocking wires for clocked molecular quantum-dot cellular automata. *Journal of Computational Electronics*, 9(1):49–55, March 2010.
- [5] B. Bosi, G. Bois, and Y. Savaria. Reconfigurable pipelined 2-D convolvers for fast digital signal processing. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 7(3):299–308, September 1999.

- [6] Owen Callanan, Andy Nisbet, Emre Ozer, James Sexton, and David Gregg. FPGA implementation of a lattice quantum chromodynamics algorithm using logarithmic arithmetic. In *IPDPS '05: Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05) - Workshop 3*, page 146.2, Washington, DC, USA, 2005. IEEE Computer Society.
- [7] R. P. Cowburn and M. E. Welland. Room temperature magnetic quantum cellular automata. *Science*, 287(5457):1466–1468, 2000.
- [8] Mike Cowlinshaw, editor. *754-2008 IEEE Standard for Floating-Point Arithmetic*. IEEE Computer Society, 445 Hoes Lane, P, 2008.
- [9] M. Crocker, Xiaobo Sharon Hu, M. Niemier, Minjun Yan, and G. Bernstein. PLAs in quantum-dot cellular automata. *IEEE Trans. Nanotechnol.*, 7(3):376–386, May 2008.
- [10] Yong Dou, S. Vassiliadis, G. K. Kuzmanov, and G. N. Gaydadjiev. 64-bit floating-point FPGA matrix multiplication. In *Proceedings of the 2005 ACM/SIGDA 13th international symposium on Field-programmable gate arrays*, FPGA '05, pages 86–95, New York, NY, USA, 2005. ACM.
- [11] E. N. Ganesh, Lal Kishore, and M. J. S. Rangachar. Study of complex gate structures in quantum cellular automata technology for FPGA applications. In *Proceedings of the IET-UK International Conference on Information and Communication Technology in Electrical Sciences (ICTES 2007)*, pages 789–794, December 2007.
- [12] Akila Gothandaraman, Gregory D. Peterson, G. Lee Warren,



- Robert J. Hinde, and Robert J. Harrison. FPGA acceleration of a quantum monte carlo application. *Parallel Computing*, 34(4-5):278–291, 2008.
- [13] Yongfeng Gu, Tom Van Court, and Martin C. Herbordt. Improved interpolation and system integration for FPGA-based molecular dynamics simulations. In *Proceedings of the 2006 International Conference on Field Programmable Logic and Applications (FPL)*, pages 1–8. IEEE, August 2006.
- [14] Yongfeng Gu and Martin C. Herbordt. FPGA-based multi-grid computation for molecular dynamics simulations. In *FCCM '07: Proceedings of the 15th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, pages 117–126, Washington, DC, USA, 2007. IEEE Computer Society.
- [15] Yongfeng Gu, Tom VanCourt, and Martin C. Herbordt. Explicit design of FPGA-based coprocessors for short-range force computations in molecular dynamics simulations. *Parallel Comput.*, 34(4-5):261–277, 2008.
- [16] Scott Hauck and André DeHon, editors. *Reconfigurable Computing. Systems on Silicon*. Morgan Kaufman, 30 Corporate Drive, Suite 400, Burlington, MA 01803, 2008.
- [17] K. S. Hemmert and K. D. Underwood. An Analysis of the Double-Precision Floating-Point FFT on FPGAs. In *FCCM '05: Proceedings of the 13th Annual IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM'05)*, pages 171–180, Washington, DC, USA, 2005. IEEE Computer Society.
- [18] Jing Huang, Mariam Momenzadeh, and Fabrizio Lombardi. Defect

- tolerance of qca tiles. In *DATE '06: Proceedings of the conference on Design, automation and test in Europe*, pages 774–779, 3001 Leuven, Belgium, Belgium, 2006. European Design and Automation Association.
- [19] Brad Hutchings, Peter Bellows, Joseph Hawkins, Scott Hemmert, Brent Nelson, and Mike Rytting. A CAD suite for high-performance FPGA design. In *FCCM '99: Proceedings of the Seventh Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, page 12, Washington, DC, USA, 1999. IEEE Computer Society.
- [20] Nachiket Kapre and André DeHon. Accelerating SPICE Model-Evaluation using FPGAs. *Field-Programmable Custom Computing Machines, Annual IEEE Symposium on*, 0:37–44, 2009.
- [21] V. Kindratenko and D. Pointer. A case study in porting a production scientific supercomputing application to a reconfigurable computer. In *FCCM '06: Proceedings of the 14th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, pages 13–22, Washington, DC, USA, April 2006. IEEE Computer Society.
- [22] R. K. Kummamuru, A. O. Orlov, R. Ramasubramaniam, C. S. Lent, G. H. Bernstein, and G. L. Snider. Operation of a quantum-dot cellular automata (qca) shift register and analysis of errors. *IEEE Transactions on Electron Devices*, 50(9):1906–1913, August 2003.
- [23] Timothy D. Lantz and Eric R. Peskin. A QCA implementation of a configurable logic block for an FPGA. In *Proceedings of the*

*Third International Conference on Reconfigurable Computing and FPGAs (ReConFig 2006)*, pages 132–141, September 2006.

- [24] C. S. Lent, P. D. Tougaw, W. Porod, and G. H. Bernstein. Quantum cellular automata. *Nanotechnology*, 4(1):49–57, January 1993.
- [25] Craig S. Lent, Sarah E. Frost, and Peter M. Kogge. Reversible computation with quantum-dot cellular automata (qca). In *CF '05: Proceedings of the 2nd conference on Computing frontiers*, page 403, New York, NY, USA, 2005. ACM.
- [26] Craig S. Lent, Mo Liu, and Yuhui Lu. Bennett clocking of quantum-dot cellular automata and the limits to binary logic scaling. *Nanotechnology*, 17(16):4240+, August 2006.
- [27] Antonio Lopes and George Constantinides. A High Throughput FPGA-Based Floating Point Conjugate Gradient Implementation. In Roger Woods, Katherine Compton, Christos Bouganis, and Pedro Diniz, editors, *Reconfigurable Computing: Architectures, Tools and Applications*, volume 4943 of *Lecture Notes in Computer Science*, chapter 10, pages 75–86. Springer Berlin / Heidelberg, Berlin, Heidelberg, 2008.
- [28] Yuhui Lu, Mo Liu, and Craig Lent. Molecular quantum-dot cellular automata: From molecular structure to circuit dynamics. *Journal of Applied Physics*, 102(3):034311, 2007.
- [29] Josh Model and Martin C. Herbordt. Discrete event simulation of molecular dynamics with configurable logic. In Koen Bertels,

- Walid A. Najjar, Arjan J. van Genderen, and Stamatis Vassiliadis, editors, *FPL 2007, International Conference on Field Programmable Logic and Applications, Amsterdam, The Netherlands, 27-29 August 2007*, pages 151–158. IEEE, 2007.
- [30] Tamer S. Mohamed and Wael Badawy. Integrated Hardware-Software Platform for Image Processing Applications. *System-on-Chip for Real-Time Applications, International Workshop on*, 0:145–148, 2004.
- [31] Gordon E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8):114 ff., April 1965.
- [32] G. R. Morris and V. K. Prasanna. An FPGA-Based Floating-Point Jacobi Iterative Solver. In *Parallel Architectures, Algorithms and Networks, 2005. ISPAN 2005. Proceedings. 8th International Symposium on*, pages 420–427, 2005.
- [33] Michael Thaddeus Niemier, Arun Francis Rodrigues, and Peter M. Kogge. A potentially implementable FPGA for quantum dot cellular automata. In *Proceedings of the First Workshop on Non-Silicon Computation (NSC-1)*, pages 38–45, 2002.
- [34] A. O. Orlov, I. Amlani, G. Toth, C. S. Lent, G. H. Bernstein, and G. L. Snider. Experimental demonstration of a binary wire for quantum-dot cellular automata. *Applied Physics Letters*, 74(19):2875–2877, 1999.
- [35] Alexei O. Orlov, Ravi K. Kumamuru, Rajagopal Ramasubramaniam, Geza Toth, Craig S. Lent, Gary H. Bernstein, and Gregory L.

- Snider. Experimental demonstration of a latch in clocked quantum-dot cellular automata. *Applied Physics Letters*, 78(11):1625–1627, 2001.
- [36] Marco Ottavi, Luca Schiano, Fabrizio Lombardi, and Douglas Tougaw. HDLQ: A HDL environment for QCA design. *J. Emerg. Technol. Comput. Syst.*, 2(4):243–261, 2006.
- [37] S. Perri, M. Lanuzza, P. Corsonello, and G. Cocorullo. A high-performance fully reconfigurable FPGA-based 2D convolution processor. *Microprocessors and Microsystems*, 29(8-9):381–391, November 2005.
- [38] Anders K. Roger and Roger D. Hersch. Designing a Halftoning Coprocessor. In *8th Eurographics Workshop on Graphics Hardware*. Swiss Federal Institute of Technology, September 1993.
- [39] R. Scrofano, M. Gokhale, F. Trouw, and V. K. Prasanna. Hardware/software approach to molecular dynamics on reconfigurable computers. In *FCCM '06: Proceedings of the 14th Annual IEEE Symposium on Field-Programmable Custom Computing Machines*, volume 0, pages 23–34, Los Alamitos, CA, USA, April 2006. IEEE Computer Society.
- [40] Ronald Scrofano, Maya B. Gokhale, Frans Trouw, and Viktor K. Prasanna. Accelerating molecular dynamics simulations with reconfigurable computers. *IEEE Trans. Parallel Distrib. Syst.*, 19(6):764–778, 2008.
- [41] B. Sen, M. Dalui, and B. K. Sikdar. Introducing universal qca logic

- gate for synthesizing symmetric functions with minimum wire-crossings. In *ICWET '10: Proceedings of the International Conference and Workshop on Emerging Trends in Technology*, pages 828–833, New York, NY, USA, 2010. ACM.
- [42] G. L. Snider, A. O. Orlov, I. Amlani, X. Zuo, G. H. Bernstein, C. S. Lent, J. L. Merz, and W. Porod. Quantum-dot cellular automata. In *Papers from the 45th National Symposium of the American Vacuum Society*, volume 17, pages 1394–1398. AVS, 1999.
- [43] Yong Tang, Alexei O. Orlov, Gregory L. Snider, and Patrick J. Fay. Radio frequency operation of clocked quantum-dot cellular automata latch. *Applied Physics Letters*, 95(19):193109+, 2009.
- [44] Himanshu Thapliyal and Nagarajan Ranganathan. Conservative qca gate (cqca) for designing concurrently testable molecular qca circuits. In *VLSID '09: Proceedings of the 2009 22nd International Conference on VLSI Design*, pages 511–516, Washington, DC, USA, 2009. IEEE Computer Society.
- [45] Chia-Ching Tung. Implementation of multi-clb designs using quantum-dot cellular automata. Master’s thesis, Rochester Institute of Technology, Rochester, NY, February 2010.
- [46] Chia-Ching Tung, Ruchi B. Rungta, and Eric R. Peskin. Simulation of a QCA-based CLB and a multi-CLB application. In *Proceedings of the 2009 International Conference on Field-Programmable Technology (FPT'09)*, pages 62–69. IEEE, December 2009.
- [47] Keith Underwood. Fpgas vs. cpus: trends in peak floating-point performance. In *FPGA '04: Proceedings of the 2004 ACM/SIGDA*

*12th international symposium on Field programmable gate arrays*, pages 171–180, New York, NY, USA, 2004. ACM.

- [48] K. Walus. QCADesigner — Microsystems and Nanotechnology Group (MiNa). <http://www.mina.ubc.ca/qcadesigner>, June 2009.
- [49] K. Walus, T. J. Dysart, G. A. Jullien, and R. A. Budiman. QCADesigner: a rapid design and simulation tool for quantum-dot cellular automata. *IEEE Trans. Nanotechnol.*, 3(1):26–31, March 2004.
- [50] K. Walus, G. Schulhof, Mazur, and G. A. Jullien. Simple 4-bit processor based on quantum-dot cellular automata (QCA). In *ASAP '05: Proceedings of the 2005 IEEE International Conference on Application-Specific Systems, Architecture Processors*, pages 288–293, Washington, DC, USA, 2005. IEEE Computer Society.
- [51] Yuliang Wang and M. Lieberman. Thermodynamic behavior of molecular-scale quantum-dot cellular automata (QCA) wires and logic devices. *IEEE Trans. Nanotechnol.*, 3(3):368–376, September 2004.
- [52] Tongquan Wei, Kaijie Wu, Ramesh Karri, and Alex Orailoglu. Fault tolerant quantum cellular array (qca) design using triple modular redundancy with shifted operands. In *ASP-DAC '05: Proceedings of the 2005 Asia and South Pacific Design Automation Conference*, pages 1192–1195, New York, NY, USA, 2005. ACM.
- [53] Xilinx, Inc., 2100 Logic Drive, San Jose, CA, 95124. *Virtex-6 Family Overview*.

# Appendix A

## run\_bistable\_simulation C Code Listing

The following is the C code for the run\_bistable\_simulation function from QCADesigner.

```

1 //-----//
2 // — this is the main simulation procedure — //
3 //-----//
4 simulation_data *run_bistable_simulation (int SIMULATION_TYPE, DESIGN *
      design, bistable_OP *options, VectorTable *pvt)
5 {
6   int i, j, k, l, total_cells = 0 ;
7   int icLayers, icCellsInLayer;
8   time_t start_time, end_time;
9   simulation_data *sim_data = NULL ;
10  // optimization variables //
11  int number_of_cell_layers = 0, *number_of_cells_in_layer = NULL ;
12  QCADCell ***sorted_cells = NULL ;
13  double clock_shift = (options->clock_high + options->clock_low)/2 +
      options->clock_shift;
14  double clock_prefactor = (options->clock_high - options->clock_low) *
      options->clock_amplitude_factor;
15  double four_pi_over_number_samples = 4.0 * PI / (double) options->
      number_of_samples;
16  double two_pi_over_number_samples = 2.0 * PI / (double) options->
      number_of_samples;
17  int idxMasterBitOrder = -1 ;
18  int max_iterations_per_sample = ((bistable_OP *)options)->
      max_iterations_per_sample;
19  BUSLAYOUT_ITER bli ;
20 #ifdef REDUCE_DEREF
21   // For dereference reduction
22   int sim_data_number_samples = 0, pvt_vectors_icUsed = 0,

```



```

23     design_bus_layout_outputs_icUsed = 0,
        design_bus_layout_inputs_icUsed = 0, pvt_inputs_icUsed = 0 ;
24 int number_of_cells_in_current_layer = 0 ;
25 EXP_ARRAY *pvt_inputs = NULL ;
26 EXP_ARRAY *pvt_vectors = NULL ;
27 EXP_ARRAY *design_bus_layout_inputs = NULL ;
28 EXP_ARRAY *design_bus_layout_outputs = NULL ;
29 BUSLAYOUT *design_bus_layout = NULL ;
30 #endif
31 // For randomization
32 int Nix, Nix1, idxCell1, idxCell2 ;
33 QCADCell *swap = NULL ;
34 // — these used to be inside run_bistable_iteration — //
35 int q, iteration = 0;
36 int stable = FALSE;
37 double old_polarization;
38 double new_polarization;
39 double tolerance = ((bistable_OP *)options)->convergence_tolerance;
40 double polarization_math;
41 bistable_model *current_cell_model = NULL ;
42 QCADCell *cell;
43
44 STOP_SIMULATION = FALSE;
45
46 // — get the starting time for the simulation — //
47 if((start_time = time (NULL)) < 0)
48     fprintf(stderr, "Could not get start time\n");
49
50 // Create per-layer cell arrays to be used by the engine
51 simulation_inproc_data_new (design, &number_of_cell_layers, &
        number_of_cells_in_layer, &sorted_cells) ;
52
53 for(i = 0; i < number_of_cell_layers; i++)
54     {
55 #ifdef REDUCEDEREF
56     number_of_cells_in_current_layer = number_of_cells_in_layer[i] ;
57     for(j = 0; j < number_of_cells_in_current_layer ; j++)
58 #else
59     for(j = 0; j < number_of_cells_in_layer[i] ; j++)
60 #endif
61     {
62     // attach the model parameters to each of the simulation cells //
63     current_cell_model = g_malloc0 (sizeof(bistable_model)) ;
64     sorted_cells[i][j]->cell_model = current_cell_model;
65
66     // — Clear the model pointers so they are not dangling — //

```

```

67     current_cell_model->neighbours = NULL;
68     current_cell_model->Ek = NULL;
69
70     // — set polarization in cell model for fixed cells since they
       are set with actual dot charges by the user — //
71     if(QCAD_CELL_FIXED == sorted_cells[i][j]->cell_function)
72         current_cell_model->polarization =
           qcad_cell_calculate_polarization(sorted_cells[i][j]);
73
74     total_cells++;
75     }
76 }
77
78 // if we are performing a vector table simulation we consider only the
       activated inputs //
79 if(SIMULATION_TYPE == VECTOR_TABLE)
80     {
81     for (Nix = 0 ; Nix < pvt->inputs->icUsed ; Nix++)
82         if (!exp_array_index_1d (pvt->inputs, VT.INPUT, Nix).active_flag)
83             exp_array_index_1d (pvt->inputs, VT.INPUT, Nix).input->
               cell_function = QCAD_CELLNORMAL ;
84     }
85
86 // write message to the command history window //
87 command_history_message (_("Simulation found %d inputs %d outputs %d
       total cells\n"), design->bus_layout->inputs->icUsed, design->
       bus_layout->outputs->icUsed, total_cells) ;
88
89 command_history_message(_("Starting initialization\n"));
90 set_progress_bar_visible (TRUE) ;
91 set_progress_bar_label (_("Bistable simulation:")) ;
92
93 // — Initialize the simulation data structure — //
94 sim_data = g_malloc0 (sizeof(simulation_data));
95 sim_data->number_of_traces = design->bus_layout->inputs->icUsed +
       design->bus_layout->outputs->icUsed;
96 sim_data->number_samples = options->number_of_samples;
97 sim_data->trace = g_malloc0 (sizeof (struct TRACEDATA) * sim_data->
       number_of_traces);
98
99 // create and initialize the inputs into the sim data structure //
100 for (i = 0; i < design->bus_layout->inputs->icUsed; i++)
101     {
102     sim_data->trace[i].data_labels = g_strdup (qcad_cell_get_label (
           QCAD_CELL (exp_array_index_1d (design->bus_layout->inputs,
           BUSLAYOUT_CELL, i).cell))) ;

```

```

103     sim_data->trace[i].drawtrace = TRUE;
104     sim_data->trace[i].trace_function = QCAD_CELLINPUT;
105     sim_data->trace[i].data = g_malloc0 (sizeof (double) * sim_data->
        number_samples);
106     }
107
108 // create and initialize the outputs into the sim data structure //
109 for (i = 0; i < design->bus_layout->outputs->icUsed; i++)
110     {
111     sim_data->trace[i + design->bus_layout->inputs->icUsed].data_labels
        = g_strdup (qcad_cell_get_label(QCAD_CELL(exp_array_index_1d (
        design->bus_layout->outputs, BUSLAYOUT_CELL, i).cell)));
112     sim_data->trace[i + design->bus_layout->inputs->icUsed].drawtrace =
        TRUE;
113     sim_data->trace[i + design->bus_layout->inputs->icUsed].
        trace_function = QCAD_CELLOUTPUT;
114     sim_data->trace[i + design->bus_layout->inputs->icUsed].data =
        g_malloc0 (sizeof (double) * sim_data->number_samples);
115     }
116
117 // create and initialize the clock data //
118 sim_data->clock_data = g_malloc0 (sizeof (struct TRACEDATA) * 4);
119
120 for (i = 0; i < 4; i++)
121     {
122     sim_data->clock_data[i].data_labels = g_strdup_printf ("CLOCK %d", i
        );
123     sim_data->clock_data[i].drawtrace = 1;
124     sim_data->clock_data[i].trace_function = QCAD_CELLFIXED; // Abusing
        the notation here
125
126     sim_data->clock_data[i].data = g_malloc0 (sizeof (double) * sim_data
        ->number_samples);
127
128     if (SIMULATION_TYPE == EXHAUSTIVE_VERIFICATION)
129         for (j = 0; j < sim_data->number_samples; j++)
130             {
131             sim_data->clock_data[i].data[j] = clock_prefactor * cos (((
        double)(1 << design->bus_layout->inputs->icUsed)) * (double)
        j * four_pi_over_number_samples - PI * i / 2) + clock_shift
        ;
132             sim_data->clock_data[i].data[j] = CLAMP (sim_data->clock_data[i
        ].data[j], options->clock_low, options->clock_high) ;
133             }
134     else
135 //     if (SIMULATION_TYPE == VECTOR_TABLE)

```

```

136     for (j = 0; j < sim_data->number_samples; j++)
137     {
138         sim_data->clock_data[i].data[j] = clock_prefactor * cos (((
                double)pvt->vectors->icUsed) * (double)j *
                two_pi_over_number_samples - PI * i / 2) + clock_shift ;
139         sim_data->clock_data[i].data[j] = CLAMP (sim_data->clock_data[i
                ].data[j], options->clock_low, options->clock_high) ;
140     }
141 }
142
143 // — refresh all the kink energies to all the cells neighbours within
        the radius of effect — //
144 bistable_refresh_all_Ek (number_of_cell_layers ,
        number_of_cells_in_layer , sorted_cells , options);
145 //this function takes up pretty much the whole initialization time
146
147 // randomize the cells in the design so as to minimize any numerical
        problems associated //
148 // with having cells simulated in some predefined order. //
149 // randomize the order in which the cells are simulated //
150 //if (options->randomize_cells)
151 // for each layer ...
152 for (Nix = 0 ; Nix < number_of_cell_layers ; Nix++)
153     // ...perform as many swaps as there are cells therein
154     for (Nix1 = 0 ; Nix1 < number_of_cells_in_layer[Nix] ; Nix1++)
155     {
156         idxCell1 = rand () % number_of_cells_in_layer[Nix] ;
157         idxCell2 = rand () % number_of_cells_in_layer[Nix] ;
158
159         swap = sorted_cells[Nix][idxCell1] ;
160         sorted_cells[Nix][idxCell1] = sorted_cells[Nix][idxCell2] ;
161         sorted_cells[Nix][idxCell2] = swap ;
162     }
163
164 // — get and print the total initialization time — //
165 if((end_time = time (NULL)) < 0)
166     fprintf(stderr, "Could not get end time\n");
167
168 command_history_message("Total initialization time: %g s\n", (double)(
        end_time - start_time));
169
170 command_history_message("Starting Simulation\n");
171
172 set_progress_bar_fraction (0.0) ;
173
174 // perform the iterations over all samples //

```

```

175 #ifndef REDUCE_DEREF
176 // Dereference some structures now so we don't do it over and over in
    the loop
177 sim_data_number_samples = sim_data->number_samples ;
178 pvt_inputs = pvt->inputs ;
179 pvt_inputs_icUsed = pvt_inputs->icUsed ;
180 pvt_vectors = pvt->vectors ;
181 pvt_vectors_icUsed = pvt->vectors->icUsed ;
182 design_bus_layout = design->bus_layout ;
183 design_bus_layout_inputs = design_bus_layout->inputs ;
184 design_bus_layout_inputs_icUsed = design_bus_layout_inputs->icUsed ;
185 design_bus_layout_outputs = design_bus_layout->outputs ;
186 design_bus_layout_outputs_icUsed = design_bus_layout_outputs->icUsed ;
187 #else
188 #define sim_data_number_samples sim_data->number_samples
189 #define pvt_inputs pvt->inputs
190 #define pvt_inputs_icUsed pvt_inputs->icUsed
191 #define pvt_vectors pvt->vectors
192 #define pvt_vectors_icUsed pvt->vectors->icUsed
193 #define design_bus_layout design->bus_layout
194 #define design_bus_layout_inputs design_bus_layout->inputs
195 #define design_bus_layout_inputs_icUsed design_bus_layout_inputs->
    icUsed
196 #define design_bus_layout_outputs design_bus_layout->outputs
197 #define design_bus_layout_outputs_icUsed design_bus_layout_outputs->
    icUsed
198 #endif
199 for (j = 0; j < sim_data_number_samples ; j++)
200 {
201     if (j % 100 == 0)
202     {
203         // write the completion percentage to the command history window
        //
204         set_progress_bar_fraction ((float) j / (float)
            sim_data_number_samples) ;
205         // redraw the design if the user wants it to appear animated //
206         if (options->animate_simulation)
207         {
208             // update the charges to reflect the polarizations so that they
            can be animated //
209             for (icLayers = 0; icLayers < number_of_cell_layers; icLayers++)
210             {
211 #ifndef REDUCE_DEREF
212                 number_of_cells_in_current_layer = number_of_cells_in_layer[
                    icLayers] ;

```

```

213         for(icCellsInLayer = 0; icCellsInLayer <
                number_of_cells_in_current_layer; icCellsInLayer++)
214 #else
215         for(icCellsInLayer = 0; icCellsInLayer <
                number_of_cells_in_layer [icLayers]; icCellsInLayer++)
216 #endif
217         qcad_cell_set_polarization(sorted_cells [icLayers][
                icCellsInLayer],((bistable_model *)sorted_cells [icLayers
                ][icCellsInLayer]->cell_model)->polarization);
218     }
219 #ifdef DESIGNER
220     redraw_async(NULL);
221     gdk_flush ();
222 #endif /* def DESIGNER */
223     }
224 }
225
226 // — for each of the (VECTOR_TABLE => active?) inputs — //
227 if (EXHAUSTIVE_VERIFICATION == SIMULATION_TYPE)
228     for (idxMasterBitOrder = 0, design_bus_layout_iter_first (
                design_bus_layout, &bli, QCAD_CELLINPUT, &i) ; i > -1 ;
                design_bus_layout_iter_next (&bli, &i), idxMasterBitOrder++)
229         ((bistable_model *)exp_array_index_1d (design_bus_layout_inputs ,
                BUSLAYOUT_CELL, i).cell->cell_model)->polarization =
230         sim_data->trace [i].data [j] = (-1 * sin (((double)(1 <<
                idxMasterBitOrder)) * (double)j * FOUR_PI / (double)
                sim_data.number_samples) > 0) ? 1 : -1 ;
231 else
232 //     if (VECTOR_TABLE == SIMULATION_TYPE)
233     for (design_bus_layout_iter_first (design_bus_layout, &bli,
                QCAD_CELLINPUT, &i) ; i > -1 ; design_bus_layout_iter_next (&
                bli, &i))
234         if (exp_array_index_1d (pvt_inputs, VT_INPUT, i).active_flag)
235             ((bistable_model *)exp_array_index_1d (pvt_inputs, VT_INPUT, i
                ).input->cell_model)->polarization =
236             sim_data->trace [i].data [j] = exp_array_index_2d (pvt_vectors
                , gboolean, (j * pvt_vectors.icUsed) /
                sim_data.number_samples, i) ? 1 : -1 ;
237
238 // randomize the order in which the cells are simulated to try and
                minimize numerical errors
239 // associated with the imposed simulation order.
240 if(options->randomize_cells)
241     // for each layer ...
242     for (Nix = 0 ; Nix < number_of_cell_layers ; Nix++)
243         {

```

```

244     // ...perform as many swaps as there are cells therein
245 #ifdef REDUCE_DEREF
246     number_of_cells_in_current_layer = number_of_cells_in_layer[Nix]
        ;
247     for (Nix1 = 0 ; Nix1 < number_of_cells_in_current_layer ; Nix1
        ++)
248 #else
249     for (Nix1 = 0 ; Nix1 < number_of_cells_in_layer[Nix] ; Nix1++)
250 #endif
251     {
252 #ifdef REDUCE_DEREF
253         idxCell1 = rand () % number_of_cells_in_current_layer ;
254         idxCell2 = rand () % number_of_cells_in_current_layer ;
255 #else
256         idxCell1 = rand () % number_of_cells_in_layer[Nix] ;
257         idxCell2 = rand () % number_of_cells_in_layer[Nix] ;
258 #endif
259
260         swap = sorted_cells[Nix][idxCell1] ;
261         sorted_cells[Nix][idxCell1] = sorted_cells[Nix][idxCell2] ;
262         sorted_cells[Nix][idxCell2] = swap ;
263     }
264 }
265
266 // — run the iteration with the given clock value — //
267 // — iterate until the entire design has stabalized — //
268 iteration = 0;
269 stable = FALSE;
270 while (!stable && iteration < max_iterations_per_sample)
271     {
272     iteration++;
273     // — assume that the circuit is stable — //
274     stable = TRUE;
275
276     for (icLayers = 0; icLayers < number_of_cell_layers; icLayers++)
277     {
278 #ifdef REDUCE_DEREF
279         number_of_cells_in_current_layer = number_of_cells_in_layer[
            icLayers] ;
280         for (icCellsInLayer = 0 ; icCellsInLayer <
            number_of_cells_in_current_layer ; icCellsInLayer++)
281 #else
282         for (icCellsInLayer = 0 ; icCellsInLayer <
            number_of_cells_in_layer[icLayers] ; icCellsInLayer++)
283 #endif
284         {

```

```

285     cell = sorted_cells[icLayers][icCellsInLayer] ;
286
287     if (!(QCAD_CELLINPUT == cell->cell_function) ||
288         (QCAD_CELLFIXED == cell->cell_function))
289     {
290         current_cell_model = ((bistable_model *)cell->cell_model) ;
291         old_polarization = current_cell_model->polarization;
292         polarization_math = 0;
293
294         for (q = 0; q < current_cell_model->number_of_neighbours; q
295             ++
296             polarization_math += (current_cell_model->Ek[q] * ((
297                 bistable_model *)current_cell_model->neighbours[q]->
298                 cell_model)->polarization);
299
300         // math = math / 2 * gamma
301         polarization_math /= (2.0 * sim_data->clock_data[cell->
302             cell_options.clock].data[j]);
303
304         // --- calculate the new cell polarization --- //
305         // if math < 0.05 then math/sqrt(1+math^2) ~ math with
306         // error <= 4e-5
307         // if math > 100 then math/sqrt(1+math^2) ~ +-1 with error
308         // <= 5e-5
309         new_polarization =
310             (polarization_math > 1000.0) ? 1
311             :
312             (polarization_math < -1000.0) ? -1
313             :
314             (fabs (polarization_math) < 0.001) ?
315                 polarization_math :
316                 polarization_math / sqrt (1 + polarization_math *
317                 polarization_math) ;
318
319         // --- set the polarization of this cell --- //
320         current_cell_model->polarization = new_polarization;
321
322         // If any cells polarization has changed beyond this
323         // threshold
324         // then the entire circuit is assumed to have not converged.
325         stable = (fabs (new_polarization - old_polarization) <=
326             tolerance) && stable ;
327     }
328 }
329 }
330 }
331 }
332 }
333 }
334 }
335 }
336 }
337 }
338 }
339 }
340 }
341 }
342 }
343 }
344 }
345 }
346 }
347 }
348 }
349 }
350 }
351 }
352 }
353 }
354 }
355 }
356 }
357 }
358 }
359 }
360 }
361 }
362 }
363 }
364 }
365 }
366 }
367 }
368 }
369 }
370 }
371 }
372 }
373 }
374 }
375 }
376 }
377 }
378 }
379 }
380 }
381 }
382 }
383 }
384 }
385 }
386 }
387 }
388 }
389 }
390 }
391 }
392 }
393 }
394 }
395 }
396 }
397 }
398 }
399 }
400 }
401 }
402 }
403 }
404 }
405 }
406 }
407 }
408 }
409 }
410 }
411 }
412 }
413 }
414 }
415 }
416 }
417 }
418 }
419 }
420 }
421 }
422 }
423 }
424 }
425 }
426 }
427 }
428 }
429 }
430 }
431 }
432 }
433 }
434 }
435 }
436 }
437 }
438 }
439 }
440 }
441 }
442 }
443 }
444 }
445 }
446 }
447 }
448 }
449 }
450 }
451 }
452 }
453 }
454 }
455 }
456 }
457 }
458 }
459 }
460 }
461 }
462 }
463 }
464 }
465 }
466 }
467 }
468 }
469 }
470 }
471 }
472 }
473 }
474 }
475 }
476 }
477 }
478 }
479 }
480 }
481 }
482 }
483 }
484 }
485 }
486 }
487 }
488 }
489 }
490 }
491 }
492 }
493 }
494 }
495 }
496 }
497 }
498 }
499 }
500 }
501 }
502 }
503 }
504 }
505 }
506 }
507 }
508 }
509 }
510 }
511 }
512 }
513 }
514 }
515 }
516 }
517 }
518 }
519 }
520 }
521 }
522 }
523 }
524 }
525 }
526 }
527 }
528 }
529 }
530 }
531 }
532 }
533 }
534 }
535 }
536 }
537 }
538 }
539 }
540 }
541 }
542 }
543 }
544 }
545 }
546 }
547 }
548 }
549 }
550 }
551 }
552 }
553 }
554 }
555 }
556 }
557 }
558 }
559 }
560 }
561 }
562 }
563 }
564 }
565 }
566 }
567 }
568 }
569 }
570 }
571 }
572 }
573 }
574 }
575 }
576 }
577 }
578 }
579 }
580 }
581 }
582 }
583 }
584 }
585 }
586 }
587 }
588 }
589 }
590 }
591 }
592 }
593 }
594 }
595 }
596 }
597 }
598 }
599 }
600 }
601 }
602 }
603 }
604 }
605 }
606 }
607 }
608 }
609 }
610 }
611 }
612 }
613 }
614 }
615 }
616 }
617 }
618 }
619 }
620 }
621 }
622 }
623 }
624 }
625 }
626 }
627 }
628 }
629 }
630 }
631 }
632 }
633 }
634 }
635 }
636 }
637 }
638 }
639 }
640 }
641 }
642 }
643 }
644 }
645 }
646 }
647 }
648 }
649 }
650 }
651 }
652 }
653 }
654 }
655 }
656 }
657 }
658 }
659 }
660 }
661 }
662 }
663 }
664 }
665 }
666 }
667 }
668 }
669 }
670 }
671 }
672 }
673 }
674 }
675 }
676 }
677 }
678 }
679 }
680 }
681 }
682 }
683 }
684 }
685 }
686 }
687 }
688 }
689 }
690 }
691 }
692 }
693 }
694 }
695 }
696 }
697 }
698 }
699 }
700 }
701 }
702 }
703 }
704 }
705 }
706 }
707 }
708 }
709 }
710 }
711 }
712 }
713 }
714 }
715 }
716 }
717 }
718 }
719 }
720 }
721 }
722 }
723 }
724 }
725 }
726 }
727 }
728 }
729 }
730 }
731 }
732 }
733 }
734 }
735 }
736 }
737 }
738 }
739 }
740 }
741 }
742 }
743 }
744 }
745 }
746 }
747 }
748 }
749 }
750 }
751 }
752 }
753 }
754 }
755 }
756 }
757 }
758 }
759 }
760 }
761 }
762 }
763 }
764 }
765 }
766 }
767 }
768 }
769 }
770 }
771 }
772 }
773 }
774 }
775 }
776 }
777 }
778 }
779 }
780 }
781 }
782 }
783 }
784 }
785 }
786 }
787 }
788 }
789 }
790 }
791 }
792 }
793 }
794 }
795 }
796 }
797 }
798 }
799 }
800 }
801 }
802 }
803 }
804 }
805 }
806 }
807 }
808 }
809 }
810 }
811 }
812 }
813 }
814 }
815 }
816 }
817 }
818 }
819 }
820 }
821 }
822 }
823 }
824 }
825 }
826 }
827 }
828 }
829 }
830 }
831 }
832 }
833 }
834 }
835 }
836 }
837 }
838 }
839 }
840 }
841 }
842 }
843 }
844 }
845 }
846 }
847 }
848 }
849 }
850 }
851 }
852 }
853 }
854 }
855 }
856 }
857 }
858 }
859 }
860 }
861 }
862 }
863 }
864 }
865 }
866 }
867 }
868 }
869 }
870 }
871 }
872 }
873 }
874 }
875 }
876 }
877 }
878 }
879 }
880 }
881 }
882 }
883 }
884 }
885 }
886 }
887 }
888 }
889 }
890 }
891 }
892 }
893 }
894 }
895 }
896 }
897 }
898 }
899 }
900 }
901 }
902 }
903 }
904 }
905 }
906 }
907 }
908 }
909 }
910 }
911 }
912 }
913 }
914 }
915 }
916 }
917 }
918 }
919 }
920 }
921 }
922 }
923 }
924 }
925 }
926 }
927 }
928 }
929 }
930 }
931 }
932 }
933 }
934 }
935 }
936 }
937 }
938 }
939 }
940 }
941 }
942 }
943 }
944 }
945 }
946 }
947 }
948 }
949 }
950 }
951 }
952 }
953 }
954 }
955 }
956 }
957 }
958 }
959 }
960 }
961 }
962 }
963 }
964 }
965 }
966 }
967 }
968 }
969 }
970 }
971 }
972 }
973 }
974 }
975 }
976 }
977 }
978 }
979 }
980 }
981 }
982 }
983 }
984 }
985 }
986 }
987 }
988 }
989 }
990 }
991 }
992 }
993 }
994 }
995 }
996 }
997 }
998 }
999 }
1000 }

```



```

319
320     if (VECTOR.TABLE == SIMULATION.TYPE)
321         for (design_bus_layout_iter_first (design_bus_layout, &bli,
322             QCAD.CELLINPUT, &i) ; i > -1 ; design_bus_layout_iter_next (&
323             bli, &i))
324             if (!exp_array_index_1d (pvt_inputs, VT.INPUT, i).active_flag)
325                 sim_data->trace[i].data[j] = ((bistable_model *)
326                     exp_array_index_1d (pvt_inputs, VT.INPUT, i).input->
327                     cell_model)->polarization;
328
329     // — collect all the output data from the simulation — //
330     for (design_bus_layout_iter_first (design_bus_layout, &bli,
331         QCAD.CELLOUTPUT, &i) ; i > -1 ; design_bus_layout_iter_next (&
332         bli, &i)){
333         sim_data->trace[design_bus_layout_inputs_icUsed + i].data[j] = ((
334             bistable_model *)exp_array_index_1d (design_bus_layout_outputs
335             , BUSLAYOUT.CELL, i).cell->cell_model)->polarization;
336     }
337     // — if the user wants to stop the simulation then exit. — //
338     if(TRUE == STOP.SIMULATION)
339         j = sim_data_number_samples ;
340     }//for number of samples
341
342     // Free the neighbours and Ek array introduced by this simulation//
343     for (k = 0; k < number_of_cell_layers; k++)
344     {
345     #ifdef REDUCE.DEREF
346         number_of_cells_in_current_layer = number_of_cells_in_layer[k] ;
347         for (l = 0 ; l < number_of_cells_in_current_layer ; l++)
348     #else
349         for (l = 0 ; l < number_of_cells_in_layer[k] ; l++)
350     #endif
351     {
352         g_free(((bistable_model *)sorted_cells[k][l]->cell_model)->
353             neighbours);
354         g_free(((bistable_model *)sorted_cells[k][l]->cell_model)->
355             neighbour_layer);
356         g_free(((bistable_model *)sorted_cells[k][l]->cell_model)->Ek);
357     }
358     }
359
360     simulation_inproc_data_free (&number_of_cell_layers, &
361         number_of_cells_in_layer, &sorted_cells) ;
362
363     // Restore the input flag for the inactive inputs
364     if (VECTOR.TABLE == SIMULATION.TYPE)

```

```

354     for (i = 0 ; i < pvt_inputs_icUsed ; i++)
355         exp_array_index_1d (pvt_inputs , VT_INPUT, i).input->cell_function
           = QCAD_CELLINPUT ;
356
357 // — get and print the total simulation time — //
358     if ((end_time = time (NULL)) < 0)
359         fprintf(stderr , "Could not get end time\n");
360
361     command_history_message ("Total simulation time: %g s\n", (double)(
           end_time - start_time));
362     set_progress_bar_visible (FALSE) ;
363
364 #ifndef REDUCE_DEREF
365     #undef sim_data_number_samples
366     #undef pvt_inputs
367     #undef pvt_inputs_icUsed
368     #undef pvt_vectors
369     #undef pvt_vectors_icUsed
370     #undef design_bus_layout
371     #undef design_bus_layout_inputs
372     #undef design_bus_layout_inputs_icUsed
373     #undef design_bus_layout_outputs
374     #undef design_bus_layout_outputs_icUsed
375 #endif
376
377     return sim_data ;
378 } // run_bistable

```

Listing A.1: Bistable simulation engine software code

## Appendix B

# run\_bistable\_simulation\_hardware C Code Listing

The following is the code for the version of run\_bistable\_simulation created to use the hardware

```

1 //calls the hardware to run an alternate version of the bistable
  simulation engine
2 simulation_data *run_bistable_simulation_hardware (int SIMULATION_TYPE,
  DESIGN *design, bistable_OP *options, VectorTable *pvt)
3 {
4   int i, j, k, l, total_cells = 0 ;
5   int icLayers, icCellsInLayer;
6   time_t start_time, end_time;
7   simulation_data *sim_data = NULL ;
8   // optimization variables //
9   int number_of_cell_layers = 0, *number_of_cells_in_layer = NULL ;
10  QCADCell ***sorted_cells = NULL ;
11  float clock_shift = (float)((options->clock_high + options->clock_low)
  /2 + options->clock_shift);
12  float clock_prefactor = (float)(options->clock_high - options->
  clock_low) * options->clock_amplitude_factor;
13  float four_pi_over_number_samples = 4.0 * (float)PI / (float) options
  ->number_of_samples;
14  float two_pi_over_number_samples = 2.0 * (float)PI / (float) options->
  number_of_samples;
15  int idxMasterBitOrder = -1 ;
16  int max_iterations_per_sample = ((bistable_OP *)options)->
  max_iterations_per_sample;
17  BUSLAYOUT_ITER bli ;
18 #ifdef REDUCE_DEREF
19 // For dereference reduction
20 int sim_data_number_samples = 0, pvt_vectors_icUsed = 0,
21 design_bus_layout_outputs_icUsed = 0,
  design_bus_layout_inputs_icUsed = 0, pvt_inputs_icUsed = 0 ;

```

```

22     int number_of_cells_in_current_layer = 0 ;
23     EXP_ARRAY *pvt_inputs = NULL ;
24     EXP_ARRAY *pvt_vectors = NULL ;
25     EXP_ARRAY *design_bus_layout_inputs = NULL ;
26     EXP_ARRAY *design_bus_layout_outputs = NULL ;
27     BUS_LAYOUT *design_bus_layout = NULL ;
28 #endif
29     // — these used to be inside run_bistable_iteration — //
30     int Nix;
31     int q, iteration, calculatedCellNum= 0;
32     int stable = FALSE;
33     float old_polarization;
34     float new_polarization;
35     float tolerance = (float)((bistable_OP *)options)->
        convergence_tolerance;
36     float temp; //just used to move data formats between DWORDS and floats
37     bistable_model *current_cell_model = NULL ;
38     QCADCell *cell;
39
40     //wildcard initialization
41     WC_RetCode rc = WC_SUCCESS;
42     WC_TestInfo TestInfo;
43     char **TestLoc = NULL;
44     DWORD control[8];
45     WC4_DmaHandle hDMA_Src;
46     WC4_DmaHandle hDMA_Dst;
47     DWORD dGrantedDwords;
48     DWORD dSourceHardwareAddress;
49     DWORD dDestinationHardwareAddress;
50     float *DMABufSrc;
51     float *DMABufDest;
52     boolean intStatus;
53     DWORD debugReg[8];
54     DWORD maxNeighbours_DMA = 0x400; //max number of neighbours
        supported, used to allocate source buffer (1024)
55     DWORD numCells_DMA = 0; //number of cells to be read from the PE;
        sets DMA destination buffer; calculated later
56     DWORD totem = 0xBF000000; //totem value used to signal end of things
57
58     TestInfo.bVerbose      = DEFAULT_VERBOSITY;
59     TestInfo.DeviceNum    = DEFAULT_SLOT_NUMBER;
60     TestInfo.dIterations  = DEFAULT_ITERATIONS;
61     TestInfo.fClkFreq     = DEFAULT_FREQUENCY;
62
63     STOP_SIMULATION = FALSE;
64

```

```

65 // -- get the starting time for the simulation -- //
66 if((start_time = time (NULL)) < 0)
67     fprintf(stderr, "Could not get start time\n");
68 // Create per-layer cell arrays to be used by the engine
69 simulation_inproc_data_new (design, &number_of_cell_layers, &
    number_of_cells_in_layer, &sorted_cells);
70 for(i = 0; i < number_of_cell_layers; i++)
71     {
72 #ifdef REDUCEDEREF
73     number_of_cells_in_current_layer = number_of_cells_in_layer[i];
74     for(j = 0; j < number_of_cells_in_current_layer; j++)
75 #else
76     for(j = 0; j < number_of_cells_in_layer[i]; j++)
77 #endif
78     {
79 // attach the model parameters to each of the simulation cells //
80 current_cell_model = g_malloc0 (sizeof(bistable_model));
81 sorted_cells[i][j]->cell_model = current_cell_model;
82
83 // -- Clear the model pointers so they are not dangling -- //
84 current_cell_model->neighbours = NULL;
85 current_cell_model->Ek = NULL;
86
87 // -- set polarization in cell model for fixed cells since they
    are set with actual dot charges by the user -- //
88 if(QCAD_CELL_FIXED == sorted_cells[i][j]->cell_function)
89     current_cell_model->polarization =
        qcad_cell_calculate_polarization(sorted_cells[i][j]);
90
91 total_cells++;
92
93 if (!(QCAD_CELL_INPUT == sorted_cells[i][j]->cell_function) ||
    QCAD_CELL_FIXED == sorted_cells[i][j]->cell_function))
94     numCells_DMA++; //how many cells will be calculated by the
        hardware
95     }
96     }
97
98 // if we are performing a vector table simulation we consider only the
    activated inputs //
99 if(SIMULATION_TYPE == VECTOR_TABLE)
100     {
101     for (Nix = 0; Nix < pvt->inputs->icUsed; Nix++)
102         if (!exp_array_index_1d (pvt->inputs, VT.INPUT, Nix).active_flag)
103             exp_array_index_1d (pvt->inputs, VT.INPUT, Nix).input->
                cell_function = QCAD_CELL_NORMAL;

```

```

104     }
105
106     // write message to the command history window //
107     command_history_message (_("Simulation found %d inputs %d outputs %d
        total cells\n"), design->bus_layout->inputs->icUsed, design->
        bus_layout->outputs->icUsed, total_cells) ;
108
109     command_history_message(_("Starting initialization\n"));
110     command_history_message(_("Opening card.."));
111     rc = WC_Open (TestInfo.DeviceNum, 0);
112     if (rc != WCSUCCESS) {
113         command_history_message("Failed to open card.\n");
114         return NULL;
115     }
116
117
118     command_history_message(_("Initializing hardware..."));
119     //initialize and reset the PE
120     rc = wc4_qca_init( &TestInfo);
121     if (rc != WCSUCCESS) {
122         command_history_message("initialization failed\n");
123         return NULL;
124     }
125     //CHECK_RC(rc);
126
127     command_history_message(_("resetting hardware..."));
128     rc = WC_PeReset (TestInfo.DeviceNum, TRUE);
129     if (rc != WCSUCCESS) {
130         command_history_message("failed\n");
131         return NULL;
132     }
133     //CHECK_RC(rc);
134
135     rc = WC_PeReset (TestInfo.DeviceNum, FALSE);
136     if (rc != WCSUCCESS) {
137         command_history_message("failed\n");
138         return NULL;
139     }
140     //CHECK_RC (rc);
141     command_history_message("Done\n");
142
143     set_progress_bar_visible (TRUE) ;
144     set_progress_bar_label (_("Bistable simulation:")) ;
145
146     // --- Initialize the simulation data structure --- //
147     sim_data = g_malloc0 (sizeof(simulation_data));

```

```

148  sim_data->number_of_traces = design->bus_layout->inputs->icUsed +
      design->bus_layout->outputs->icUsed;
149  sim_data->number_samples = options->number_of_samples;
150  sim_data->trace = g_malloc0 (sizeof (struct TRACEDATA) * sim_data->
      number_of_traces);
151
152  // create and initialize the inputs into the sim data structure //
153  for (i = 0; i < design->bus_layout->inputs->icUsed; i++)
154  {
155      sim_data->trace[i].data_labels = g_strdup (qcad_cell_get_label (
          QCAD_CELL (exp_array_index_1d (design->bus_layout->inputs,
          BUSLAYOUT_CELL, i).cell)) );
156      sim_data->trace[i].drawtrace = TRUE;
157      sim_data->trace[i].trace_function = QCAD_CELLINPUT;
158      sim_data->trace[i].data = g_malloc0 (sizeof (double) * sim_data->
          number_samples);
159  }
160
161  // create and initialize the outputs into the sim data structure //
162  for (i = 0; i < design->bus_layout->outputs->icUsed; i++)
163  {
164      sim_data->trace[i + design->bus_layout->inputs->icUsed].data_labels
          = g_strdup (qcad_cell_get_label(QCAD_CELL(exp_array_index_1d (
          design->bus_layout->outputs, BUSLAYOUT_CELL, i).cell)));
165      sim_data->trace[i + design->bus_layout->inputs->icUsed].drawtrace =
          TRUE;
166      sim_data->trace[i + design->bus_layout->inputs->icUsed].
          trace_function = QCAD_CELLOUTPUT;
167      sim_data->trace[i + design->bus_layout->inputs->icUsed].data =
          g_malloc0 (sizeof (double) * sim_data->number_samples);
168  }
169
170  // create and initialize the clock data //
171  sim_data->clock_data = g_malloc0 (sizeof (struct TRACEDATA) * 4);
172
173  for (i = 0; i < 4; i++)
174  {
175      sim_data->clock_data[i].data_labels = g_strdup_printf ("CLOCK %d", i
          );
176      sim_data->clock_data[i].drawtrace = 1;
177      sim_data->clock_data[i].trace_function = QCAD_CELLFIXED; // Abusing
          the notation here
178
179      sim_data->clock_data[i].data = g_malloc0 (sizeof (double) * sim_data
          ->number_samples);
180

```

```

181     if (SIMULATION_TYPE == EXHAUSTIVE_VERIFICATION)
182         for (j = 0; j < sim_data->number_samples; j++)
183             {
184                 sim_data->clock_data[i].data[j] = clock_prefactor * cos (((
                    double)(1 << design->bus_layout->inputs->icUsed)) * (double)
                    j * four_pi_over_number_samples - PI * i / 2) + clock_shift
                    ;
185                 sim_data->clock_data[i].data[j] = CLAMP (sim_data->clock_data[i
                    ].data[j], options->clock_low, options->clock_high) ;
186             }
187     else
188 //     if (SIMULATION_TYPE == VECTOR_TABLE)
189         for (j = 0; j < sim_data->number_samples; j++)
190             {
191                 sim_data->clock_data[i].data[j] = clock_prefactor * cos (((
                    double)pvt->vectors->icUsed) * (double)j *
                    two_pi_over_number_samples - PI * i / 2) + clock_shift ;
192                 sim_data->clock_data[i].data[j] = CLAMP (sim_data->clock_data[i
                    ].data[j], options->clock_low, options->clock_high) ;
193             }
194     }
195
196     command_history_message(" Allocating DMA buffers...");
197     rc = WC4_DmaMemAlloc (TestInfo.DeviceNum, &hDMASource,
        maxNeighbours_DMA, (DWORD **)&DMABufSource);
198     if (rc != WC_SUCCESS) {
199         command_history_message("\nError Allocating DMA source buffer\n");
200         WC4_DmaMemFree(TestInfo.DeviceNum, hDMASource);
201         WC_PeDeprogram (TestInfo.DeviceNum);
202         WC4_PeInitiateAutoProgram (TestInfo.DeviceNum);
203         WC_Close (TestInfo.DeviceNum);
204         return NULL;
205     }
206     rc = WC4_DmaMemAlloc (TestInfo.DeviceNum, &hDMADestination,
        numCells_DMA, (DWORD **)&DMABufDest);
207     if (rc != WC_SUCCESS) {
208         command_history_message("\nError Allocating DMA destination buffer\n
        " );
209         WC4_DmaMemFree(TestInfo.DeviceNum, hDMASource);
210         WC4_DmaMemFree(TestInfo.DeviceNum, hDMADestination);
211         WC_PeDeprogram (TestInfo.DeviceNum);
212         WC4_PeInitiateAutoProgram (TestInfo.DeviceNum);
213         WC_Close (TestInfo.DeviceNum);
214         return NULL;
215     }
216     printf("DONE\n");

```



```

217
218 printf(" Binding DMA buffers\t\t\t");
219 rc = WC4_DmaBind (TestInfo.DeviceNum, hDMASource, (DWORD *)
    DMABufSource, maxNeighbours_DMA, &dGrantedDwords, &
    dSourceHardwareAddress );
220 if (rc != WC_SUCCESS) {
221     printf ("\nError binding DMA source buffer\n" );
222     WC4_DmaMemFree (TestInfo.DeviceNum, hDMASource);
223     WC4_DmaMemFree (TestInfo.DeviceNum, hDMADestination);
224     WC_PeDeprogram (TestInfo.DeviceNum);
225     WC4_PeInitiateAutoProgram (TestInfo.DeviceNum);
226     WC_Close (TestInfo.DeviceNum);
227     return NULL;
228 } else if (maxNeighbours_DMA != dGrantedDwords) {
229     printf ("\nError could not bind full source buffer!\n");
230     WC4_DmaMemFree (TestInfo.DeviceNum, hDMASource);
231     WC4_DmaMemFree (TestInfo.DeviceNum, hDMADestination);
232     WC_PeDeprogram (TestInfo.DeviceNum);
233     WC4_PeInitiateAutoProgram (TestInfo.DeviceNum);
234     WC_Close (TestInfo.DeviceNum);
235     return NULL;
236 }
237
238 rc = WC4_DmaBind (TestInfo.DeviceNum, hDMADestination, (DWORD *)
    DMABufDest, numCells_DMA, &dGrantedDwords, &
    dDestinationHardwareAddress );
239 if (rc != WC_SUCCESS) {
240     printf ("\nError binding DMA destination buffer\n");
241     WC4_DmaUnbind (TestInfo.DeviceNum, hDMASource);
242     WC4_DmaMemFree (TestInfo.DeviceNum, hDMADestination);
243     WC4_DmaMemFree (TestInfo.DeviceNum, hDMASource);
244     WC_PeDeprogram (TestInfo.DeviceNum);
245     WC4_PeInitiateAutoProgram (TestInfo.DeviceNum);
246     WC_Close (TestInfo.DeviceNum);
247     return NULL;
248 } else if (numCells_DMA != dGrantedDwords) {
249     printf ("\nError could not bind full destination buffer!\n");
250     WC4_DmaUnbind (TestInfo.DeviceNum, hDMASource);
251     WC4_DmaMemFree (TestInfo.DeviceNum, hDMADestination);
252     WC4_DmaMemFree (TestInfo.DeviceNum, hDMASource);
253     WC_PeDeprogram (TestInfo.DeviceNum);
254     WC4_PeInitiateAutoProgram (TestInfo.DeviceNum);
255     WC_Close (TestInfo.DeviceNum);
256     return NULL;
257 }
258 printf("DONE\n");

```

```

259
260 //send DMA control data to PE
261 control[0] = dDestinationHardwareAddress;
262 control[1] = dSourceHardwareAddress;
263 control[2] = numCells_DMA;
264 command_history_message(" Sending DMA control data\n");
265 rc = WC_PeRegWrite( TestInfo.DeviceNum, CTRL_REG_BASE, 3, control);
266 if (rc != WC_SUCCESS) {
267     command_history_message(" Register transfer failed\n");
268     wildcard_cleanup( TestInfo.DeviceNum, &hDMA_Src, &hDMA_Dst);
269     return NULL;
270 }
271 WC_PeRegRead( TestInfo.DeviceNum, CTRL_REG_BASE+4,4, debugReg);
272 command_history_message(" state: %x\n", debugReg[0]);
273 command_history_message(" Calculating kink energies\n");
274 command_history_message("DMA To PE Addr: %x\n",dSourceHardwareAddress)
    ;
275 // — refresh all the kink energies to all the cells neighbours within
    the radius of effect — //
276 bistable_refresh_all_Ek_hardware ( number_of_cell_layers ,
    number_of_cells_in_layer , sorted_cells , options , DMA_Buf_Src,
    TestInfo.DeviceNum);
277 //Using TestInfo.DeviceNum instead of TestInfo->DeviceNum. Don't know
    why it has to be like this. Make sure this doesn't cause issues
278 command_history_message(" done calculating kink energies , waiting for
    interrupt\n");
279 WC_IntQueryStatus( TestInfo.DeviceNum, &intStatus);
280 command_history_message(" interrupt status: %d\n",intStatus);
281 rc=WC_IntWait( TestInfo.DeviceNum, INT_TIMEOUT_MS); //wait for
    interrupt for the last set of Eks to be confirmed as received
282 //might want to find a different way of testing for the interrupt ,
    because this seems cumbersome.
283 if(rc != WC_SUCCESS) {
284     WC_PeRegRead( TestInfo.DeviceNum, CTRL_REG_BASE, 8, debugReg);
285     command_history_message(" Interrupt timed out; int Counter ,
        numNeighbours , dram_addr , dma_incount , state , in_fifo_out , DRAM
        data , DMA_outCount: %x %x %x %x %x %x %x %x\n" ,
286         debugReg[0] , debugReg[1] , debugReg[2] , debugReg[3] , debugReg[4] ,
        debugReg[5] , debugReg[6] , debugReg[7]);
287     wildcard_cleanup( TestInfo.DeviceNum, &hDMA_Src, &hDMA_Dst);
288     return NULL;
289 }
290 command_history_message(" interrupt received\n");
291 rc=WC_IntEnable( TestInfo.DeviceNum, FALSE); //reset interrupt
292 rc=WC_IntReset( TestInfo.DeviceNum);
293 rc=WC_IntEnable( TestInfo.DeviceNum, TRUE);

```

```

294     command_history_message("interrupt reset\n");
295     WC_PeRegWrite(TestInfo.DeviceNum, CTRL_REG_BASE+3, 1, &totem); //send
        a totem value to tell the PE that all Ek data has been sent
296
297     // — get and print the total initialization time — //
298     if((end_time = time (NULL)) < 0)
299         fprintf(stderr, "Could not get end time\n");
300
301     //there was a cell randomization sequence here, but the hardware
        accelerator makes it unnecessary
302
303     command_history_message("Total initialization time: %g s\n", (double)(
        end_time - start_time));
304
305     command_history_message("Starting Simulation\n");
306     set_progress_bar_fraction (0.0) ;
307
308     // perform the iterations over all samples //
309 #ifdef REDUCEDEREF
310     // Dereference some structures now so we don't do it over and over in
        the loop
311     sim_data_number_samples = sim_data->number_samples ;
312     pvt_inputs = pvt->inputs ;
313     pvt_inputs_icUsed = pvt_inputs->icUsed ;
314     pvt_vectors = pvt->vectors ;
315     pvt_vectors_icUsed = pvt->vectors->icUsed ;
316     design_bus_layout = design->bus_layout ;
317     design_bus_layout_inputs = design_bus_layout->inputs ;
318     design_bus_layout_inputs_icUsed = design_bus_layout_inputs->icUsed ;
319     design_bus_layout_outputs = design_bus_layout->outputs ;
320     design_bus_layout_outputs_icUsed = design_bus_layout_outputs->icUsed ;
321 #else
322     #define sim_data_number_samples sim_data->number_samples
323     #define pvt_inputs pvt->inputs
324     #define pvt_inputs_icUsed pvt_inputs->icUsed
325     #define pvt_vectors pvt->vectors
326     #define pvt_vectors_icUsed pvt->vectors->icUsed
327     #define design_bus_layout design->bus_layout
328     #define design_bus_layout_inputs design_bus_layout->inputs
329     #define design_bus_layout_inputs_icUsed design_bus_layout_inputs->
        icUsed
330     #define design_bus_layout_outputs design_bus_layout->outputs
331     #define design_bus_layout_outputs_icUsed design_bus_layout_outputs->
        icUsed
332 #endif
333     for (j = 0; j < sim_data_number_samples ; j++)

```

```

334     {
335     if (j % 100 == 0)
336     {
337         // write the completion percentage to the command history window
338         //
339         set_progress_bar_fraction ((float) j / (float)
340             sim_data.number_samples) ;
341         // redraw the design if the user wants it to appear animated //
342     }
343
344     // — for each of the (VECTOR_TABLE => active?) inputs — //
345     if (EXHAUSTIVE_VERIFICATION == SIMULATION_TYPE)
346     for (idxMasterBitOrder = 0, design_bus_layout_iter_first (
347         design_bus_layout, &bli, QCAD_CELLINPUT, &i) ; i > -1 ;
348         design_bus_layout_iter_next (&bli, &i), idxMasterBitOrder++)
349     ((bistable_model *)exp_array_index_1d (design_bus_layout_inputs,
350         BUSLAYOUT_CELL, i).cell->cell_model)->polarization =
351     sim_data->trace[i].data[j] = (-1 * sin (((double)(1 <<
352         idxMasterBitOrder)) * (double)j * FOUR_PI / (double)
353         sim_data.number_samples) > 0) ? 1 : -1 ;
354
355     else
356     // if (VECTOR_TABLE == SIMULATION_TYPE)
357     for (design_bus_layout_iter_first (design_bus_layout, &bli,
358         QCAD_CELLINPUT, &i) ; i > -1 ; design_bus_layout_iter_next (&
359         bli, &i))
360     if (exp_array_index_1d (pvt_inputs, VT_INPUT, i).active_flag)
361     ((bistable_model *)exp_array_index_1d (pvt_inputs, VT_INPUT, i
362         ).input->cell_model)->polarization =
363     sim_data->trace[i].data[j] = exp_array_index_2d (pvt_vectors
364         , gboolean, (j * pvt_vectors.icUsed) /
365         sim_data.number_samples, i) ? 1 : -1 ;
366
367     // — run the iteration with the given clock value — //
368     // — iterate until the entire design has stabilized — //
369     iteration = 0;
370     stable = FALSE;
371     while (!stable && iteration < max_iterations_per_sample)
372     {
373         iteration++;
374         // — assume that the circuit is stable — //
375         stable = TRUE;
376
377         for (icLayers = 0; icLayers < number_of_cell_layers; icLayers++)
378         {
379 #ifdef REDUCE_DEREF

```

```

367     number_of_cells_in_current_layer = number_of_cells_in_layer [
        icLayers] ;
368     for (icCellsInLayer = 0 ; icCellsInLayer <
        number_of_cells_in_current_layer ; icCellsInLayer++)
369 #else
370     for (icCellsInLayer = 0 ; icCellsInLayer <
        number_of_cells_in_layer [icLayers] ; icCellsInLayer++)
371 #endif
372     {
373         cell = sorted_cells [icLayers] [icCellsInLayer] ;
374
375         if (!( (QCAD_CELLINPUT == cell->cell_function) ||
376             (QCAD_CELL_FIXED == cell->cell_function) ))
377             {
378                 current_cell_model = ((bistable_model *) cell->cell_model) ;
379
380                 for (q = 0; q < current_cell_model->number_of_neighbours; q
                    ++))
381             {
382                 DMABufSource[q] = ((bistable_model *) current_cell_model->
                    neighbours [q]->cell_model)->polarization;
383             }
384
385         command_history_message(" setting gamma...");
386         temp = (float)(sim_data->clock_data [ cell->cell_options.clock].
            data[j]);
387         control[0] = *(DWORD *)&temp; //casting necessary to put a float
            into a DWORD unaltered
388         command_history_message(" setting neighbours...\n");
389         control[1] = (DWORD)(current_cell_model->number_of_neighbours);
390         command_history_message(" set\n");
391         if(icCellsInLayer != 0) //no interrupt comes for the first cell
392         //we may have to add a condition for the first layer as well...
393         { command_history_message(" waiting for calculation\n");
394           rc=WC.IntWait(TestInfo.DeviceNum, INT_TIMEOUT_MS); //wait for
            interrupt, signifying current calculation is done
395           if(rc != WC_SUCCESS) { //debug output
396             WC_PeRegRead(TestInfo.DeviceNum, CTRL_REG_BASE, 8, debugReg);
397             command_history_message(" Interrupt timed out; int Counter,
                numNeighbours, dram_addr, dma_incount, state, in_fifo_out,
                DRAM data, DMA_outCount: %x %x %x %x %x %x %x %x\n",
398             debugReg[0], debugReg[1], debugReg[2], debugReg[3], debugReg
                [4], debugReg[5], debugReg[6], debugReg[7]);
399             wildcard_cleanup(TestInfo.DeviceNum, &hDMASource, &
                hDMADestination);
400             return NULL;

```

```

401     }
402     command_history_message("interrupt received\n"); //debugging
           information
403     rc=WC_IntEnable(TestInfo.DeviceNum, FALSE); //reset interrupt
404     command_history_message("interrupt disabled\n");
405     rc=WC_IntReset(TestInfo.DeviceNum);
406     command_history_message("interrupt reset\n");
407     rc=WC_IntEnable(TestInfo.DeviceNum, TRUE);
408     command_history_message("interrupt enabled\n");
409 }
410 if(j==0 && iteration == 1)
411     control[1] = 0x80000000 | (DWORD)(current_cell_model->
           number_of_neighbours); //tell the PE this is not the first
           sample
412
413     command_history_message("control data: %x %x\n", (int)control[0],
           (int)control[1]);
414     command_history_message("sending polarizations\n");
415     //get debugging data
416     WC_PeRegRead(TestInfo.DeviceNum, CTRL_REG_BASE, 8, debugReg);
417     command_history_message("Ek in, read count, write count, gamma
           in, state, in_fifo_out, DRAM fifo data, out_fifo_in: %x %x %
           x %x %x %x %x %x\n",
418         debugReg[0], debugReg[2], debugReg[7], debugReg[1], debugReg
           [4], debugReg[5], debugReg[6], debugReg[3]);
419     rc = WC_PeRegWrite(TestInfo.DeviceNum, CTRL_REG_BASE+3, 2, control
           ); //send gamma and numNeighbours
420     if (rc != WC_SUCCESS) {
421         command_history_message("Register transfer failed\n");
422         wildcard_cleanup(TestInfo.DeviceNum, &hDMASource, &
           hDMADestination);
423         return NULL;
424     }
425
426     command_history_message("waiting for transfer\n");
427     rc=WC_IntWait(TestInfo.DeviceNum, INT_TIMEOUT_MS); //wait for
           interrupt, signifying all polarizations have been received.
           Loop around to begin gathering next set
428     if(rc != WC_SUCCESS) {
429         WC_PeRegRead(TestInfo.DeviceNum, CTRL_REG_BASE, 8, debugReg);
430         command_history_message("Interrupt timed out; int Counter,
           numNeighbours, dram_addr, dma_incount, state, in_fifo_out,
           DRAM data, DMA_outCount: %x %x %x %x %x %x %x %x\n",
431             debugReg[0], debugReg[1], debugReg[2], debugReg[3], debugReg
           [4], debugReg[5], debugReg[6], debugReg[7]);

```

```

432     wildcard_cleanup(TestInfo.DeviceNum, &hDMASource, &
         hDMADestination);
433     return NULL;
434 }
435 command_history_message("interrupt received\n");
436 rc=WC_IntEnable(TestInfo.DeviceNum, FALSE); //reset interrupt
437 command_history_message("interrupt disabled\n");
438 rc=WC_IntReset(TestInfo.DeviceNum);
439 command_history_message("interrupt reset\n");
440 rc=WC_IntEnable(TestInfo.DeviceNum, TRUE);
441 command_history_message("interrupt enabled\n");
442 //go on to next cell (continue with loop):
443 }
444 }
445
446 command_history_message("waiting for calculation (after loop)\n");
447 rc=WC_IntWait(TestInfo.DeviceNum, INT_TIMEOUT_MS); //wait for
         interrupt, signifying last polarization has been calculated
448 if(rc != WCSUCCESS) {
449     WC_PeRegRead(TestInfo.DeviceNum, CTRL_REG_BASE, 8, debugReg);
450     command_history_message("Interrupt timed out; int Counter,
         numNeighbours, dram_addr, dma_incount, state, in_fifo_out,
         DRAM data, DMA_outCount: %x %x %x %x %x %x %x %x\n",
451         debugReg[0], debugReg[1], debugReg[2], debugReg[3], debugReg
         [4], debugReg[5], debugReg[6], debugReg[7]);
452     wildcard_cleanup(TestInfo.DeviceNum, &hDMASource, &
         hDMADestination);
453     return NULL;
454 }
455 rc=WC_IntEnable(TestInfo.DeviceNum, FALSE); //reset interrupt
456 rc=WC_IntReset(TestInfo.DeviceNum);
457 rc=WC_IntEnable(TestInfo.DeviceNum, TRUE);
458 WC_PeRegWrite(TestInfo.DeviceNum, CTRL_REG_BASE+3, 1, &totem); //
         send a totem value to tell the PE that all cells have been
         calculated
459 if (rc != WCSUCCESS) {
460     command_history_message("Register transfer failed\n");
461     wildcard_cleanup(TestInfo.DeviceNum, &hDMASource, &
         hDMADestination);
462     return NULL;
463 }
464
465 calculatedCellNum = 0; //this variable is used to generate the
         index for the relevant data in DMABufDest
466 //since it is a smaller size than number_of_cells_in_layer
467 command_history_message("sorting out results\n");

```

```

468         for (icCellsInLayer = 0 ; icCellsInLayer <
              number_of_cells_in_layer[icLayers] ; icCellsInLayer++)
469     { cell = sorted_cells[icLayers][icCellsInLayer] ;
470       if (!(QCAD_CELLINPUT == cell->cell_function) || (QCAD_CELLFIXED
              == cell->cell_function))
471     { //this condition must be here, as it was for these lines in the
        original version
472       //tolerance testing could be done in hardware fairly easily, but
        doing alongside the data replacement makes sense
473       current_cell_model = ((bistable_model *)cell->cell_model) ;
474       old_polarization = current_cell_model->polarization;
475       current_cell_model->polarization = DMABufDest[calculatedCellNum
              ]; //eliminated need for *(float *)& casting by changing
              buffer's type
476       command_history_message("new polarization: %e\n", DMABufDest[
              calculatedCellNum]);
477       stable = (fabs(DMABufDest[calculatedCellNum] - old_polarization)
              <= tolerance) && stable;
478       DMABufDest[calculatedCellNum] = 1;
479       calculatedCellNum++;
480     }
481   }
482 }
483 }//WHILE !STABLE
484
485 if (VECTOR_TABLE == SIMULATION_TYPE)
486   for (design_bus_layout_iter_first (design_bus_layout, &bli,
        QCAD_CELLINPUT, &i) ; i > -1 ; design_bus_layout_iter_next (&
        bli, &i))
487     if (!exp_array_index_1d (pvt_inputs, VT_INPUT, i).active_flag)
488       sim_data->trace[i].data[j] = ((bistable_model *)
        exp_array_index_1d (pvt_inputs, VT_INPUT, i).input->
        cell_model)->polarization;
489
490 // — collect all the output data from the simulation — //
491 for (design_bus_layout_iter_first (design_bus_layout, &bli,
        QCAD_CELLOUTPUT, &i) ; i > -1 ; design_bus_layout_iter_next (&
        bli, &i))
492   sim_data->trace[design_bus_layout_inputs_icUsed + i].data[j] = ((
        bistable_model *)exp_array_index_1d (design_bus_layout_outputs
        , BUSLAYOUT_CELL, i).cell->cell_model)->polarization;
493
494 // — if the user wants to stop the simulation then exit. — //
495 if(TRUE == STOP_SIMULATION)
496   j = sim_data_number_samples ;
497

```



```

498     }//for number of samples
499
500     // Free the neighbours and Ek array introduced by this simulation//
501     for (k = 0; k < number_of_cell_layers; k++)
502     {
503 #ifdef REDUCE_DEREF
504         number_of_cells_in_current_layer = number_of_cells_in_layer[k] ;
505         for (l = 0 ; l < number_of_cells_in_current_layer ; l++)
506 #else
507         for (l = 0 ; l < number_of_cells_in_layer[k] ; l++)
508 #endif
509         {
510             g_free(((bistable_model *)sorted_cells[k][l]->cell_model)->
                    neighbours);
511             g_free(((bistable_model *)sorted_cells[k][l]->cell_model)->
                    neighbour_layer);
512             g_free(((bistable_model *)sorted_cells[k][l]->cell_model)->Ek);
513         }
514     }
515
516     simulation_inproc_data_free (&number_of_cell_layers , &
                    number_of_cells_in_layer , &sorted_cells) ;
517
518     wildcard_cleanup (TestInfo.DeviceNum, &hDMASource, &hDMADestination);
519
520 // Restore the input flag for the inactive inputs
521 if (VECTOR_TABLE == SIMULATION_TYPE)
522     for (i = 0 ; i < pvt.inputs_icUsed ; i++)
523         exp_array_index_1d (pvt.inputs , VT_INPUT, i).input->cell_function
                    = QCAD_CELLINPUT ;
524
525 // — get and print the total simulation time — //
526 if ((end_time = time (NULL)) < 0)
527     fprintf(stderr, "Could not get end time\n");
528
529 command_history_message ("Total simulation time: %g s\n", (double)(
                    end_time - start_time));
530 set_progress_bar_visible (FALSE) ;
531
532 #ifndef REDUCE_DEREF
533 #undef sim_data_number_samples
534 #undef pvt_inputs
535 #undef pvt_inputs_icUsed
536 #undef pvt_vectors
537 #undef pvt_vectors_icUsed
538 #undef design_bus_layout

```

```
539  #undef design_bus_layout_inputs
540  #undef design_bus_layout_inputs_icUsed
541  #undef design_bus_layout_outputs
542  #undef design_bus_layout_outputs_icUsed
543 #endif
544
545     return sim_data;
546 }//run_bistable_hardware
```

Listing B.1: Bistable simulation engine software code with hardware interaction

# Appendix C

## qcad VHDL Code Listing

The following is the code for the top module of the QCA simulating hardware module.

```
1 library IEEE;
2 use IEEE.std_logic_arith.all;
3 use IEEE.std_logic_1164.all;
4 use IEEE.std_logic_unsigned.all;
5
6 library WILDCARD4LIB;
7 use WILDCARD4LIB.pe_package.all;
8 use WILDCARD4LIB.lad_tools_package.all;
9
10 library UNISIM;
11 use UNISIM.vcomponents.all;
12
13 architecture qcad of pe is
14
15 ---component declarations
16 component p_math
17     Port ( pol : in  STD_LOGIC_VECTOR (31 downto 0);
18           gamma : in  STD_LOGIC_VECTOR (31 downto 0);
19           Ek : in  STD_LOGIC_VECTOR (31 downto 0);
20           start : in  STD_LOGIC;
21           clk : in  STD_LOGIC;
22           ce : in  std_logic;
23           new_cell : in  std_logic;
24           reset : in  STD_LOGIC;
25           new_pol : out  STD_LOGIC_VECTOR (31 downto 0);
26           done : out  STD_LOGIC);
27 end component;
28
29 component fifo
30     port (
31         clk : IN  std_logic;
```

```

32   din: IN std_logic_VECTOR(31 downto 0);
33   rd_en: IN std_logic;
34   rst: IN std_logic;
35   wr_en: IN std_logic;
36   dout: OUT std_logic_VECTOR(31 downto 0);
37   empty: OUT std_logic;
38   full: OUT std_logic);
39 end component;
40
41 component out_fifo
42   port (
43     clk: IN std_logic;
44     din: IN std_logic_VECTOR(31 downto 0);
45     rd_en: IN std_logic;
46     rst: IN std_logic;
47     wr_en: IN std_logic;
48     dout: OUT std_logic_VECTOR(31 downto 0);
49     empty: OUT std_logic;
50     full: OUT std_logic);
51 end component;
52
53 component DRAM_fifo
54   port (
55     din: IN std_logic_VECTOR(127 downto 0);
56     rd_clk: IN std_logic;
57     rd_en: IN std_logic;
58     rst: IN std_logic;
59     wr_clk: IN std_logic;
60     wr_en: IN std_logic;
61     dout: OUT std_logic_VECTOR(127 downto 0);
62     empty: OUT std_logic;
63     full: OUT std_logic;
64     prog_empty: OUT std_logic);
65 end component;
66
67 ---wildcard interface signals
68 constant STARTUP_WAIT      : natural                := 8000;
69                               --- Use for synthesis
70 ---constant STARTUP_WAIT   : natural                := 16;
71                               --- Use for simulation
72 signal reset                : std_logic := '0';
73 signal clocks_in            : clocks_in_type;
74 signal clocks_out           : clocks_out_type;
75 signal lad_in               : lad_in_type;
76 signal lad_out              : lad_out_type;

```

```

76 --data input/output for p_math unit
77 signal Ek_in, gamma_in : std_logic_vector(31 downto 0);
78
79 --control signals for various aspects
80 signal p_math_nd, cell_done, p_math_con, p_math_en : std_logic;
81
82 --state machine for core operation
83 type core_states is (idle, setup_Ek_DMA, Ek_DMA_init1, Ek_DMA_init2,
84   Ek_to_DRAM,
85   next_Ek_DMA, last_Ek_to_DRAM, get_regs, totem_wait, setup_pol_DMA,
86   pol_DMA_init1,
87   pol_DMA_init2, DRAM_in_start, dumbState, calculate1, calculate2,
88   next_cell, nextCell_int,
89   setup_deliver, deliver_init);
90 signal core_state : core_states;
91
92 --DRAM signals
93 signal dram_in : dram_in_type;
94 signal dram_out : dram_out_type;
95 signal dram_read, dram_read_sel : std_logic;
96 signal dram_write : std_logic;
97 signal dram_buf_write, dram_super_write : std_logic; --write control for
98   dram input buffer
99 signal dram_addr : std_logic_vector(24 downto 0);
100 signal dram_write_addr, dram_read_addr : std_logic_vector(22 downto 0);
101 signal dram_data_out : std_logic_vector(127 downto 0);
102 signal dram_data_in_valid : std_logic;
103 signal dram_data_in : std_logic_vector(127 downto 0);
104 signal dram_write_rdy : std_logic;
105 signal dram_read_rdy : std_logic;
106
107 signal dram_fifo_rd_cnt : std_logic_vector(8 downto 0);
108
109 --state machine control signals
110 signal smBus_req, smDMA_init, DMAfromPE : std_logic;
111 signal smLAD_data_out : std_logic_vector(31 downto 0);
112 signal DMA_From_Pe_Addr : std_logic_vector(29 downto 0) := (others
113   => '0');
114 signal DMA_To_Pe_Addr : std_logic_vector(29 downto 0) := (others
115   => '0');
116 signal DMA_inCount, DMA_outCount, numCells_DMA, numNeighbours_DMA :
117   std_logic_vector(31 downto 0);
118 signal DMA_inCount_clr, DMA_outCount_clr : std_logic;
119 signal wait_count : std_logic;
120 signal wait_counter : std_logic_vector(5 downto 0);

```

```

115 signal dram_addr_inc, dram_addr_clr, dram_realAddr_inc : std_logic;
116 signal dram_tailAddr_inc : std_logic;
117 signal pN_delay : std_logic_vector(3 downto 0); --perNeighbour delay
118 signal pN_count, pN_clr : std_logic;
119 signal i_clock_prev : std_logic; --maybe bad practice, but used to
    detect i_clock edges
120 --signal sram_addr_inc, sram_addr_clr : std_logic;
121
122 --other control signals
123 type lad_register_vector is array (natural range<>) of std_logic_vector
    (31 downto 0);
124 signal control_register : lad_register_vector(0 to 7);
125 signal debug_reg : lad_register_vector(0 to 7);
126 signal reg_strobe_out : std_logic;
127 signal reg_LAD_out : std_logic_vector(31 downto 0);
128 signal Reg_Index : natural := 0;
129 signal Start_DMA, gamma_set : std_logic;
130 signal generate_interrupt, pci_rdy, Interrupt, Interrupt_Done :
    std_logic;
131 signal clear_interrupt : std_logic;
132 signal int_counter : std_logic_vector(4 downto 0);
133
134 --controls to and from the fifo units
135 signal in_fifo_rd, in_fifo_empty, in_fifo_full : std_logic;
136 signal in_fifo_out : std_logic_vector(31 downto 0);
137 signal out_fifo_rd, out_fifo_empty, out_fifo_full, out_fifo_prog_full :
    std_logic;
138 signal out_fifo_out, out_fifo_in : std_logic_vector(31 downto 0);
139 signal dram_fifo_rd, dram_fifo_empty, dram_fifo_prog_empty,
    dram_fifo_full : std_logic;
140 signal dram_fifo_clr : std_logic;
141 signal dram_fifo_out : std_logic_vector(127 downto 0);
142
143 alias p_clock          : std_logic is clocks_in.p_clock.clock;
144 alias i_clock          : std_logic is lad_in.clock_in;
145
146 constant REG_BASE      : std_logic_vector(15 downto 0) := x"0100";
147 constant REG_MASK      : std_logic_vector(15 downto 0) := x"FFFF";
148 constant REG_BASE_ADDRESS : std_logic_vector(15 downto 0) := x
    "1000";
149 constant REG_BASE_MASK  : std_logic_vector(15 downto 0) := x"
    FFF8";
150
151 begin
152
153 polarization_math : p_math

```

```

154 Port map( pol => in_fifo_out ,
155     gamma => gamma_in, --from register
156     Ek => Ek_in, --from DRAM buffer thing
157     start => p_math_nd,
158     clk => i_clock ,
159     ce => p_math_en ,
160     new_cell => p_math_con ,
161     reset => reset ,
162     new_pol => out_fifo_in ,
163     done => cell_done);
164
165 in_fifo : fifo --fifo storage for incoming data; size is 32x1024
166     port map (
167         clk => i_clock ,
168         din => lad_in.data_in ,
169         rd_en => in_fifo_rd ,
170         rst => reset ,
171         wr_en => lad_in.DMA_strobe,
172         dout => in_fifo_out , --goes to DRAM buffer thing (sort out which
173             part of the 128 bit word to use)
174         empty => in_fifo_empty ,
175         full => in_fifo_full);
176
177 away : out_fifo --fifo for outgoing data; size is 32x256
178 --size was enlarged from 32x32 to deal with laggy full signals and
179 prevent data overflow
180     port map (
181         clk => i_clock ,
182         din => out_fifo_in ,
183         rd_en => out_fifo_rd ,
184         rst => reset ,
185         wr_en => cell_done ,
186         dout => out_fifo_out ,
187         empty => out_fifo_empty ,
188         full => out_fifo_full);
189
190 out_fifo_rd <= (not out_fifo_empty) and lad_in.bus_gnt and lad_in.
191     pci_rdy;
192
193 DRAM_output_fifo : DRAM_fifo --fifo for DRAM output; size is 128x256
194     port map (
195         din => dram_data_in ,
196         rd_clk => i_clock ,
197         rd_en => dram_fifo_rd , --controlled from state machine when data
198             is needed

```

```

195     rst => dram_fifo_clr, --controlled in the state machine so empty
        flag resets
196     wr_clk => p_clock,
197     wr_en => dram_data_in_valid,
198     dout => dram_fifo_out, --goes to DRAM buffer thing (the thing that
        sorts out which part of the 128 bit word to use)
199     empty => dram_fifo_empty,
200     full => dram_fifo_full,
201     prog_empty => dram_fifo_prog_empty); --set to turn on when less
        than 50 data are present
202     --this is to ensure this fifo communicates effectively with the
        DRAM control
203     --to ensure data does not over- or underflow
204     --because data must always be present to meet timing
205     --and the 173 cycle delay on DRAM read means
206     --there must be a sizable buffer zone between stopping the read
        and the fifo being full
207
208 --DRAM output buffer thing (between DRAM output fifo and p_math input):
209 --governing which part of the DRAM word goes to the p_math core based on
        the address bits
210 --this may need to be modified if we are using two p_math cores
211 Ek_in <= dram_fifo_out(127 downto 96) when (dram_addr(1 downto 0) =
        "11") else
212     dram_fifo_out(95 downto 64) when (dram_addr(1 downto 0) = "10")
        else
213     dram_fifo_out(63 downto 32) when (dram_addr(1 downto 0) = "01")
        else
214     dram_fifo_out(31 downto 0) when (dram_addr(1 downto 0) = "00")
        else
215     (others => '0');
216
217 --dram input buffer thing: DRAM write only happens once at the beginning
        , in sequence,
218 --so every time the last two address bits are 11, the previous three
        data slots have been filled,
219 --and we have a 128 bit word to write.
220 --the last set of data will need to have a special case to be sure it's
        written
221 --this also means that real address 0 will have no data.
222 DRAM_input_buffer: process(i_clock, reset)
223 begin
224     if(reset = '1') then
225         dram_data_out <= (others => '0');
226         dram_write <= '0';
227     else

```



```

228     if rising_edge(i_clock) then
229         if(dram_buf_write = '1') then
230             case dram_addr(1 downto 0) is
231                 when "00" =>
232                     dram_data_out(31 downto 0) <= in_fifo_out;
233                 when "01" =>
234                     dram_data_out(63 downto 32) <= in_fifo_out;
235                 when "10" =>
236                     dram_data_out(95 downto 64) <= in_fifo_out;
237                 when "11" =>
238                     dram_data_out(127 downto 96) <= in_fifo_out;
239                 when others =>
240                     dram_data_out <= dram_data_out;
241             end case;
242
243             if((dram_addr(1 downto 0) = "11") or (dram_super_write = '1'))
244                 then
245                     dram_write <= '1';
246                 else
247                     dram_write <= '0';
248                 end if;
249             else
250                 dram_write <= '0';
251                 dram_data_out <= dram_data_out;
252             end if;
253         end if;
254     end process;
255
256     -----
257     --DMA control stuff
258     -----
259     lad_reg_control_proc: process (reset, i_clock) is
260     begin
261         if (reset = '1') then
262             Start_DMA <= '0';
263             gamma_set <= '0';
264             control_register <= (others => (others => '0'));
265             reg_LAD_out <= (others => '0');
266             reg_strobe_out <= '0';
267         elsif rising_edge (i_clock) then --make sure this is the correct clock
268             to use
269             if (lad_in.reg_strobe = '1') then
270                 if(lad_in.write = '1') then
271                     control_register(Reg_Index) <= lad_in.data_in;
272                     if(Reg_Index = 4) then

```

```

272         Start_DMA <= '1'; --start DMA on third transfer
273     elsif(Reg_Index = 3) then
274         gamma_set <= '1';
275     end if;
276     else
277         reg_LAD_out <= debug_reg(Reg_Index);
278         reg_strobe_out <= '1';
279     end if;
280     else
281         gamma_set <= '0';
282         start_DMA <= '0';
283         reg_strobe_out <= '0';
284         reg_LAD_out <= (others => '0');
285     end if;
286 end if;
287 end process lad_reg_control_proc;
288
289 Reg_Index    <= CONV_INTEGER('0' & lad_in.Addr(2 downto 0));
290
291 DMA_From_PE_Addr <= control_register(0)(31 downto 2);
292 DMA_To_PE_Addr   <= control_register(1)(31 downto 2);
293 numCells_DMA <= control_register(2); --DMA from PE count; the number of
      cells
294 gamma_in    <= control_register(3);
295 numNeighbours_DMA <= X"7fffffff" and control_register(4); --DMA to PE
      count; the number of neighbours a cell has
296 --the gamma and neighbours registers are arranged in this order to
      streamline the transfer
297 --of gamma followed by the number of neighbours for starting a
      polarization transfer
298 --ANDed with this value because of dumb things. (see comment for
      dumbState in the transitions process)
299
300 --debug_reg(2) <= dram_fifo_empty & control_register(4)(31) & gamma_set
      & start_DMA & "000" & dram_addr;
301 debug_reg(4) <= X"00000000" when core_state = idle else
302     X"00000001" when core_state = setup_Ek_DMA else
303     X"00000002" when core_state = Ek_DMA_init1 else
304     X"00000003" when core_state = Ek_DMA_init2 else
305     X"00000004" when core_state = Ek_to_DRAM else
306     X"00000005" when core_state = next_Ek_DMA else
307     X"00000006" when core_state = last_Ek_to_DRAM else
308     X"00000007" when core_state = get_regs else
309     X"00000008" when core_state = totem_wait else
310     X"00000009" when core_state = setup_pol_DMA else
311     X"0000000a" when core_state = pol_DMA_init1 else

```

```

312 X"0000000b" when core_state = pol_DMA_init2 else
313 X"0000000c" when core_state = DRAM_in_start else
314 X"0000000d" when core_state = dumbState else
315 X"0000000e" when core_state = calculate1 else
316 X"0000000f" when core_state = calculate2 else
317 X"00000010" when core_state = next_cell else
318 X"00000011" when core_state = nextCell_int else
319 X"00000012" when core_state = setup_deliver else
320 X"00000013" when core_state = deliver_init else
321 X" ffffffff ";
322
323 debug_out : process(reset , i_clock) is
324 begin
325     if (reset = '1') then
326         debug_reg(0) <= (others => '0');
327         debug_reg(1) <= (others => '0');
328         debug_reg(5) <= (others => '0');
329         debug_reg(3) <= (others => '0');
330         debug_reg(2) <= (others => '0');
331         debug_reg(7) <= (others => '0');
332
333         dram_fifo_rd_cnt <= (others => '0');
334     elsif(rising_edge(i_clock)) then
335         if(core_state = calculate1) then
336             debug_reg(0) <= Ek_in;
337             debug_reg(2) <= "000000000000000000000000" & dram_fifo_rd_cnt;
338             debug_reg(7) <= (others => '0');
339             debug_reg(1) <= gamma_in;
340             debug_reg(5) <= in_fifo_out;
341             debug_reg(6) <= dram_fifo_out(127 downto 96);
342         end if;
343
344         if(core_state = nextCell_int) then
345             debug_reg(3) <= out_fifo_in;
346         end if;
347
348         if(dram_fifo_rd = '1') then
349             dram_fifo_rd_cnt <= dram_fifo_rd_cnt + 1;
350         elsif(dram_fifo_clr = '1') then
351             dram_fifo_rd_cnt <= (others => '0');
352         end if;
353     end if;
354 end process;
355
356 --state machine transitions
357 core_transitions : process (i_clock , reset)

```

```

358 begin
359   if(reset = '1') then
360     core_state <= idle;
361   elsif rising_edge(i_clock) then
362     case core_state is
363     when idle => --this state should only happen once, when the core is
                 reset
364       --wait for signal to start DMA transfer of Ek data
365       if(Start_DMA = '1') then
366         core_state <= setup_Ek_DMA;
367       end if;
368     when setup_Ek_DMA => --this is the state we should loop back to for
                 repeating the DMA transfer
369       if (lad_in.bus_gnt = '1') then
370         core_state <= Ek_DMA_init1;
371       end if;
372     when Ek_DMA_init1 =>
373       core_state <= Ek_DMA_init2;
374     when Ek_DMA_init2 =>
375       core_state <= Ek_to_DRAM;
376     when Ek_to_DRAM =>
377       --shove Ek values from the FIFO into DRAM
378       --Max address limit is irrelevant because the amount of data is
                 never that high
379       --get number of neighbours from the host in the RegWrite that
                 triggers things
380       --no check for dram_write_rdy here because in_fifo is read only
                 when that signal is high
381       --so in_fifo_empty will activate after the complete set of *
                 successful* DRAM writes
382       if((DMA_inCount >= numNeighbours_DMA) and (in_fifo_empty = '1'))
                 then
383         core_state <= next_Ek_DMA; --start new DMA for more Ek data...
384       end if;
385     when next_Ek_DMA =>
386       --will need to send an interrupt here
387       if(Start_DMA = '1') then
388         core_state <= setup_Ek_DMA;
389       elsif(gamma_set = '1') then
390         core_state <= last_Ek_to_DRAM;
391       end if;
392     when last_Ek_to_DRAM =>
393       if(dram_write_rdy = '1') then --if dram_write_rdy is 1, then the
                 thing has happened and we can move on
394         core_state <= setup_deliver;
395       end if;

```

```

396   when setup_deliver =>
397   --open the channel to deliver data back to host (DMA from the SRAM)
398     if(lad_in.bus_gnt = '1') then
399       core_state <= deliver_init;
400     end if;
401   when deliver_init =>
402     core_state <= totem_wait;
403   when totem_wait => --a wait state for between iterations
404     if(gamma_in /= X"BF000000" and wait_counter >= X"4") then
405       core_state <= get_regs;
406     end if;
407   when get_regs => --get gamma and neighbour values from host for
      current cell in register transfer
408   --this is the state we start from every time after we have Ek data
409   --wait until we have gamma and neighbours
410     if(DMA_outCount = X"00000000" or start_DMA = '1') then
411       core_state <= setup_pol_DMA;
412     elsif(gamma_set = '1' and gamma_in = X"BF000000") then
413       --test here for whether the last cell is done or not
414       --test must be done in the first transfer or it will cause a false
      positive
415       core_state <= setup_deliver;
416     end if;
417   when setup_pol_DMA => --this is for v1 of the hardware
      --wait for the bus grant, like in the Ek DMA setup
418     if (lad_in.bus_gnt = '1') then
419       core_state <= Pol_DMA_init1;
420     end if;
421   --it would be nice to keep fetching more polarization data as soon
      as all the current stuff goes in
422   --but I think it would require extracting the DMA control into a
      seperate machine
423   when pol_DMA_init1 =>
424     core_state <= pol_DMA_init2;
425   when pol_DMA_init2 =>
426     core_state <= DRAM_in_start;
427   when DRAM_in_start => --we have all Ek, gamma, and Polarizations are
      on their way, so we can start calculating the new polarization
428   --this state should request data from DRAM, then move to a wait
      state
429   --DRAM address in this and subsequent states should be incrementing
      continuously until some limit...
430   if(dram_fifo_empty = '0' and interrupt_done = '1') then
431     --make sure DRAM data is ready and that polarization data has come
432     through the fifo

```

```

433     --the particulars of the logic at play here mean that the in_fifo
         must be at least
434     --as large as the largest DMA input transaction
435     if(control_register(4)(31) = '1' or DMA_outCount /= X"0") then
436     --do dumb things on the first cell of any sample but the first.
437         core_state <= calculate1;
438     else
439         core_state <= dumbState;
440     end if;
441 end if;
442 when dumbState =>
443     --there is a dumb thing that happens where the DRAM doesn't actually
         read the address it's told to read...
444     --Upon doing read(1) on anything but the first sample, the first
         data_in_valid would be associated
445     --with all zeroes on the data line (presumably read(0), but I don't
         know).
446     --in order to compensate for this dumb thing that happens, this dumb
         state was added
447     --to flush that first worthless word out of the dram_fifo when
         needed.
448     --The need to do this is indicated by having a zero in the MSB of
         control_register(4) when sent
449     --from the host. This means also that the MSB of control_register(4)
         (which recieves numNeighbours)
450     --must be anded with 7fffffff in order to make numNeighbours a
         normal value for use in
451     --DMA initialization
452     --Anyway this is dumb. I don't know why it needs to happen like this
         , I can't explain the simulated
453     --behaviour of the DRAM, and it seems to happen the same in hardware
         ...
454     --it may have something to do with the write.
455     core_state <= calculate1;
456     when calculate1 =>
457     --first data ready from DRAM fifo , start calculation
458     --if(dram_fifo_empty = '0' and in_fifo_empty = '0') then --make
         sure data is received before moving on
459     --the functionality of the p_math core is dependent upon precise
         and reliable timing
460     --ZERE CAN BE NO VAITINK!
461     --(*ahem* there can be no waiting) Although extra waiting around
         between cells is acceptable
462     --the question of the moment is: is it safe to assume the DRAM
         will reliably have new data ready?
463     core_state <= calculate2;

```

```

464     --end if;
465     when calculate2 =>
466     --this is the wait state; it should have two cycles
467     if(wait_counter = X"1") then
468     --how to determine when to move on to the next cell...
469     if(in_fifo_empty = '1') then
470     core_state <= next_cell;
471     else
472     core_state <= calculate1;
473     end if;
474     end if;
475     when next_cell =>
476     --p.math core prepped for next stuff...
477     --new_pol from p.math should be automatically sent when core is
478     --rather than being controlled explicitly here
479     if(cell_done = '1') then --wait around until calculation is done
480     core_state <= nextCell_int;
481     end if;
482     when nextCell_int => --state for generating interrupt
483     if(interrupt_done = '1' and start_DMA = '0') then
484     core_state <= get_regs;
485     end if;
486     when others =>
487     core_state <= idle;
488     end case;
489     end if;
490 end process;
491
492 core_control : process (core_state , DMA_To_Pe_Addr , numNeighbours_DMA ,
493     dram_write_rdy , in_fifo_empty , dram_addr , dram_read_rdy ,
494     dram_fifo_full ,
495     DMA_inCount , DMA_From_Pe_Addr , DMA_outCount , numCells_DMA ,
496     out_fifo_full ,
497     out_fifo_empty , cell_done , dram_fifo_prog_empty)
498 begin
499     dram_addr_inc <= '0';
500     dram_realAddr_inc <= '0';
501     dram_tailAddr_inc <= '0';
502     dram_addr_clr <= '0';
503     dram_buf_write <= '0';
504     dram_super_write <= '0';
505     dram_read <= '0';
506     dram_fifo_rd <= '0';
507     dram_fifo_clr <= '0';
508     p_math_nd <= '0';

```

```

507 p_math_en <= '0';
508 p_math_con <= '0';
509 smbus_req <= '0';
510 smdma_init <= '0';
511 smLAD_Data_out <= (others => '0');
512 DMA_inCount_clr <= '0';
513 DMA_outCount_clr <= '0';
514 in_fifo_rd <= '0';
515 pn_clr <= '0';
516 pn_count <= '0';
517 generate_interrupt <= '0';
518 clear_interrupt <= '0';
519 wait_count <= '0';
520 dram_read_sel <= '0';
521
522 case core_state is
523 when idle =>
524     clear_interrupt <= '1';
525     DMA_inCount_clr <= '1';
526     DMA_outCount_clr <= '1';
527     dram_fifo_clr <= '1';
528 when setup_Ek_DMA =>
529     smBus_Req <= '1'; --request LAD bus for DMA initialization
530 when Ek_DMA_init1 =>
531     smBus_Req           <= '1';
532     smdma_init          <= '1';
533     smLAD_Data_Out     <= DMA_TO_PE_Addr & "00";
534 when Ek_DMA_init2 =>
535     smbus_req <= '1';
536     smdma_init <= '1';
537     smLAD_Data_Out <= numNeighbours_DMA;
538     DMA_inCount_clr <= '1';
539     clear_interrupt <= '1';
540 when Ek_to_DRAM => --shove Ek values from the FIFO into DRAM
541 --we could extract this out to be an autonomous control structure, so
542 --that it could continue
543 --running while new data is being fetched, but "in theory" new data
544 --should be stored immediately
545 --so there should be no waiting
546 --data has to be written 128 bits at a time, this is handled by the
547 --DRAM input buffer
548 if (DRAM_write_rdy = '1' and in_fifo_empty = '0') then
549     dram_buf_write <= '1'; --control sent to the DRAM input buffer
550     in_fifo_rd <= '1';
551 end if;
552 when next_Ek_DMA => --wait for next totem, and fire interrupt

```



```

550     generate_interrupt <= '1';
551 when last_Ek_to_DRAM => --this is for making sure the last set of data
    in the DRAM input buffer gets written
552 --because if it does not end on a "11" it will otherwise just stay in
    the buffer and not get written
553     if(dram_addr(1 downto 0) /= "00" and DRAM_write_rdy = '1') then
554     --indicating that the last address was not ending in "11"
555         dram_super_write <= '1';
556         dram_buf_write <= '1';
557         dram_realAddr_inc <= '1';
558         in_fifo_rd <= '1';
559     end if;
560 when setup_deliver =>
561 --deliver data back to host (DMA from the SRAM)
562     smBus_req <= '1';
563     clear_interrupt <= '1';
564     dram_read_sel <= '1';
565 when deliver_init =>
566     smBus_Req           <= '1';
567     smdma_init         <= '1';
568     smLAD_Data_Out    <= DMA_From_PE_Addr & "01";
569     dram_read_sel <= '1';
570 when totem_wait =>
571     wait_count <= '1';
572     DMA_outCount_clr <= '1';
573     dram_fifo_clr <= '1';
574     dram_addr_clr <= '1'; --all data has been put in DRAM, reset the
        address for reading
575     dram_read_sel <= '1';
576 when get_regs =>
577     p_math_en <= '1';
578     p_math_con <= '1';
579     dram_read_sel <= '1';
580 when setup_pol_DMA =>
581     smBus_Req <= '1'; --request LAD bus for DMA initialization
582     p_math_en <= '1';
583     dram_read_sel <= '1';
584 when pol_DMA_init1 =>
585     smBus_Req           <= '1';
586     smdma_init         <= '1';
587     smLAD_Data_Out    <= DMA_TO_PE_Addr & "00";
588     p_math_en <= '1';
589     dram_read_sel <= '1';
590 when pol_DMA_init2 =>
591     smbusr_req <= '1';
592     smdma_init <= '1';

```

```

593     smLAD_Data_Out <= numNeighbours_DMA;
594     DMA_inCount_clr <= '1';
595     p_math_en <= '1';
596     clear_interrupt <= '1';
597     dram_read_sel <= '1';
598     when DRAM_in_start => —we have all Ek, gamma, and Polarizations, so
        we can start calculating the new polarization
599     —this state should request data from DRAM, then move to a wait state
600     —DRAM address in this and subsequent states should be incrementing
        continuously until some limit...
601     if (DRAM_read_rdy = '1' and dram_fifo_prog_empty = '1') then
602         dram_read <= '1';
603     end if;
604     p_math_en <= '1';
605
606     —fire the interrupt here, allowing the host to gather the next set
        of polarizations
607     —while the calculation is being done.
608     —need to fire interrupt again when cell calculation is done to tell
        host
609     —it is ok to send new gamma and neighbours
610     if (DMA_inCount >= numNeighbours_DMA) then —this condition is used
        in a few places...
611         generate_interrupt <= '1';
612     end if;
613     dram_read_sel <= '1';
614     when dumbState =>
615         if (DRAM_read_rdy = '1' and dram_fifo_prog_empty = '1') then
616             dram_read <= '1';
617         end if;
618         dram_fifo_rd <= '1';
619         p_math_en <= '1';
620         dram_read_sel <= '1';
621     when calculate1 =>
622     —first data ready from DRAM fifo, start calculation
623     if (DRAM_read_rdy = '1' and dram_fifo_prog_empty = '1') then —be
        careful with this read_rdy signal and p_math_nd...
624         dram_read <= '1';
625     end if;
626     p_math_nd <= '1';
627     p_math_en <= '1';
628     dram_tailAddr_inc <= '1';
629     in_fifo_rd <= '1';
630     if (dram_addr(1 downto 0) = "11") then —when the fourth 32 bit word
        has been read, get next 128 bit word
631         dram_fifo_rd <= '1';

```

```

632     end if;
633     dram_read_sel <= '1';
634 when calculate2 =>
635   —currently the p_math unit is configured for a 2–latency accumulator
        ...
636   —this is the downcycle
637     if (DRAM_read_rdy = '1' and dram_fifo_prog_empty = '1') then —be
        careful with this read_rdy signal and p_math.nd....
638       dram_read      <= '1';
639     end if;
640     p_math_en <= '1';
641     wait_count <= '1';
642     dram_read_sel <= '1';
643 when next_cell =>
644     if (DRAM_read_rdy = '1' and dram_fifo_prog_empty = '1') then —be
        careful with this read_rdy signal and p_math.nd....
645       dram_read      <= '1';
646     end if;
647     p_math_en <= '1';
648     clear_interrupt <= '1';
649     dram_read_sel <= '1';
650 when nextCell_int => —interrupt to inform that the cell has been
        calculated
651     if (DRAM_read_rdy = '1' and dram_fifo_prog_empty = '1') then —this
        is just running all the time to make sure data is present
652     —stopping the DRAM transfer only after fifo_full is detected, may
        cause data loss...
653     —BECAUSE TIMING IS SUPER CRITICAL WE MUST HAVE DATA
654     dram_read      <= '1';
655     end if;
656     generate_interrupt <= '1';
657     p_math_en <= '1';
658     dram_read_sel <= '1';
659   end case;
660 end process;
661
662 process (reset, p_clock, i_clock) is
663 begin
664   if(reset = '1') then
665     dram_read_addr <= (others => '0');
666     dram_write_addr <= (others => '0');
667     dram_addr(1 downto 0) <= "00";
668   else
669     if rising_edge(p_clock) then
670       if(dram_read = '1') then
671         dram_read_addr <= dram_read_addr + 1;

```

```

672     end if;
673
674     if(dram_addr_clr = '1') then
675         --this is the address where the first data is, and it is only
           set after DRAM is filled
676         --because the nature of how the DRAM input buffer works means
           that address zero is not written
677         --so we just skip over it when we go to reads.
678         --(the global reset is enough to set the initial writing address
           to 0.)
679         dram_read_addr <= "000000000000000000000001";
680     end if;
681 end if;
682
683 if rising_edge(i_clock) then
684     if (dram_buf_write = '1' and dram_addr(1 downto 0) = "11") then
685         dram_write_addr <= dram_write_addr + 1;
686     end if;
687
688     if(dram_addr_clr = '1') then
689         dram_write_addr <= "000000000000000000000001";
690     end if;
691
692     --it is necessary to separate these two out at some point, because
           while streaming the DRAM output
693     --we must keep the real address going up as fast as possible to
           ensure presense of data,
694     --in order to keep up with the timing concerns of the p_math core.
695     --but then the two LSBs get out of sync from where we want them,
           which is pointing to
696     --the next data coming out of the DRAM fifo
697     if(dram_tailAddr_inc = '1' or dram_buf_write = '1') then
698         dram_addr(1 downto 0) <= dram_addr(1 downto 0) + 1;
699     end if;
700
701     if(dram_addr_clr = '1') then
702         dram_addr(1 downto 0) <= "00";
703     end if;
704 end if;
705 end if;
706 end process;
707
708 --this is so we can have the address controlled by two different clocks
709 dram_addr(24 downto 2) <= dram_read_addr when dram_read_sel = '1' else
710     dram_write_addr;
711

```

```

712 --be wary of clocks....
713 wait_counter_proc: process (reset, i_clock) is
714 begin
715     if (reset = '1') then
716         DMA_inCount <= (others => '0');
717         DMA_outCount <= (others => '0');
718         wait_counter <= (others => '0');
719     elsif rising_edge (i_clock) then
720         --this is synchronized to i_clock because they are tracking DMA
           transactions over the LAD bus
721         if (lad_in.DMA_strobe = '1') then
722             DMA_inCount <= DMA_inCount + 1;
723         elsif (DMA_inCount_clr = '1') then
724             DMA_inCount <= (others => '0');
725         end if;
726
727         if (DMAfromPE = '1') then
728             DMA_outCount <= DMA_outCount + 1;
729         elsif (DMA_outCount_clr = '1') then
730             DMA_outCount <= (others => '0');
731         end if;
732
733         if (wait_count = '1') then
734             wait_counter <= wait_counter + 1;
735         else
736             wait_counter <= (others => '0');
737         end if;
738
739         DMAfromPE <= out_fifo_rd; --similar structure as the DMA example for
           outputting data
740     end if;
741 end process wait_counter_proc;
742
743 -----
744 --DRAM signal registration
745 -----
746 register_signals_proc: process (p_clock, reset) is
747 begin
748     if (reset = '1') then
749         dram_out.read      <= '0';
750         dram_out.write    <= '0';
751         dram_out.addr     <= (others => '0');
752         dram_out.data_out <= (others => '0');
753         dram_data_in      <= (others => '0');
754         dram_data_in_valid <= '0';
755         dram_write_rdy    <= '0';

```

```

756     dram_read_rdy         <= '0';
757   elsif rising_edge (p_clock) then
758     dram_out.read         <= dram_read;
759     dram_out.write        <= dram_write;
760     dram_out.addr         <= dram_addr(24 downto 2); --lower two bits are
        used to decide which 32 bit word to select
761     dram_out.data_out     <= dram_data_out;
762     dram_data_in          <= dram_in.data_in;
763     dram_data_in_valid    <= dram_in.data_in_valid;
764     dram_write_rdy        <= dram_in.write_rdy;
765     dram_read_rdy         <= dram_in.read_rdy;
766   end if;
767 end process register_signals_proc;
768
769 lad_out_proc: process (reset, i_clock) is
770 begin
771   if (reset = '1') then
772     lad_out.Data_Out       <= (others => '0');
773     lad_out.Strobe_Out     <= '0';
774     lad_out.dma_init       <= '0';
775     lad_out.bus_req        <= '0';
776     lad_out.pe_rdy         <= '0';
777     lad_out.int_req        <= '0';
778     pci_rdy                <= '0';
779   elsif rising_edge (i_clock) then
780     if (DMAFromPe = '1') then --DMAfromPe will need to be controlled
        somewhere, figure that out
781       lad_out.Data_Out <= out_fifo_out; --SRAM data output goes here,
        somehow
782     elsif (reg_strobe_out = '1') then
783       lad_out.Data_out <= Reg.LAD_out;
784     else
785       lad_out.Data_Out <= smLAD_Data_Out;
786     end if;
787
788     -- Generate a strobe on the LAD bus for any of the following 2 cases
        : --
789     -- Regs_Strobe_Out = '1' : Register read from the control register
        --
790     -- DMAFromPe = '1'           : DMA-From-PE transaction data strobe
        --
791
792     lad_out.Strobe_Out <= DMAFromPe or Reg_strobe_out;
793
794     -- Allow the state machine to drive the dma_init line when
        necessary. --

```

```

795 |-----
796 |   lad_out.dma_init <= smdma_init;
797 |
798 |   lad_out.bus_req <= ((not out_fifo_empty) and lad_in.pci_rdy) or
      |   DMAFromPE or smBus_Req;
799 |
800 |   -- In addition to performing the DMA-to-PE initialization
      |   transactions --
801 |   -- the PE also must inform the PCI controller when it is able to
      |   accept --
802 |   -- DMA data. This is done through the lad_out.pe_rdy line.
803 |   lad_out.pe_rdy <= not in_fifo_full;
804 |
805 |   lad_out.int_req <= Interrupt;
806 |
807 |   -- Delay PCI ready by one clock to make better timing.
808 |   pci_rdy <= lad_in.pci_rdy;
809 |   end if;
810 | end process lad_out_proc;
811 |
812 |-----
813 | --Interrupt Control
814 |-----
815 | int_proc: process(i_clock , reset)
816 | begin
817 |   if (reset = '1') then
818 |     Int_Counter <= (others => '0');
819 |   elsif rising_edge (i_clock) then
820 |     if (Generate_Interrupt = '1' and Interrupt_Done = '0' and lad_in.
      |       dma_in_progress = '0') then
821 |       Int_Counter <= Int_Counter + 1;
822 |     end if;
823 |
824 |     if (clear_interrupt = '1') then
825 |       int_counter <= (others => '0');
826 |     end if;
827 |   end if;
828 | end process int_proc;
829 |
830 | Interrupt_Done <= Int_Counter(4);
831 | Interrupt <= Int_Counter(4) or Int_Counter(3) or Int_Counter(2) or
      |   Int_Counter(1) or Int_Counter(0);
832 |
833 | -- Tie off unused Clock signals
834 | clocks_out.p_clock.reset <= '0';
835 |

```

```

836 -----
837 --instantiate control interfaces
838 -----
839 dram_interface_inst: dram_interface
840 generic map (
841   STARTUP_WAIT    => STARTUP_WAIT
842 ) port map (
843   pads             => pads.dram,
844   reset            => reset,
845   p_clock          => clocks_in.p_clock,
846   lad_clock        => i_clock,
847   user_in          => dram_in,
848   user_out         => dram_out
849 );
850
851 default_interface_inst: default_interface
852 port map (
853   pads             => pads.default,
854   reset_out        => reset,
855   clocks_in        => clocks_in,
856   clocks_out       => clocks_out,
857   lad_in           => lad_in,
858   lad_out          => lad_out
859 );
860 -----
861 --end instantiate control interfaces
862 -----
863
864 -- Do not delete/modify this line
865 pads    <= init_pe_pads;
866
867 end qcad;

```

Listing C.1: QCA hardware simulator code



## Appendix D

### p\_math VHDL Code Listing

The listing that follows is the VHDL code for the main mathematical core.

```

1  -- use instructions: first cycle - turn on ce, provide gamma value,
   -- provide Ek and Pol values, pulse start signal for one cycle
2  --           leave start signal off for one cycle
3  --           provide next Ek and Pol values and pulse start signal for
   -- one cycle
4  --           leave start signal off for one cycle
5  --           repeat until all Ek and Pol values have been sent in
6  --           after pulsing start signal for the last time,
7  --           wait for (mult + add + 1) latency, then pulse new_cell to
   -- start a new cell
8  --           repeat instructions from first cycle
9  --           wait until done signal activates, retrieve data.
10
11 library IEEE;
12 use IEEE.STD_LOGIC_1164.ALL;
13 use IEEE.STD_LOGIC_ARITH.ALL;
14 use IEEE.STD_LOGIC_UNSIGNED.ALL;
15
16 entity p_math is
17     Port ( pol : in  STD_LOGIC_VECTOR (31 downto 0); --parallel data
           -- inputs for polarization and Ek
18           gamma : in  STD_LOGIC_VECTOR (31 downto 0);
19           Ek : in  STD_LOGIC_VECTOR (31 downto 0);
20           start : in  STD_LOGIC; --pulse this with each new set of data
           -- input
21           clk : in  STD_LOGIC;
22           ce : in  std_logic; --clock enable, stays on until "done"
           -- triggers
23           new_cell : in  std_logic; --pulse this when ready to send in new
           -- train of cell data
24           reset : in  STD_LOGIC;

```

```

25         new_pol : out  STD_LOGIC_VECTOR (31 downto 0); --result
           output
26         done : out  STD_LOGIC); --computation complete signal
27 end p_math;
28
29 architecture Behavioral of p_math is
30
31 --some of these floating point units could have latencies reduced to
           free up hardware
32 component multiply
33     port (
34         a: IN std_logic_VECTOR(31 downto 0);
35         b: IN std_logic_VECTOR(31 downto 0);
36         operation_nd: IN std_logic;
37         clk: IN std_logic;
38         sclr: IN std_logic;
39         ce: IN std_logic;
40         result: OUT std_logic_VECTOR(31 downto 0);
41         rdy: OUT std_logic);
42 end component;
43
44 component add
45     port (
46         a: IN std_logic_VECTOR(31 downto 0);
47         b: IN std_logic_VECTOR(31 downto 0);
48         operation_nd: IN std_logic;
49         clk: IN std_logic;
50         sclr: IN std_logic;
51         ce: IN std_logic;
52         result: OUT std_logic_VECTOR(31 downto 0);
53         rdy: OUT std_logic);
54 end component;
55
56 component fastadder --floating point adder with latency of 1 for low-
           latency accumulation
57     port (
58         a: IN std_logic_VECTOR(31 downto 0);
59         b: IN std_logic_VECTOR(31 downto 0);
60         operation_nd: IN std_logic;
61         clk: IN std_logic;
62         sclr: IN std_logic;
63         ce: IN std_logic;
64         result: OUT std_logic_VECTOR(31 downto 0);
65         rdy: OUT std_logic);
66 end component;
67

```

```

68 component sqrt
69   port (
70     a: IN std_logic_VECTOR(31 downto 0);
71     operation_nd: IN std_logic;
72     clk: IN std_logic;
73     sclr: IN std_logic;
74     ce: IN std_logic;
75     result: OUT std_logic_VECTOR(31 downto 0);
76     rdy: OUT std_logic);
77 end component;
78
79 component divider
80   port (
81     a: IN std_logic_VECTOR(31 downto 0);
82     b: IN std_logic_VECTOR(31 downto 0);
83     operation_nd: IN std_logic;
84     clk: IN std_logic;
85     sclr: IN std_logic;
86     ce: IN std_logic;
87     result: OUT std_logic_VECTOR(31 downto 0);
88     rdy: OUT std_logic);
89 end component;
90
91 component shift_reg --32 bit shift register for use as a delay line
92   Generic ( depth : integer := 50);
93   Port ( data_in : in  STD_LOGIC_VECTOR (31 downto 0);
94         data_out : out STD_LOGIC_VECTOR (31 downto 0);
95         clk : in  STD_LOGIC);
96 end component;
97
98 component onebit_shift --1 bit delay line
99   Generic ( depth : integer := 10);
100   Port ( data_in : in  STD_LOGIC;
101         data_out : out STD_LOGIC;
102         clk : in  STD_LOGIC);
103 end component;
104
105 signal PxEk, sum : std_logic_vector(31 downto 0);
106 signal two_gamma, pol_math, p_math_squared : std_logic_vector(31 downto
107   0);
107 signal pol_math_delayed : std_logic_vector(31 downto 0);
108 signal square_plus_1, SoP, den : std_logic_vector(31 downto 0);
109 signal in_counter, calc_counter : std_logic_vector(9 downto 0);
110
111 signal add_done, pol_math_rdy, square_done, plus1_rdy : std_logic;
112 signal den_rdy, accum_rdy, initial, mpy_done : std_logic;

```

```

113
114 begin
115
116 PxEk_mult : multiply --this latency must be greater than adder+1 in
    order for everything to work
117     port map (
118         a => pol, --pol and Ek from main inputs
119         b => Ek,
120         operation_nd => start, --signal from main inputs, should be pulsed
            with each new pol/Ek
121         clk => clk,
122         sclr => reset,
123         ce => ce,
124         result => PxEk,
125         rdy => mpy_done);
126
127 --adder latency is reduced to 1 cycles
128 accum : add --accumulator, stage 1b (multiplier above is stage 1a)
129     port map (
130         a => PxEk,
131         b => SoP,
132         operation_nd => mpy_done,
133         clk => clk,
134         sclr => reset,
135         ce => ce,
136         result => sum,
137         rdy => add_done);
138
139 process(clk, reset, new_cell)
140 begin
141
142 --here the accumulation process is handled by careful timing and control
143 if(reset = '1' or new_cell = '1') then
144     Sop <= (others => '0');
145     calc_counter <= (others => '0');
146     initial <= '1'; --this signal is an indicator to whether we are at the
        beginning
147
148     accum_rdy <= '0';
149 elseif(clk'event and clk = '1') then
150     if(initial = '1' and mpy_done = '1') then
151         initial <= '0';
152     elseif(add_done = '1') then
153         SoP <= sum; --running total updated manually, easier to manage than
            a simple loop
154         calc_counter <= calc_counter + 1;

```

```

155     elsif ((calc_counter = in_counter) and (initial = '0')) then
156         --needed to keep accum_rdy from going into a permanent high
157         calc_counter <= calc_counter + 1;
158     end if;
159
160     --this is where "initial" is helpful, because the two counters are
161     also equal at the beginning
162     if ((calc_counter = in_counter) and (initial = '0')) then
163         accum_rdy <= '1';
164     else
165         accum_rdy <= '0';
166     end if;
167 end if;
168 end process;
169
170 --this process controls "in_counter" which counts how many sets of pol/
171 Ek come in
172 --"in_counter" is compared to "calc_counter", controlled above, to
173 determine when
174 --the accumulation is finished.
175 process(clk, reset, new_cell)
176 begin
177     if(reset = '1' or new_cell = '1') then
178         in_counter <= (others => '0');
179     elsif(clk'event and clk = '1') then
180         if(start = '1') then
181             in_counter <= in_counter + 1;
182         end if;
183     end if;
184 end if;
185
186 end process;
187
188 --trick to double a floating point number, to avoid using an extra
189 multiplier
190 --this method will only cause issue when gamma > 1.7e38, which it never
191 is.
192 --saves about 200 slices
193 twoGamma : process(clk, reset) --gamma doubling, stage 1c
194 begin
195     if (reset = '1') then
196         two_gamma <= (others => '0');
197     elsif (clk'event and clk = '1') then
198         if (start = '1') then --add 1 to exponent

```

```

195     two_gamma <= gamma(31) & (gamma(30 downto 23)+1) & gamma(22 downto
        0);
196     end if;
197     end if;
198 end process;
199
200 divvy : divider --sum of products divided by gamma, stage 2
201     port map (
202         a => SoP, --result from accumulation
203         b => two_gamma,
204         operation_nd => accum_rdy, --signal from accumulator control to
            say it's done
205         clk => clk,
206         sclr => reset,
207         ce => ce,
208         result => pol_math,
209         rdy => pol_math_rdy);
210
211 pol_mathDelay : shift_reg --delay line from p_math to final divider
212     Generic map( depth => 6) --depth is the combined latency of parts on
        other path (multiplier, adder, sqrt)
213     Port map( data_in => pol_math,
214             data_out => pol_math_delayed,
215             clk => clk);
216
217 square : multiply --square the divider result, stage 3
218     port map (
219         a => pol_math,
220         b => pol_math,
221         operation_nd => pol_math_rdy,
222         clk => clk,
223         sclr => reset,
224         ce => ce,
225         result => p_math_squared,
226         rdy => square_done);
227
228 plus_one : add --add one to squared pol_math, stage 4
229     port map (
230         a => p_math_squared,
231         b => X"3F800000", --floating point 1
232         operation_nd => square_done, --_nd lines dominoing from rdy lines
            ...
233         clk => clk,
234         sclr => reset,
235         ce => ce,
236         result => square_plus_1,

```

```
237     rdy => plus1_rdy);
238
239 make_den : sqrt ---square root function for denominator, stage 5
240     port map (
241         a => square_plus_1 ,
242         operation_nd => plus1_rdy ,
243         clk => clk ,
244         sclr => reset ,
245         ce => ce ,
246         result => den ,
247         rdy => den_rdy);
248
249 final : divider ---give final result , stage 6, depends on results from
      sqrt (stage 5) and divvy (stage 2)
250     port map (
251         a => pol_math_delayed ,
252         b => den ,
253         operation_nd => den_rdy ,
254         clk => clk ,
255         sclr => reset ,
256         ce => ce ,
257         result => new_pol ,
258         rdy => done); ---connected to main output
259
260 end Behavioral;
```

Listing D.1: Polarization math arithmetic core hardware code