Computational Modelling of Treg Networks in Experimental Autoimmune Encephalomyelitis

Richard Brian Greaves

Submitted for the degree of Master of Science by Research.

University of York
Department of Computer Science

September 2011

Abstract

Recent experiments have demonstrated the value of rigorously validated simulation in furthering our understanding of immunology. We seek to expand on a body of existing work at York in which we simulate the mouse disease, Experimental Autoimmune Encephalomyelitis (EAE) which is a model for Multiple Sclerosis. We use a locally developed EAE simulation which was designed using the CoSMoS (Complex Systems Modelling and Simulation) process. The CoSMoS process was conceived with the aim of promoting the development of rigorous complex system models.

In the model of EAE employed herein, there are two populations of regulatory T-cells (Treg), CD4 Treg and CD8 Treg. The CD4 Treg serve to stimulate dendritic cells to express a protein called Qa-1. Qa-1 permits CD8 Treg to bind to dendritic cells and be activated by them, thus facilitating regulation of autoimmunity.

Previous experimentation demonstrated a large increase in the population of CD8 Treg upon abrogation of the CD4 Treg from the simulation providing that dendritic cells were made capable of constitutively expressing Qa-1.We use simulation to explore two hypotheses proposed to account for this observation. The hypotheses explored are:

i) the timing of Qa-1 expression is influential in determining the population of CD8 Treg.
ii) removal of spatial competition between Treg sub-types favours expansion of the CD8 Treg population.

We demonstrate that both hypotheses are significant in explaining the observed experimental result.

We subsequently investigate addition of a further regulatory mechanism to the existing model. This additional mode of regulation is poorly understood, but has been suggested by an expert immunologist to be an important mechanism of disease regulation. We augmented our model to include this pathway so we could make it more closely resemble the real world murine immune system in EAE. The model implemented proved to cause an overly severe reduction in the activation of T-cells, demonstrating the potential influence of this pathway in disease regulation and that the behaviour of this pathway warrants further investigation.

Table of Contents

List of Figures

List of Tables

Acknowledgements

Throughout the course of this MSc, the author has received a great deal of support from various people and would like to take this opportunity to acknowledge my gratitude to them for this.

Thanks are due to:

Professor Jon Timmis for acting as my Supervisor and for patiently guiding me through the process of the MSc.

Dr Fiona Polack for acting as my Assessor and for discussion on the nature and conduct of an MSc by Research.

Dr Mark Read for the extensive help he gave me in introducing me to EAE and to the EAE Simulator, for advice and discussion of experimental design and not least for reading through the thesis and suggesting ways to strengthen its presentation.

Dr Paul Andrews for his help in setting up my laptop and rescuing it when Windows overwrote the boot-sector of the hard disk!

Richard Williams for providing me with a big head start on my literature review and for conducting the research on which my work builds.

Daniel Moyo, Kieran Alden, Tiong Hoo Lim, Jenny Owen and Nikolai Bode with who I shared RCH/332 during my studies and who played a big part in providing practical assistance and moral support.

Dr Carl Ritson and Dr Fred Bishop who kindly assisted me, firstly in allowing me to run simulations on the CoSMoS cluster at the University of Kent and secondly in providing timely advice and assistance in using the cluster.

To Alex and Tom, the Department's support desk team, for help with numerous little hardware and software niggles such as a dead mouse and trying to install DropBox on my PC.

And last, but not least, to all at YCCSA for support, encouragement, discussions and laughter (listed in no particular order).

A Note on Basic Immunological Notation

When a cell type is preceded by a molecule name e.g. CD4 Treg or CD4CD25FoxP3 Treg, this is taken to mean that the named molecules are expressed by the named cell type and appear on the surfaces of these cells. The expressed protein's identity may also be optionally followed by a '+' (e.g. CD4+ Treg) or '$^{+}$' (e.g. CD8^{+} T-cell) reinforcing the fact that these cells express this protein. Cell sub-populations are often distinguished from each other by the combinations of cell surface proteins that they express.

The notation CD (the so-called 'Cluster of Differentiation') is a protocol used for identifying the surface molecules of white blood cells i.e. of specifying their immuno-phenotype.

Conversely the notation where a protein name is followed by $^{-/-}$ is taken to indicate that a cell type or organism does not possess genes that would allow it to express or synthesise the named molecule. e.g. CD200$^{-/-}$ mice are mice that are genetically incapable of producing the CD200 signalling protein. The use of two minus signs is a convention meant to denote the mouse is homozygous in this deficiency i.e. the mouse genome which should contain two working copies of each gene, one from each parent, is in fact missing both copies.

This thesis uses the simplest convention of listing expressed proteins without additional '+' signs when distinguishing cell sub-populations.

The chemical messengers that facilitate communication between cells ('cytokines') are often numbered and prefixed by 'IL' e.g. IL-2. In this instance IL denotes the fact that the cytokine in question is a molecule of the interleukin family. These are all immuno-modulators that signal changes of behaviour in certain cell types.

Author Declaration

The work contained within this thesis is solely the work of the author, unless otherwise stated.

To my family and friends
who have always treated me with
patient forbearance

Chapter 1: Introduction

We live in an age in which there has been an unprecedented explosion in the volume of biological data available to researchers (from large biological studies such as the Human Genome Project[1], for example). However, this has not automatically translated into rapid progress in the fight against disease as it is difficult to extract knowledge and understanding from the masses of raw data [Germain *et al.* 2011, Kitano 2004, 2002b] – a situation that is complicated by the fact that the data available invariably concerns low-level functions such as gene expression, protein sequences, etc. when what we require is a system (organ / organism) level understanding of disease.

Systems biology has developed from the need to organise and integrate available biological data into some form whereby true understanding can be derived from it [Germain *et al.* 2011, Kitano 2001]. There are two principal fields of endeavour which tackle this problem: bioinformatics concerns itself with extracting patterns or trends from data ('data mining') whereas simulation attempts to create a working computational model of the system (or some abstraction of it) based on representations of its cellular and molecular components. The interested reader can find more detail about systems biology in the work of Kitano [Kitano 2001, 2002a, 2002b, 2004] and in the following review [Germain *et al.* 2011].

Simulation offers us the potential of being able to observe system-wide effects emerge from the interactions of the system's components, effectively linking behaviour at the molecular level to that at the system level [Germain *et al.* 2011]. There are numerous methods available for implementing simulations and models. Of these, we are principally interested in agent-based modelling as this permits the explicit representation of individual heterogeneous cells as individual software agents [Walker and Southgate 2009] and provides a natural description of the immune system.

With greater computational power becoming more readily available, agent-based modelling and simulation has grown in usage in the biological domain [Walker and Southgate 2009, An 2008, Forrest and Beauchemin 2007, Bauer *et al.* 2009]. Use of agent-based models has enhanced understanding of immunology and the mechanisms of disease [Bauer *et al.* 2009] via facilitation of hypothesis testing and the integration of data across multiple experiments [Forrest and Beauchemin 2007]. The focus of this thesis is on agent-based simulations of the murine immune system in the disease Experimental Autoimmune Encephalomyelitis (EAE).

Despite the growing use of computer-based simulation in biology, the uptake of usage has been slow. This may be due, in part, to a perceived 'lack of trust' in the methodology [Polack *et al.* 2010], potentially arising due to the inherent limitations of model building. Firstly, the number of cells in the real system is far greater than even the largest simulation is able to realistically incorporate – potentially leading to

[1] http://www.wellcome.ac.uk/genome

scale-related effects [Kleinstein and Seiden 2000]. Secondly, the model is, at best, an abstraction of the real system – that is certain details are disregarded for the sake of making the computational model tractable and understandable. This entails approximations and assumptions being made [Andrews *et al.* 2010]. These considerations give rise to the need for simulations to be rigorously validated [Andrews *et al.* 2008]. The CoSMoS process [Andrews *et al.* 2010] outlines a rigorous approach to developing models and simulations of complex systems such as the immune system. By providing a principled approach to model building, CoSMoS permits the documentation and communication of models (and the assumptions made during their design and implementation) and hence, the promotion of trust in simulation as a research tool.

The EAE Simulator has been rigorously designed following the CoSMoS approach and has been implemented and calibrated in a principled manner in conjunction with an immunological domain expert [Read *et al.* 2009a, 2009b, 2011, Read 2011]. The work presented here employs augmentations of the existing simulation, the design of which also adhere to the CoSMoS process, to test hypotheses and to model a poorly understood regulatory mechanism.

Throughout the body of the thesis, there is a need to compare experimental results and to be able to assess the significance of any changes brought about by specific changes in parameterization. We employ non-parametric statistical tests to compare the simulation results from pairs of directly comparable experiments i.e. those that differ only in the value of a particular parameter of interest.

1.1 Thesis Goal

EAE is a mouse model of Multiple Sclerosis. This neuro-degenerative condition is caused when immune system cells recognise and attack components of the myelin sheath which protects nerve fibres. The cells which recognise the myelin components are a sub-population of the cells called T-helper cells. The disease is regulated by a further sub-population of T-cells called regulatory T-cells (Treg). These express the protein CD8 on their surface (CD8 Treg). Treg which express the molecule CD4 (CD4 Treg) also exist and these serve to stimulate cells (dendritic cells) which activate CD8 Treg to carry out their regulatory function.

The driving aim of the thesis is to increase understanding of the disease regulation mechanisms in EAE via *in silico* experimentation. This translates into a two-fold goal: at the immunological level, we wish to identify relevant questions that we may direct to the domain expert to address. At the computational level, we seek to gain further insight into the operation of the simulation under hypothetical situations within the EAE system. These are experiments we wish to understand more fully and that, at least at present, are not easily implemented *in vivo* or *in vitro*.

A full exploration of the simulation is beyond the scope of this thesis and since the relevant detail exists elsewhere, interested readers are referred to the Doctoral Thesis of Dr Mark Read [Read 2011].

The thesis is structured as follows:

Chapter 2: Presents a literature review of the relevant domain knowledge i.e. immunology, autoimmune disease as exemplified by Multiple Sclerosis and a model for it, EAE. Relevant computational aspects of the work are also reviewed, for example, systems biology, documenting software models and previous simulations of the immune system. We conclude with a description of the CoSMoS process [Andrews *et al.* 2010] and a synopsis of the existing work carried out using the EAE Simulator [Read *et al.* 2009a, 2009b, 2011; Read 2011, Williams 2010b, Williams *et al.* 2011].

Chapter 3: Presents the experimentation conducted to explain observations from previous simulation within which the CD4 Treg population was abrogated. These cells are required to stimulate the dendritic cells which in turn activate CD8 Treg. We conduct a series of experiments to assess the relative impact of two hypotheses on the peak population of CD8 Treg in this simulation:

i) that the timing of Qa-1 expression by dendritic cells is influential in determining the population size for CD8 Treg
ii) removal of spatial competition with CD4 Treg favours population expansion by CD8 Treg.

Chapter 4: Presents an initial investigation of a poorly understood disease regulatory pathway (the CD200-CD200R axis). Our aim is to incorporate the axis into our model in order to make it a closer approximation to the real-world system and to gain understanding of how this axis might function *in vivo*. The inclusion of the additional regulatory pathway entails an investigation of how it interacts with the mechanism of regulation already implemented in the simulation (i.e. CD8 Treg mediated regulation) and ultimately allows us to rebalance the influence of the two pathways to produce a more realistic model of the EAE system.

Chapter 5: Presents a brief summary of the major conclusions that we have drawn from our experimental evidence and of the contribution made by this work. The chapter subsequently presents work still outstanding from the proposed experimentation and also proposals for future experimentation that build on and further validate the work presented.

Chapter 2: General Overview of the Computational and Immunological Context of the Project

2.1 Introduction

In the previous chapter we stated that the research presented within this thesis represents a body of work in the field of Computational Immunology [Forrest and Beauchemin 2007]. The purpose of this chapter is to provide a foundation in the concepts required to understand the work presented in the later chapters. i.e. the domain knowledge relevant to the simulations carried out, the reasons for doing the research the way that we have and to facilitate sufficient understanding to appreciate what the results of the simulation are telling us.

We adopt the CoSMoS process [Andrews *et al.* 2010] to implement our experimentation. This chapter also follows a CoSMoS-like approach to reviewing the relevant literature. We therefore begin with a description of our domain, which necessitates a brief grounding in immunology (Section 2.2) and an appreciation of the impact of autoimmunity as exemplified by Multiple Sclerosis described in Section 2.3. We then progress to discussion of the specific domain of interest – a model of EAE in mice (described in Section 2.4) [Smith and Kumar 2008a, 2008b, Tang *et al.* 2005] and utilised in the laboratory of our domain expert, Professor Vipin Kumar at the Torrey Pines Institute of Molecular Studies, San Diego.

Having outlined the state of domain understanding upon which we have based our models, we then proceed to discuss the field of Systems Biology (Section 2.5), grounding our work within this context. We outline the potential motivations for conducting simulation-based research in Section 2.6 and take a brief overview of some of the research already carried out in this discipline (Section 2.7). Section 2.8 forms an interlude, and provides brief details of notations that are commonly employed for documenting software models employed in agent-based simulations. Section 2.9 then describes the CoSMoS process which was employed in the design and implementation of the EAE simulator which is described in Section 2.10.

The thesis builds on extensive local work [Read *et al.* 2009a, 2009b, 2011, Read 2011, Williams 2010b, Williams *et al.* 2011] developing and utilising the EAE Simulator. In Chapter 3 we employ the simulator in testing hypotheses proposed by the domain expert to explain the results of previous *in silico* experimentation and in Chapter 4 we seek to add further complexity into our EAE model so as to make it a more realistic representation of the real-world system. The further complexity consists of an additional regulatory pathway which the domain expert believes to be important, but which is poorly understood at present. We seek to employ our *in silico* experimentation to gain further insight into our EAE model and the *in vivo* disease.

2.2 The Immune System

Our domain of interest is the murine immune system in the EAE disease state. It is therefore necessary for us to have some basic knowledge of what the immune system is, what it does and how it carries out its function. This understanding aids us to construct a reasonable domain model from which we may implement simulations.

The immune system comprises two interacting sets of mechanisms (innate and adaptive immunity) that protect the body from pathogens: bacteria, fungi, viruses and protozoan parasites [Garrett and Grisham 2004, Janeway *et al.* 2008].

The protection offered by the innate immune system is rapid and responds to general features of pathogens [Janeway *et al.* 2008]. This is an evolutionarily ancient system [Berg *et al.* 2007] and is invariant across all individuals of a given species. It comprises a set of molecular and cellular responses that are deployed in the early stages of infection. An innate immune response mediated by macrophages is triggered by a breach of the body's physical defences (e.g. the skin) [Janeway *et al.* 2008].

Macrophages are phagocytic or ‘scavenger’ cells and are responsible for recognizing and removing dead cells and cellular debris from the body [Playfair and Lydyard 1995]. As part of this function they engulf and digest invading pathogens [Janeway *et al.* 2008]. Macrophages induce inflammation [Janeway *et al.* 2008], which they achieve by secreting certain signalling molecules called cytokines [Playfair and Lydyard 1995]. Inflammation activates other immune system cells and recruits them to the immune response [Janeway *et al.* 2008]. Sometimes when inoculating against a disease, an injection of vaccine is accompanied by a substance ('adjuvant') e.g. a toxin, which provokes a more aggressive immune response from the phagocytic cells.

A second class of cells that possess phagocytic capability are the dendritic cells. Dendritic cells (DC) form in the bone marrow and then migrate to the periphery where they monitor for pathogens [Janeway *et al.* 2008]. As infection is set up within the body, DCs recognise pathogens via a set of molecules that commonly appear in their cell walls: Pathogen Associated Molecular Patterns (PAMPs). PAMP binding stimulates the DC to phagocytose ('phagocytosis' literally 'cell eating process') the invading cell and present antigens (proteins or fragments of digested proteins that can trigger an immune response) on the cell surface [Janeway *et al.* 2008]. DCs or macrophages presenting antigens in this manner are called Antigen Presenting Cells (APCs) [Kindt *et al.* 2007, Janeway *et al.* 2008].

The DC then migrates from the site of infection to the lymph nodes (which are represented as Secondary Lymphoid Organs (SLO) in the simulation) where it matures and presents the antigens to 'naïve' or 'unprimed' T-cells (T-cells which have not yet encountered the antigen for which they are specific – the 'cognate antigen'). DCs are capable of activating naïve T-cells via presentation of antigens, delivery of an additional excitation signal ('costimulation') and secretion of cytokines [Janeway *et al.* 2008, Kindt *et al.* 2007].

The presentation of antigens to naïve T-cells triggers the adaptive immune response [Janeway *et al.* 2008]. The inflammation initiated by macrophages ensures a more potent immune response by recruiting other immune cell types to the site of infection [Janeway *et al.* 2008]. A full description of innate immunity lies outside the scope of this thesis as the innate response is not explicitly included in our model of EAE.

In contrast to innate immunity, adaptive immunity is highly antigen specific and each individual possesses a unique adaptive immune system dependent on heredity and the history of infections that the individual has been exposed to. This mechanism relies on the ability of the lymphocytes to generate a diversity of antigen receptors that allows them to recognise new potential threats [Berg *et al.* 2007]. The adaptive immune response is more efficient than the innate immune response, but cannot act as swiftly; it requires several days to build up a population of effector cells large enough to mount an effective attack on the invading pathogen [Janeway *et al.* 2008].

Broadly speaking the adaptive immune system also consists of two inter-connected systems: humoral and cellular (‘cell mediated’) immunity, both being effected by lymphocytes. Lymphocytes are a subpopulation of white blood cells [Playfair and Lydyard 1995, Kindt *et al.* 2007], which arise in the bone marrow and which can leave the circulatory system and take up positions in the intercellular spaces. They mediate the recognition of pathogens via specific cell surface proteins [Voet and Voet 2004]. There are two types of lymphocyte: B-cells and T-cells [Janeway *et al.* 2008].

Humoral immunity is mediated by B-cells which originate and mature in the bone marrow. They express B-cell Receptors (BCR) [Kindt *et al.* 2007] which each recognise and bind a specific antigen [Garrett and Grisham 2004]. When a B-cell recognises its cognate antigen via the BCR, the cell is activated and can then differentiate into effector cells ('plasma cells') or memory cells. The plasma cell can secrete a soluble form of the antigen receptor called an antibody [Kindt *et al.* 2007, Playfair and Lydyard 1995] and this molecule then binds to the antigen and marks it for destruction by the body's phagocytic cells. The diversity of antibodies observed is generated by gene re-arrangements of the BCR protein genes [Garrett and Grisham 2004]. Many features of B-cell biology are analogous to those of the T-cells so this description has been provided for the sake of completeness only. B-cells are not included in the domain model of EAE at all.

Both cellular and humoral immunity are triggered by the presence of lymphocyte cognate antigens presented by APCs [Kindt *et al.* 2007]. Cellular immunity is mediated by T-cells which recognise antigens via a diversity of T-cell Receptors (TCRs) [Voet and Voet 2004]. T-cells originate in the bone marrow but mature in the thymus [Voet and Voet 2004]. They have distinctive membrane-bound antigen receptors (TCRs) [Kindt *et al.* 2007, Janeway *et al.* 2008] which recognise a processed fragment of the antigen that is presented as a complex with a Major Histocompatibility Complex (MHC) molecule (see Appendix A.1) on the surface of an APC [Kindt *et al.* 2007, Janeway *et al.* 2008]. The diversity (‘repertoire’) of TCRs is also generated by gene rearrangements ensuring that a great diversity of antigens can be recognised by T-cells [Garrett and Grisham 2004].

The T-cell population is divided into subpopulations which express either CD4 or CD8 co-receptors. Those expressing CD8 only recognise antigen presented on MHC class I compounds, whereas those expressing CD4 recognise antigen presented on MHC class II compounds [Kindt *et al.* 2007, Janeway *et al.* 2008].

If a naïve T-cell does not recognise its antigen within a certain time period it dies of neglect. However, should it encounter its cognate antigen presented by APCs in the SLO it is induced to enter the rest of its cell cycle i.e. it undergoes a proliferative burst and then matures into one of the effector cell types described below. Only T-cells that are responsive to the invading pathogen are produced in great quantity ('clonal selection') [Alberts *et al.* 2008, Voet and Voet 2004].

Typically there are 10^9 lymphocytes in circulation and so the diversity of their receptors ('repertoire') is immense [Janeway *et al.* 2008]. Given the diversity of antigen-specificities available among TCRs, it is likely that at least some of these are reactive towards self-antigens. However, the number of these potentially damaging cells is kept to a minimum by a process of 'selection' during development in the thymus [Kindt *et al.* 2007, Janeway *et al.* 2008].

Once a naïve T-cell has recognised its antigen and become activated, it needs to receive a costimulating signal from a DC to allow it to start proliferating [Janeway *et al.* 2008]. Thus, APCs bound to T-cells regulate T-cell differentiation and proliferation [Voet and Voet 2004]. The T-cells then proliferate and differentiate into effector and memory cells [Kindt *et al.* 2007, Janeway *et al.* 2008]. The process is further stimulated by appropriate cytokine signals from the environment [Janeway *et al.* 2008].

The effector cells of the T-cell lineage are diverse. The CD4 bearing effectors are the T-helper (Th) cells, the CD8 effectors are Cytotoxic T-Lymphocytes (CTLs). The Th cells outnumber the CTL cells by a ratio of approximately 2:1 [Janeway *et al.* 2008].

There are two main subtypes of Th cell, Th1 and Th2. The subtype ('polarization') that a T-helper cell adopts is dependent upon the mix of cytokines in the environment. There are other subtypes of T-helper cell, but these are less well studied and therefore lie outside the scope of this description. The Th1 and Th2 subtypes of CD4 T-helper cell secrete different cytokines. Th1 polarised cells secrete cytokines which perform cell-mediated inflammatory roles, whereas Th2 cells secrete cytokines which have roles in B-cell activation [Kindt *et al.* 2007].

T-helper cells 'help' in a number of distinct ways; they stimulate the differentiation and proliferation of B-cells and 'killer' T-cells [Berg *et al.* 2007, Janeway *et al.* 2008], they provide additional stimuli to activate macrophages [Janeway *et al.* 2008, Kindt *et al.* 2007] and they amplify the immune response via cytokine production [Voet and Voet 2004]. A further role of T-helper cells is in providing assistance to ('licensing') APCs to prime CD8 T-cells. Usually this licensing takes the form of stimulating the APC to express certain stimulatory cell surface proteins [Janeway *et al.* 2008].

During adaptive immune response, CTLs serve to apoptose (kill) infected cells and CD4 T-helper cells stimulate B-cells in the Secondary Lymphoid Organs to produce antibodies against the infectious agent.

CD8 T-cells (Tc) mature into CTLs whose role is to eliminate virus and tumour-infected cells. [Kindt *et al.* 2007, Janeway *et al.* 2008]. CD8 T-cells bind to antigen-MHC Class I complexes presented on APCs [Voet and Voet 2004, Kindt *et al.* 2007, Janeway *et al.* 2008] and in so doing, stimulate the APC's antibacterial mechanisms [Janeway *et al.* 2008]. CD8 T-cells require more activation than CD4 T-cells and must be activated by a mature DC as well as receiving help from CD4 effectors bound to the same APC. This strict requirement for activation acts as a form of safeguard against autoimmunity as the antigen has, in effect, to be recognised twice (once by CD4 T-cells and once by CD8 T-cells) before it can provoke the highly CTLs to act [Kindt *et al.* 2007]. Activated CD8 T-cells effectors kill infected cells by inducing programmed cell death (‘apoptosis’) [Berg *et al.* 2007, Alberts *et al.* 2008, Voet and Voet 2004]. The only CD8 Tc included in our model is the CD8 Treg.

The effector cells described migrate from their site of maturation into the periphery, where they encounter their cognate antigen and are activated without the requirement for costimulatory signals [Janeway *et al.* 2008]. These cells then return to the circulation via the lymphatic system which comprises the body’s main sites of immune response (the lymph nodes and the spleen) [Voet and Voet 2004].

Upon resolution of the infection, the effector B- and T-cells die due to ‘death-by-neglect’, that is they no longer receive the stimulus of meeting and binding their cognate antigens. At this stage in recovery, some B- and T-cells develop into ‘memory’ cells which allow rapid response to any re-occurrence of the same infection. The immune response is finally terminated by the clearing up of cellular debris by phagocytic cells [Janeway *et al.* 2008].

2.2.1 Tolerance

One of the chief characteristics of the immune system is its ability to distinguish antigens derived from the body (‘self’) from pathogen-derived antigens ('non-self') [Voet and Voet 2004]. If the immune cells did not possess this facility then they would be able to freely mount an immune response to any antigen at all that was presented to them, which in theory would mean that the body would destroy its own tissues [Kindt *et al.* 2007]. To protect against this possibility, the immune system has two mechanisms that serve to limit the numbers of self-reactive immune cells (in particular T-cells) in circulation in the periphery – central tolerance and peripheral tolerance. Central Tolerance reduces potentially self-reactive T-cell population numbers and takes place during lymphocyte development in the thymus [Playfair and Lydyard 1995, Kindt *et al.* 2007].

However, due to the vast diversity of T-cell receptors generated, there is a danger that potentially auto-reactive T-cells mature and leave the thymus. Therefore other safeguards against un-regulated T-cell activation exist in the periphery and these

constitute ‘peripheral tolerance’. Peripheral tolerance can be exerted through several mechanisms including the induction of a state of unresponsiveness (also called ‘anergy' [Schwartz 2003]) in self-reactive T-cells outside of the primary lymphoid tissues [Playfair and Lydyard 1995]. This results in the clonal deletion of those T-cells that have recognised their cognate antigen, but have entered the anergic state due to not receiving an additional excitation signal from an APC (‘costimulation’).

Additionally, once a T-cell has been activated, it is only permitted to survive for a certain amount of time before it advances along the cell cycle and enters an apoptotic state. This is known as Activation Induced Cell Death (AICD) [Smith and Kumar 2008a].

Redmond and Sherman 2005 proposed two further mechanisms which are considered important in the tolerization of CD8 T-cells in the periphery. Firstly, the absence of inflammatory pathogens means that antigens are presented by a DC expressing low levels of costimulatory molecules. Secondly, antigen persistence appears to drive tolerance through clonal deletion, anergy or suppression. This second observation makes intuitive sense as presumably self-derived antigens might be expected to persist, whereas antigens from an infection would not.

Unlike the CD4 T-helper cells and CD8 T-Lymphocytes (CTLs) which serve to amplify the immune response, regulatory T-cells (Treg) serve to suppress immune response [Kindt *et al.* 2007]. This suppression is a further mechanism of tolerance and acts as a protection against autoimmunity [Kindt *et al.* 2007].

There are different types of Treg [Shevach 2006] and different populations appear to have different roles. For example, CD4 Treg are not reported as directly killing other T-cells, though CD4CD25Foxp3[2] T-cells are known to suppress population expansion of other T cells [Kindt *et al.* 2007, Kohm *et al.* 2002]. Treg limit the activity of quickly proliferating T-cell populations, often via direct killing mechanisms [Kindt *et al.* 2007, Janeway *et al.* 2008]. There remains much work to be done in characterizing these cell populations [Smith and Kumar 2008b]. Treg mediate their killing ability either in a contact dependent fashion employing direct cell-to-cell killing methods or via the release of cytokines [Janeway *et al.* 2008].

CD8 Treg have been demonstrated to play a pivotal role in the regulation of EAE [Kumar and Sercarz 2001, Smith and Kumar 2008a].

2.2.2 When things go wrong

As has been stated in the previous section (2.2.1), the effective regulation of T-cell activation is essential. Sometimes, for reasons not yet fully understood, the tolerance mechanism breaks down and there is a failure to distinguish 'self' from 'non-self'. When this arises, self-reactive T-cells are able to persist in the peripheral immune

[2] Foxp3 is a protein called a transcription factor that is associated with cells that possess the ability to directly apoptose other cells.

system and are then able to differentiate and proliferate if they encounter their cognate antigen [Kindt *et al.* 2007]. These kinds of inappropriate immune responses are termed 'autoimmunity'. A number of autoimmune disorders have been identified in humans. These are categorised depending on the self-antigen that gives rise to the improper immune response.

Certain autoimmune conditions are organ or tissue specific, for example insulin-dependent diabetes mellitus (IDDM). Other conditions are more generalised throughout the body ('systemic') such as Multiple Sclerosis [Kindt *et al.* 2007, Janeway *et al.* 2008]. In Section 2.3 we briefly focus on the example of Multiple Sclerosis to provide a motivation for the study of EAE described in Section 2.4.

2.3 Multiple Sclerosis

Multiple Sclerosis (MS) is a common neurological disease [Lodish and Darnell 1995]. Charcot first noted the condition in 1869 (English translation [Charcot 1877]) and postulated that it was a novel disease with a well defined progression of symptoms.

MS is a disease of young adult life [Steiner and Wirguin 2000], characterised by destruction of the myelin sheath that insulates nerve fibres in the Central Nervous System (CNS) [McDonald 1974]. As a result, the ability of nerve fibres to conduct impulses is impaired [Lodish and Darnell 1995, Robertson 1981, McDonald 1974].

Clinical symptoms include weakness, incoordination and disturbances of speech and vision. The course of the disease is usually prolonged, the term ‘multiple’ referring to the relapses and remissions that occur over many years [Dorland's 2010].

There are four recognised forms of the disease which are characterised by different patterns of relapse and remission [Dorland's 2010]. The stages are clinically defined on the basis of myelin protein loss and on the geography and extension of demyelinated plaques [Hafler 2004]. Typically, the disease progresses from a relapsing-remitting pattern to a chronic progressive pattern [Chataway 1989].

The cause of the disease is unknown, though the predominant view is that MS is ‘a complex genetic disease associated with inflammation in the CNS and thought to be mediated by auto-reactive T-cells.’ [Hafler 2004].

There are numerous factors that have confused and complicated the study of MS [Steiner and Wirguin 2000]. This lack of clarity has led to a search for ways to simplify our understanding of MS by finding animal models of the disease. One such model is discussed in Section 2.4.

2.4 Experimental Autoimmune Encephalomyelitis (EAE)

Whilst researching post-vaccination CNS dysfunction, Rivers *et al.* 1933 found that by injecting rabbit brain homogenate into the brains of experimental monkeys, they could induce a form of neurological disorder that they named Experimental Autoimmune Encephalomyelitis[3].

The disease is characterised by substantial loss of myelin in the CNS and Lublin argued that relapsing EAE could be a valuable model for human demyelinating diseases such as MS [Lublin 1985]. Full discussion of the forms of EAE lies outside the scope of this thesis however, for the interested reader the various forms of EAE are enumerated and described in detail in the review of the disease by Pender [Pender 1995].

The model of EAE which we use as the basis of our domain model is a mouse EAE model based on the current understanding of the disease of our domain expert, Professor Kumar [Kumar and Sercarz 2001, Tang *et al.* 2005, Smith and Kumar 2008a,b].

2.4.1 Progression of the Disease

The disease cannot develop spontaneously and is induced in experimental animals via inoculation of the subject with myelin-related protein fragments such as myelin basic protein (MBP) and complete Freund's adjuvant (CFA) augmented with Pertussis toxin[4] [reviewed in Pender 1995]. EAE can also be induced in a host by passive induction (transfer of lymph node material from an infected animal to a new host) [van den Bark *et al.* 1985].

After inoculation, an immune response is provoked from the host organism. DCs phagocytose the injected MBP. This is then presented on the surface of the DCs which migrate away from the periphery.

EAE is mediated by CD4 T-cells specific for MBP [van den Bark *et al.* 1985, Pender 1995]. The CD4 T-cells recognise the MBP-MHC complexes presented by the DCs in the SLO and are activated and begin to proliferate [Pender 1995]. The MBP-reactive CD4 T-helper cells are then able to migrate into the CNS and localise there [Ludowyk *et al.* 1992]. The CD4 Th express molecules [Baron *et al.* 1993] that then allow them to pass through the Blood Brain Barrier (BBB) which has been made more permeable by the Pertussis toxin in circulation [Mostarica-Stojkovic *et al.* 1992]. Once the auto-aggressive T-cells have passed the BBB, inflammation ensues. This causes the observable signs of EAE in the spinal cord. The increased permeability of the BBB during EAE also leads to oedema formation which is correlated with the clinical signs of EAE [Claudio *et al.* 1990].

[3] Sometimes also referred to as Experimental Allergic Encephalomyelitis.
[4] The neuro-toxin from *Bordetella pertussis* which causes whooping cough.

Other immune system cells are able to enter the CNS. In fact 50% of the CNS invading cells in the EAE cycle are macrophages [Huitinga *et al.* 1990] with most of the remainder being CD4 T-cells [McCombe *et al.* 1994]. Elimination of macrophage populations in the CNS during EAE, eliminates clinical signs of the disease suggesting that macrophages play a role in demyelination [Huitinga *et al.* 1990].

Microglia are the resident macrophages of the CNS. Activated microglia are detrimental to the CNS because they are capable of secreting potent neurotoxins such as Tumour Necrosis Factor-α (TNF-α) [Carson 2002, Pender 1995].

The CD4 Th1 cells that have invaded the CNS are capable of activating both the microglia [Carson 2002] and the macrophages that have also invaded the CNS [Huitinga *et al.* 1990, Pender 1995]. These macrophages then secrete TNF-α which damages neurons [Huitinga *et al.* 1990, Pender 1995, Carson 2002].

Following the onset of EAE there is a significant increase in the population of Interferon-γ (IFN-γ) secreting cells in the CNS and SLO [Mustafa *et al.* 1991]. IFN-γ is a potent activator of macrophages and so induces the production of TNF-α. IFN-γ also up-regulates the expression of MHC-II and adhesion molecules which permit cells to cross the BBB [Pender 1995].

TNF-α secretion by macrophages in the CNS causes extensive nerve damage. The apoptosed neurons are phagocytosed by macrophages which then present the MBP from the dead cells. This is recognised by MBP-reactive CD4 Th entering the CNS and the cycle of disease is perpetuated [Pender 1995].

Ultimately the CD4 T-helper cells in the CNS become apoptotic and leave the CNS [Pender 1995]. Having left the CNS, the MBP-reactive cells can be digested and processed by DC which then present peptides from the digested auto-reactive TCRs (principally the Complementarity Determining Region (CDR1/2) and the Framework region 3 (Fr3) segments from the TCR molecule) [Pender 1995, Kumar and Sercarz 2001, Tang *et al.* 2005].

2.4.2 Spontaneous Recovery

Once DCs have phagocytosed the apoptotic auto-reactive CD4 T-Helper cells, they can process and present fragments of proteins present in the T-cells including proteins from the TCR. In particular, DC present the Fr3 and CDR1/2 portions of the auto-reactive CD4 Th TCR [Pender 1995, Kumar and Sercarz 2001, Tang *et al.* 2005]. The MHC-II-Fr3 complex presented on the DC surface can be recognised and bound by CD4 Treg [Kumar *et al.* 1996] which then stimulate (or 'license') the DC for production of the MHC Class Ib compound Qa-1 [Cantor *et al.* 1978]. The licensed DC can then bind CD8 Treg which recognise the presented Qa-1 – CDR1/2 complex [Tang *et al.* 2005]. CD8 Treg are then activated, with further help from IFN-γ secretion by CD4 Treg, and proliferate [Kumar and Sercarz 2001, Pender 1995].

Effector CD8 Treg recognise Qa-1 – antigen complex [Kumar and Sercarz 2001] presented on the surface of CD4 Th and kill them [Beeston *et al.* 2010, Pender 1995]. Consequently, the CD4 Th population falls during the recovery phase and within ~30 days the population has returned to resting levels [McCombe *et al.* 1994] and within ~50 days the disease-recovery cycle is completed and the animal is fully recovered [Pender 1995].

The number of CNS T-cells declines substantially during remission with the apoptosis of T-cells reaching a peak during the recovery phase [Pender 1995]. After full recovery from acute EAE the number of T-cells falls back to near baseline levels [McCombe *et al.* 1994].

In cell-mediated recovery, CD8 Tregs participate in the apoptosis of CD4 Th [Kumar and Sercarz 2001], leading to a feedback inhibition regulatory mechanism [Tang *et al.* 2005].

Having described what it is that we wish to simulate, we now turn our attention to techniques that we may use to create a simulation, how we might choose to document the model so created.

2.5 Systems Biology

Systems Biology is the integration of experimental and computational research to explore scientific questions in biology [Kitano 2002a]. The domain has developed in response to the pressing need to integrate the vast amounts of molecular biology data and derive understanding from it [Kitano 2002a, Germain *et al.* 2011]. This process has been facilitated by the ready availability of increasingly powerful computational resources [Kitano 2002a].

Systems biology consists of two fields of investigation: bioinformatics and simulation based analysis. Simulation is used to predict the dynamics of a system or to test assumptions made about the system [Kitano 2002a]. The reader interested in how these fields inter-relate is referred to the recent review by Germain [Germain *et al.* 2011].

Systems biology aims to help us to understand biology at system (e.g. whole organism) level via modelling of the structure and dynamics of the system at a cellular or molecular level [Kitano 2002b]. However, the quality of that understanding is strongly determined by the reliability of the data on which the models are built [Kitano 2002b].

Kitano 2002a proposes that model quality is enhanced via model exchange and communication. This facilitates community validation of models used in published studies [Kitano 2002a] and promotes trust in the results obtained from them [Polack *et al.* 2010]. Model documentation is revisited in greater depth in Section 2.8 where we discuss the Unified Modelling Language.

A full review of systems biology lies outside the scope of this thesis. The definitions, spirit and practice of systems biology are reviewed in [Kitano 2001].

2.5.1 Biological Complex Systems

A general view of a complex system is a system of very many simple, homogeneous entities, each of which exhibits simple behaviour and is capable of interacting with its neighbours. Such a system of simple entities then leads to complex system-wide ('emergent') behaviour [Kitano 2002a].

However, biological systems are more difficult to understand because system level events result from complex interactions between large numbers of heterogeneous components [Kitano 2002a]. Such systems are inherently dynamic and multi-scale [Cohen 2007], with behaviour evolving in response to discrete event cues or in specific spatial locations [Bauer *et al.* 2007]. Behaviour at different spatial and temporal scales arises from local mechanisms [Forrest and Beauchemin 2007] and models typically have several 'compartments' to represent different tissue types or bodily locations to reflect this [Germain *et al.* 2011].

Adaptive immunity provides a good illustration of these concepts. Adaptive immunity is a complex process involving spatial and temporal organization of system elements [Segovia-Juarez *et al.* 2004] with many of its processes being highly dependent on specific cellular interactions and their spatial locations [Thorne *et al.* 2007]. The emergent properties of the immune system arise from the heterogeneity of the dispersed cellular interactions in the system and from stochastic events occurring within it [Germain 2001]. In this way, both the timing and the location of an event determine the result of that event.

A further example is provided by the specificity of responses to cell signalling, which are determined by the spatial and temporal dynamics of signalling networks [Kholodenko 2006]. Temporal dynamics were found to be coupled to spatial gradients of signalling activities, whereas localization of signalling proteins lead to the discrete sub-cellular location of excitation events.

2.6 Motivating Immunological Simulation

Immune system modelling and simulation is becoming more widely used as is evidenced by a number of relevant reviews [Walker and Southgate 2009, Forrest and Beauchemin 2007, Bauer *et al.* 2009].

A model is 'an abstraction that is made to aid understanding or description of something' [Andrews *et al.* 2010]. A model often makes assumptions of varying justifiability and these must be explicitly documented to facilitate validation of the model [Andrews *et al.* 2010], often with the assistance of a domain expert [Read *et al.* 2009a].

A simulation is an ‘executable model’ [Polack *et al.* 2010]. Simulation is the ‘technique of imitating the behaviour of some situation or process… by means of a suitably analogous situation or apparatus esp. for the purpose of study or personal training [*sic*]’ [Andrews *et al.* 2010].

One reason for constructing a model is to increase understanding of a system by creating an 'acceptable simplification' of that system [Andrews *et al.* 2010]. Integration of data into a coherent whole for use in modelling [Read *et al.* 2011] often helps us organise the available data in a way that allows previously unasked questions to be raised by it [Kam *et al.* 2001]. Thus, modelling becomes a means to providing explanations and making predictions [Kam *et al.* 2001].

Similarly, An 2008 saw modelling as ‘a formal means of testing, evaluating and comparing knowledge’ that is to say, a form of ‘conceptual model verification’. This view of modelling as a mode of hypothesis testing is also shared by other authors [Read *et al.* 2009a, Read *et al.* 2011]. Modelling may also be utilised as a mode of hypothesis generation once the simulation results have been interpreted in light of expert domain knowledge [Read *et al.* 2011].

Finally, at a practical level, wet laboratory experiments are often costly and / or difficult to perform – and in some cases may be simply not possible (technically or ethically) in the real world. Modelling and simulation therefore permit us to ask questions that could not easily be answered by standard laboratory techniques and serve as a useful complement to wet laboratory experimentation [Read *et al.* 2009a].

The main power of modelling is that it has the potential to produce a clearer picture of system behaviour by abstracting away low-level information. For example, the immune system is composed of $\sim 10^{12}$ cells of numerous types each consisting of various sub-populations [Seiden and Celada 1992] which interact to produce the emergent behaviour. To understand such a system in its entirety would be overwhelmingly difficult both in terms of the computational power required and in the scientist’s ability to describe the system [Read *et al.* 2011].

The methods used to build a conceptual model of a system differ according to the methods being used to study the system. As we are principally concerned with agent-based models we present a brief description on model building specific to agent-based modelling in Section 2.7.3.

Once implemented, models are parameterised using relevant biological data for the agents chosen [Macal and North 2005]. A parameterised model must then be validated against observed system behaviour before the model can be used to make predictions [Macal and North 2005]. The modelling process is usually iterative [Kitano 2002b], with the model being refined through several definition – implementation – parameterization – validation cycles [Macal and North 2005].

A number of different mathematical and computational methods have been employed in the *in silico* study of immunology. These are discussed in the following section.

2.7 Computational Techniques Employed in Modelling the Immune System

The following sections describe the techniques commonly used to model the immune system. The author aims to provide some insight into the techniques and their relative strengths and weaknesses. A short synopsis of immune system models generated via these techniques is also provided to illustrate the wider applicability of the methods.

2.7.1 Ordinary Differential Equations

Traditionally mathematical modelling of physical and biological systems has made extensive use of Ordinary Differential Equations (ODEs). Use of ODEs can prove invaluable in gaining insight into the dynamics of a system, for example the dynamics of HIV infection where AIDS occurs on a timescale of years, but yet which commences with events that occur over hours or days [Perelson and Nelson 1999]. This study successfully modelled the lag between the initial infection with HIV and the onset of full HIV/AIDS principally by accounting for changes in the circulating CD4 Th1 cell population. The interested reader is referred to Perelson's review of modelling viral infection and immune system dynamics [Perelson 2002] for a fuller exploration of the use of ODEs in modelling the immune system.

Despite the successes of using ODEs to model biological systems, such an approach does entail several perceived weaknesses. One obvious weakness becomes apparent upon reading the mathematical model; it is a daunting description for a non-mathematician to comprehend [Seiden and Celada 1992]. However, the principal weakness of the use of ODEs lies in the fact that they average space or population behaviour [Kleinstein and Seiden 2000], whereas we may be rather more interested in a spread of behaviours, for example, the diversity of functions exhibited by cells of a given population..

ODEs are continuous representations of systems that contain many discontinuities (that is, the real system is composed of discrete individuals) [Kleinstein and Seiden 2000, Seiden and Celada 1992, Bernaschi and Castiglione 2001]. The equations generated are therefore difficult to solve exactly and attempts to solve them analytically necessitate the making of potentially unacceptable approximations to the real nature of the modelled system. Often these equations need to be solved using numerical methods [Kleinstein and Seiden 2000, Seiden and Celada 1992]. It is therefore often difficult to modify the complexity of the model. For example if one chose to incorporate a new cell type into the existing model, then an entirely new set of equations would need to be derived in order to integrate the behaviour of this new cell type into the behaviour of the overall system [Kleinstein and Seiden 2000].

An ODE model assumes a population of entities with essentially identical properties that can be calculated. This is particularly problematic when simulating the immune system which consists of many cell types, some of which are heterogeneous in their

properties. One response to this might be to subdivide each class, but this would result in a good many sub-classes some of which may not be present in particularly great numbers [Kleinstein and Seiden 2000].

2.7.2 Cellular Automata

Cellular Automata (CA) are conceptually defined by a lattice of grid points which represent locations within a system of interest. Each grid point has a state associated with it and this state is changeable according to the rules of the automaton. The rules are based on the current state of the grid point of interest and the states of its immediate neighbours. Thus well defined state transitions are permitted to occur for each grid point in the lattice [Walker and Southgate 2009].

During simulation, the grid points can be populated by entities e.g. cells, and these 'move' around the grid as dictated by the changes in the states of the grid points. It is worth noting here that the agents themselves do not possess heterogeneous state.

CA have been used extensively in immune system modelling. Seiden and Celada proposed using CA to implement *in silico* modelling of the immune system as an adjunct to traditional wet lab experimentation [Seiden and Celada 1992]. In their proof of concept paper they proposed a generalised implementation of an immune system model which would be based primarily on a detailed description of the interactions between different cell types. The authors considered that their results 'reliably simulated' the immune system and that it would be appropriate to scale up the model to incorporate a greater number of cells.

Several subsequent cellular automaton models have built on this initial work, mostly focussing on increasing the number of cells that the system could incorporate by using more efficient data structures. These include IMMSIM [Kleinstein and Seiden 2000], PARIMM [Bernaschi and Castiglione 2001] and C-ImmSim [Baldazzi *et al.* 2006].

Santoni *et al.* 2008 employed C-ImmSim in their simulation of T cell differentiation during hypersensitivity reaction. The simulation was found to reproduce the essential features of hypersensitivity reaction.
Walker and Southgate provide a thorough review of approaches to modelling that use CA and continuous space ('off-lattice') agent-based modelling techniques [Walker and Southgate 2009].

The work with CA-based simulations has stimulated modelling studies by clearly demonstrating the potential of such work, whilst remaining aware of potential short-comings such as limits on the size of the system. The simulators are designed and described using language relevant to the study of biology ('domain specific language') meaning that the model is immediately comprehensible to biologists [Seiden and Celada 1992]. This conceptual simplicity offers several potential advantages as compared to ODEs (these are essentially the converse of the perceived disadvantages of ODEs).

One principal advantage of CAs is that they represent agents within a model e.g. cells explicitly, whereas ODEs represent populations of agents [Kleinstein and Seiden 2000]. Implicit in this is the fact that we can exploit stochasticity to estimate distributions of behaviour rather than average system behaviour [Kleinstein and Seiden 2000]. Solutions to ODEs are often analytical but approximate, failing to capture exactly important features of the system under examination [Hone 2009].

As stated in Section 2.6, use of simulation, offers one clear potential advantage over traditional wet-laboratory experiments. However, simulations are only capable of modelling systems of relatively small size compared to the real system and so we are forced to sacrifice detail in the model. The modeller also has to be aware of the potential for size related artefacts in the resulting behaviour of the system [Kleinstein and Seiden 1990, Seiden and Celada 1992].

Additionally, simulation introduces abstractions and assumptions to reduce model complexity. This raises the question of whether the simulation is still representative of the system being modelled. Creation of simulations in which biologists can trust requires considerable effort [Polack *et al.* 2010, Andrews *et al.* 2010]. So, although simulation is much less costly, due to the limitation of the finite system we cannot yet hope to replace wet-laboratory experimentation [Seiden and Celada 1992].

2.7.3 Agent-Based Modelling and Simulation

CA are a form of Agent-Based Modelling [Walker and Southgate 2009]. However, there are important differences between CA and continuous space ('off-lattice') ABM methodologies. In CA, the agents are represented as an aspect of the state of the grid points [Thorne 2007] and are not truly mobile [Wishart *et al.* 2004]. In ABM the agents are explicitly represented and autonomous. They possess their own heterogeneous states [Wishart *et al.* 2004].

ABMS is an approach tailored to the modelling of ‘autonomous, interacting agents’ [Macal and North 2005] and which simulates the 'actions and interactions' of these agents [Macal and North 2005]. Simulation is generally aimed at elucidating the mechanisms of ‘emergence’, the process by which the microscale behaviour of the system i.e. of the agents themselves, translates into the overall (or ‘emergent’) behaviour of the system [Macal and North 2005].

This technique offers clarity of insight into the behaviour of a system, which often cannot be gained from mathematical models such as ODEs [Kam *et al.* 2001]. For this reason, ABMS has found application in many disciplines [Macal and North 2005] where it provides information on the 'dynamic aspects of real-world systems' [Bonabeau 2002].

ABMS is particularly suited to modelling systems in which the key entities can naturally be described as separate agents and where spatial location is an important part of an object’s state e.g. cells in an organism [Macal and North 2005] and it is natural to employ ABMS when the agents exhibit complexity and stochasticity of

behaviour [Bonabeau 2002]. It is also appropriate for the simulation of multi-scale models, the results of which are heavily dependent on correctly identifying the key agents and on any simplifying assumptions made while designing the abstraction of the system [An 2008].

The term 'agent' carries a variety of subtly different connotations but in terms of modelling the immune system, we understand that agents are autonomous entities [Macal and North 2005, Mellouli *et al.* 2004, Bonabeau 2002] which can exhibit a range of heterogeneous states [Macal and North 2005] and encompass stochastic behaviour [Bonabeau 2002]. An agent should also be 'situated' [Macal and North 2005] i.e. have a well defined location within the system and able to interact with other agents [Mellouli *et al.* 2004].

Creation of an ABM of a system begins with identification of the entities which are 'agents' within the model followed by enumeration of all the possible relationships between the different agents [Macal and North 2005]. Only once these details have been decided can the user begin to draw up a strategy for developing the model, which probably includes deciding which of the several ABMS development platforms available one is going to employ. ABMS development platforms available include Netlogo[5] [Tisue and Wilensky 2004], Repast[6] [Collier *et al.* 2003], Multi-Agent Simulation of Neighbourhoods (MASON)[7] [Luke *et al.* 2003, 2004, 2005], Swarm[8] [Minar *et al*. 1996] and Flexible Large-Scale Agent-based Modelling Environment (FLAME)[9] [Holcombe 2006] *inter alia*. Once the model has been implemented the process entails iterative parameterization and validation as described in Section 2.6.

There are a growing number of ABM studies of the immune system. A comprehensive review of agent-based modelling of the immune system can be found in Forrest and Beauchemin 2007 and some influential examples are cited here.

Kam et al. 2001 created models of 'reactive systems' using the Rhapsody[10] State Chart-based modelling environment. They concluded that their multi-scale model of T-cell activation could realistically bridge the gap between reductionist and system-level views of the immune system. Building on this, Efroni *et al.* successfully modelled the effects of two gene 'knockout' mutations on the dynamics of thymocyte development in the thymus [Efroni *et al.* 2005, 2007].

In a direct comparison of simulation with experimental data, Walker *et al.* 2004 compared results from an ABM of cell growth in the epithelium to the experimentally determined growth characteristics of epithelial cells grown in monolayer culture. The authors found that the model data was a qualitatively good fit to the experimental data.

[5] http://ccl.northwestern.edu/netlogo/
[6] http://repast.sourceforge.net/api/index.html
[7] http://cs.gmu.edu/~eclab/projects/mason/
[8] http://www.swarm.org/index.php/Main_Page
[9] http://www.flame.ac.uk
[10] http://www.ilogix.com

During the same year, Segovia-Juarez *et al.* 2004 created an ABM of granuloma formation in the lungs during tuberculosis infection. The model identified several key factors in granuloma formation, which included the rate of chemokine diffusion and the location and number of T-cells within the granuloma.

More recently, Pogson *et al.* 2008 have used the FLAME modelling environment to create a model of the intracellular NF-κB signalling pathway. The model permitted the authors to predict that NF-κB inhibition is affected by the sequestering of excess NF-κB inhibitor by actin fibres in the cell cytoskeleton, a behaviour which was subsequently confirmed experimentally.

The above studies demonstrate the potential of well validated simulation to generate explanations of system behaviour and to make predictions about the system based on current domain knowledge. The extensive body of local work using ABM to investigate the mechanisms of disease in EAE are described in the final section of this chapter (Section 2.10).

Many of the perceived advantages of ABMS over ODEs are the same as those discussed for CA in Section 2.7.2. The technique also offers flexibility compared to other methodologies e.g. it is easier to add new agent classes or to tune agent behaviour [Bonabeau 2002].

Thorne 2007 identified several key challenges to ABMS experimentation in his review of the use of ABMS in biomedical science. In particular, he cautioned of the imperative need to couple such simulation and modelling to wet laboratory experimental work so as to address the issue of model parameterisation. Other challenges included finding the appropriate level of abstraction for the model and communicating completely the model implemented. Thorne also suggested the potential of exploring the simulation parameter space using sensitivity analysis.

A further consideration to bear in mind when conducting ABM is the environment in which the agents function. The environment defines the properties of the 'world' in which the agents operate and provides the conditions under which they can exist. An agent is 'of no practical use without its environment' [Odell *et al.* 2002]. Design of effective agents therefore requires careful consideration of both physical interaction with, and communication through the environment [Odell *et al.* 2002].

2.8 Documenting Models

One of the aims of Systems Biology is the communication and documentation of models [Kitano 2002a]. In the following section we discuss two widely used formalisms for documenting software implementations.

2.8.1 State Charts

State Charts were devised by David Harel as a means for visually describing complex software systems and represent an extension of the formalisms of state machines and state diagrams [Harel 1987]. State Charts introduce the 'notions of hierarchy, concurrency and communication' to the concept of state.

State Charts are considered to be 'compact and expressive' and Harel sees them as overcoming many of the perceived shortcomings of state machine diagrams. For example with State Charts we can cluster states into super-states, show refinement of states or illustrate state orthogonality.

Although the State Chart formalism promises many potential advantages to the developer of biological complex systems, it seems to have been overshadowed, perhaps rather unjustly, by the more widely used UML conventions [Booch *et al.* 2005, Larman 2005]. Indeed, the UML state machine formalism is derived from Harel's State Charts and shares many features in common with them. However, the two formalisms differ in semantics, the exact nature of which lie outside the scope of this thesis.

State Charts have principally been employed by the Harel laboratory at Rehovot University, for example, the studies cited in the discussion of biological studies using ABM in Section 2.7.3 [Kam *et al.* 2001, Efroni *et al.* 2005, 2007].

2.8.2 Unified Modelling Language (UML)

It is contemporary good-practice in the software industry to use UML notation in the description of software systems. UML notation affords the designer a standardised ‘visual language’ [Booch *et al.* 2005] which can be used to document and conceptualise the software [Larman 2005], A number of web-based UML support tools, resources, tutorials and articles are available via the Object Management Group [OMG].

UML offers a variety of diagram types for the description of software systems at different levels of abstraction, each having its own specific and well-defined purpose and each representing a different view of the system [Booch *et al.* 2005]. As well as serving to describe system architecture, UML serves as a record of implementation decisions or as a reverse engineering tool [Larman 2005].

Through its wide-spread use in the software industry, UML has gained currency and is widely used within the scientific and engineering communities. Several studies have been conducted to probe the suitability of UML formalisms for documenting biological complex systems and with a few minor reservations the authors agree that UML provides an adequately expressive language with which to fully describe these systems.

Webb and White 2004 presented a UML cell model as an illustration of the practical uses of the object-oriented programming paradigm. This was followed up by a further model of the cell which was used to elucidate the principles of object-oriented program design in complex system modelling [Webb and White 2005]. The authors concluded that simulation based on the model should be possible given that it fully describes the system in a top-down manner.

In a more recent study, Garnett *et al.* 2008 created an executable model of auxin transport in plants. Working from a high level towards more detailed representations, the authors created UML models to successfully capture the key perspectives of the system.

In a proof of concept, Williams set out a description of the glycolytic pathway and TCA described in UML models. The author clearly demonstrated the utility of creating conceptual models of biological pathways in UML notation [Williams 2010a].

We now turn our attention to the CoSMoS process, a process that assists in the rigorous conceptualization and implementation of complex system models and simulations.

2.9 The CoSMoS Process

The CoSMoS Process is a process for the modelling and simulation of complex systems [Andrews *et al.* 2010]. In this instance 'complex system' carries the connotations given to it in Section 2.5.

CoSMoS focuses on three principal objectives: to provide tools for creating models and techniques for setting up simulations for use in studying complex systems, to provide a guide to best practice in the modelling and analysis of complex systems and to help in the design of and encourage proper validation of complex system models. These goals should then allow us to "elaborate and explore science in a wider context." [Andrews *et al.* 2010].

Simulation supports mainstream scientific research e.g. by permitting experimentation *in silico* which would be costly or otherwise difficult to perform *in vitro* or *in vivo* (see Section 2.6 for the motivation behind computational simulation).

The process is constructed from principles derived from case-studies in immunology, and ecology *inter alia* and consists principally of phases, products, activities and roles.

The CoSMoS process consists of three phases termed Discovery, Development and Exploration. The phases generate five products which are the Research Context, the Domain Model, the Platform Model, the Simulation Platform and the Results Model. These are illustrated in Figure 2.1.

The phases each entail the possibility of several activities which have the effect of modifying the products of the process e.g. scoping, which might, for example, involve considering the relevance of including a new cell sub-population in the model.

The products from the three phases are inter-related and are also strongly identified with the Research Context and the Domain, which is to say that we create a simulation of a specific system to ask specific questions. These relationships are illustrated in Figure 2.2.

The Research Context captures the overall context of the project – including a motivation for the project, the questions addressed and also the validation and evaluation needs.

The Domain Model describes the domain expert's understanding of the appropriate aspects of the domain and thus is focused on current scientific knowledge of the system.

The Platform Model incorporates the design and implementation models for the simulation platform based on the Domain Model and the Research Context. It serves to document the decisions made in the construction of the model and how these might be implemented into programming code.

The Simulation Platform actually encodes the Platform Model into useable software and hardware platforms on which simulations are performed.

The Results Model describes the understanding that results from the simulation. This includes insight into Simulation Platform behaviour, results of data collection and observations of simulation runs.

The interested reader is referred to the relevant technical report [Andrews *et al.* 2010].

Figure 2.1: The Phases and Products of the CoSMoS process. In the diagram the solid arrows indicate the progress of the project through the cycle of phases and the dashed arrows indicate those products on which each phase acts. The process begins with a Discovery phase during which the current understanding of the domain is detailed with the help of a domain expert, resulting in a Domain Model. At this stage the research questions and purpose of the project are mapped out, leading to an appreciation of the Research Context. During Development the software engineer starts to implement a model in programming code (Simulation Platform), working from the domain model created during Discovery. The implemented simulation is then executed to validate it and verify the model. The simulation produces results (the Results Model) which must be interpreted in the light of domain knowledge. The results of Exploration may then feed forward into another phase of Discovery. Adapted from Andrews *et al.* 2010.

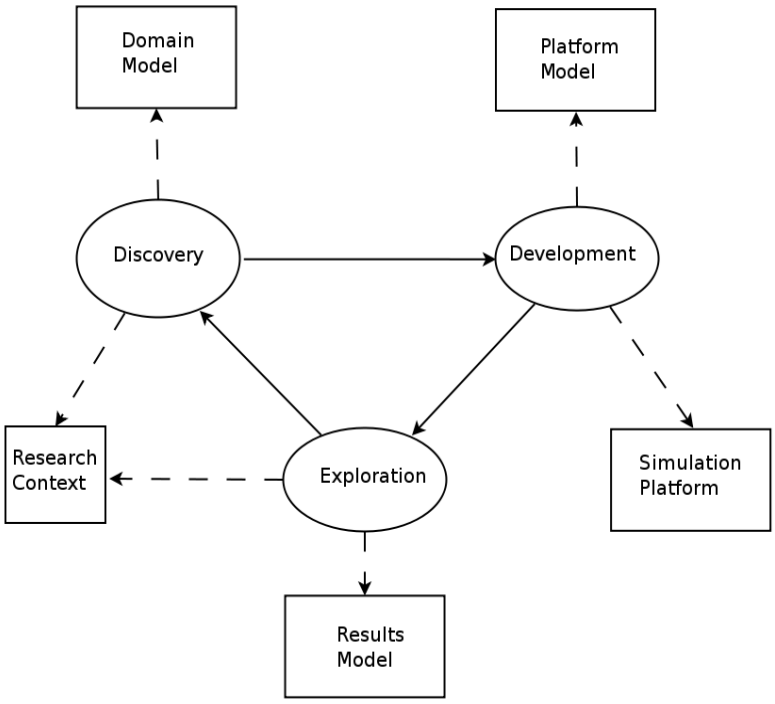

Figure 2.2: The relationships between the Products of the CoSMoS Process and the Domain. The process is iterative, but should normally begin with consideration of the domain. Discussion of the domain with an expert in the field ('domain expert'), allows a first model of the domain (‘Domain Model’) to be developed and parameterised. Only once a full model is in place can the software engineers begin to implement the model as computer code (Platform Model leading to an implemented Simulation Platform). Once the simulation is implemented it needs to be tested against the expectations of the researchers and the domain expert ('validation'). The model should also be rigorously calibrated, which may include a test of sensitivity to parameterization. The results should then be interpreted in light of what is known of the domain (Results Model). All the previous steps should be carried out with strict reference to the Research Context that is the overall context of the project – the reasons for creating the simulation, and the question(s) it was created to address. The process is intended to be flexible, but the typical flow of information involved in developing the products is indicated by the arrows in the diagram. Adapted from Andrews *et al.* 2010.

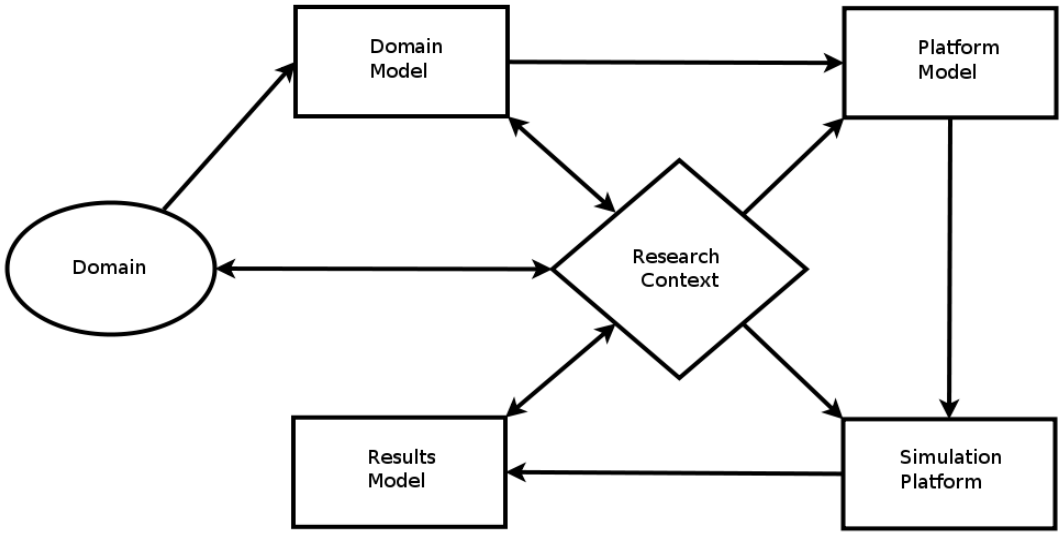

2.10 The EAE Simulator

The EAE Simulator employed in this thesis was written by Mark Read at the University of York [Read *et al.* 2009a, 2009b, 2011, Read 2011]. The simulator was designed as a CoSMoS project and as such has generated the relevant model products from the three phases of the process – Discovery, Development and Exploration. The Simulator is fully described using UML diagrams [Read *et al.* 2009a, Read 2011] and implemented using the object-oriented paradigm which is particularly suited to ABMs. Simulator design was carried out in collaboration with a domain expert, Professor Vipin Kumar of the Torrey Pines Institute of Molecular Studies (TPIMS).

Elements of the immune system in EAE correspond to objects in the ABM. These agents exhibit stochastic behaviour within the context of the model. The behaviour of the agents is determined by various model parameters derived from experimental data, with guidance from the domain expert. Because of the inherent stochasticity of the system, the outcome of simulations using the simulator could be highly dependent upon the parameterization employed. A recent analysis of simulator sensitivity to parameterization suggested that at least 500 simulation runs would be required to reduce chance uncertainties in responses to parameter changes to an acceptably small value [Read *et al.* 2011].

The Simulator has also been extensively validated [Read *et al.* 2011]; testing sensitivity to parameters and robustness using a Latin Hypercube algorithm [Mckay *et al.* 2000] to systematically measure the model response to all possible combinations of changes across the model parameters. Such an approach can yield seemingly counterintuitive results (see for example [Dancik *et al.* 2010]) and can reveal critical components and pathways or redundancies within the modelled system.

2.10.1 Domain Model

The Domain Model for the simulator provides a full diagrammatic description of the behaviour of the agent types in the simulator under different conditions or ‘states’ and at different levels of abstraction. The model consists of some 25 UML diagrams and is more fully represented in [Read 2011].

In short, the simulator incorporates a domain expert approved abstraction of the full disease-recovery cycle of EAE, incorporating all the relevant immune cell types [Read *et al.* 2009a]. A conceptual representation of the interacting disease and recovery cycles is presented in Figure 2.3 below.

2.10.2 Implementation of the Simulation

The simulator is implemented in Java using the agent-based development framework, MASON [Luke *et al.* 2003, 2004, 2005]. Many of the accompanying analysis scripts are written either in Ruby or the Matlab[11] scripting language.

[11] http://www.mathworks.com/matlab

2.10.3 Implementation Specific Details

For the interested reader, certain implementation specific description of the simulator has been provided in Appendix B of this thesis along with UML diagrams representing certain aspects of the structure of the software.

2.10.4 Results Model

The two cycles, disease and spontaneous recovery are clearly visible in the system-wide plot of T-cell populations over a 50 day simulation cycle (Figure 2.4). Firstly the CD4 Th1 cells reach a peak population at around 15 days and then start to wane. At about this time the populations of the Tregs start to climb, with CD4 Treg reaching a higher peak population than CD8 Treg at around day 27. This time marks the height of the recovery period and by day 50 the T-cell populations have returned to baseline levels and recovery is complete.

2.10.5 Summary of Results Obtained Prior to the Start of the Present Work

Significant work has been done in examining the effects of splenectomy on the system and in generating an EAE Severity Score from the calculated T-cell effector populations and levels of neuron killing [Read 2011].

Figure 2.3: A conceptual representation of the EAE Disease and Spontaneous Recovery Cycles. The disease cycle is represented by the red arrows and the recovery cycle by the green. The DCs (shown as generic APCs) play a pivotal role in both cycles, both initiating the immune response to myelin basic protein and also serving to activate the CD8 Treg which are responsible for cell mediated immune regulation. The cycle is initiated with inoculation with MBP and CFA. These stimulate the DCs which phagocytose MBP and display antigen fragments. These are recognised by MBP-reactive CD4 Th1 which are activated and begin proliferating. Due to changes in the permeability of the BBB these CD4 Th1 can enter the CNS where they stimulate the microglia to produce neurotoxic TNF-α which kills neurons. The dead neurons are phagocytosed by CNS-invading macrophages and their MBP is presented by these APCs. This activates further CD4 Th1 inside the CNS leading to a progression in the disease state. Eventually, however, the CD4 Th1 reach the end of their lifespan and become apoptotic. Then they can leave the CNS where they are phagocytosed by DCs (APCs) which then present fragments of the MBP-reactive CD4 Th1 TCR. The Fr3-MHC-II complex is recognised by CD4 Treg which stimulate DCs to express Qa-1. This is recognised by CD8 Treg which can then bind to the DCs and become activated. The activated CD8 Treg can then recognise Qa-1-CDR1/2 presented on the surface of MBP-reactive CD4 Th1 and bind the CD4 Th1 prior to apoptosing them. The apoptosed CD4 Th are then phagocytosed by DCs to complete the recovery cycle. Adapted from Read *et al.* 2009a.

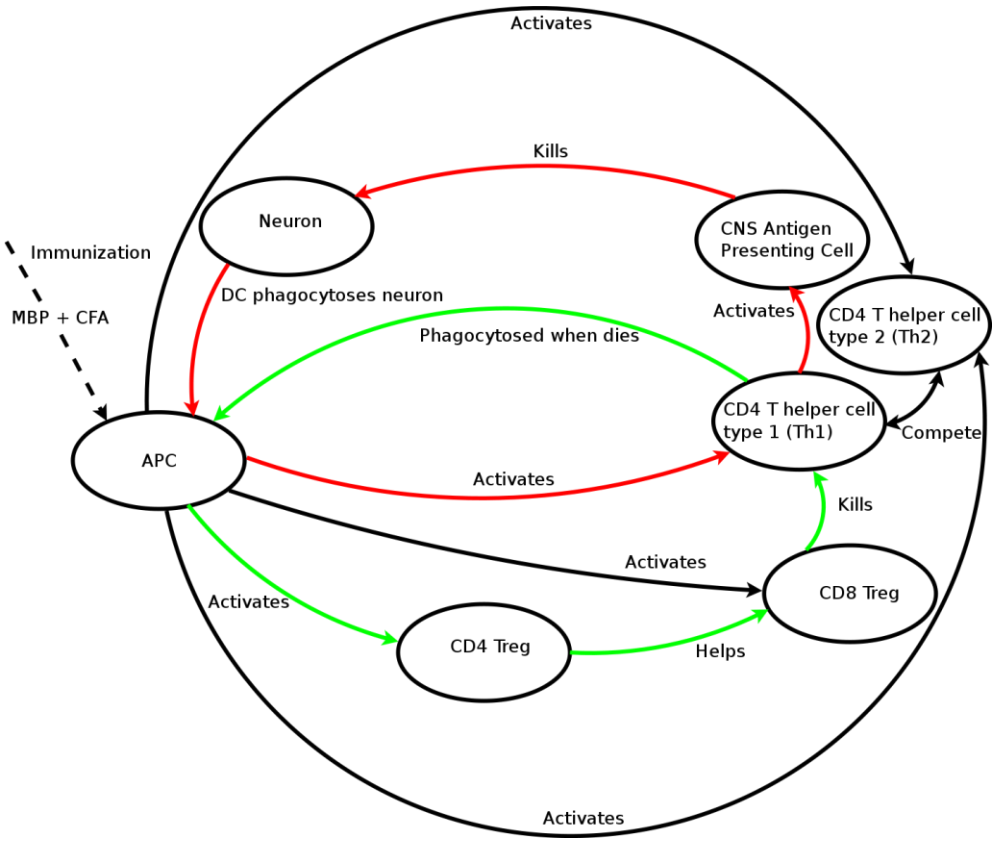

Figure 2.4: Median System Wide Effector T-cell populations in a series of 1000 runs of the baseline simulation. The curves show the progression of the system through a disease state (days 0 to 20) during which the populations of CD4 Th effectors rises to a peak and falls again. By day 20 the recovery cycle has been initiated and the CD4 Treg are approximately half-way through their response. By day 50 all the effector populations have essentially returned to resting levels. This figure was generated from data produced by re-running the baseline simulation [Read *et al.* 2009a].

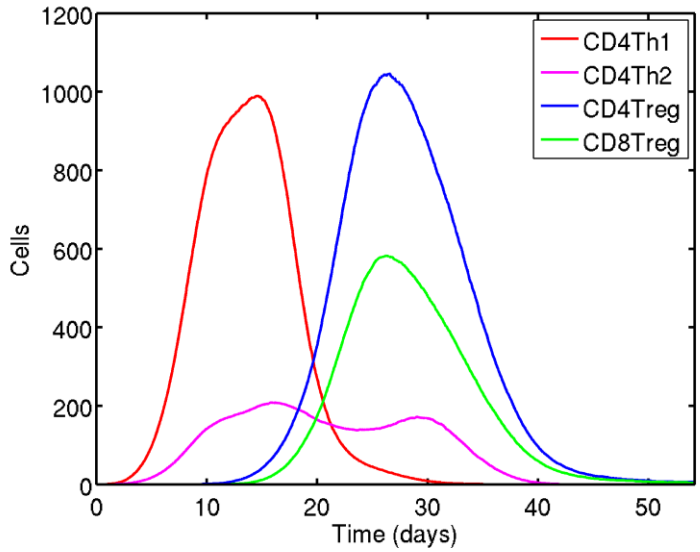

Subsequently a simulator experiment in which DCs could not simultaneously present MBP[12] and Th1-derived peptide[13] fragments ('mutually exclusive peptide presentation') demonstrated that very few DCs prime both CD4 Th1 and CD8 Treg (see Figure 2.5) [Williams 2010b, Williams *et al.* 2011]. The patterns of CD8 Treg (10% reduction) and CD4 Th1 priming and of CD4 Th1 apoptosis (5% reduction) were perturbed very little from that observed in the baseline experiment (illustrated in Figures 2.6, 2.7 and 2.8 respectively) due to the very small change in the pattern of peptide expression by the DCs. Only 5 DC were found to present both MBP and Type 1 peptides in the baseline simulation, the vast majority of DC not presenting any antigen at all.

Williams subsequently demonstrated a dramatic rise in the peak population of CD8 Treg, as compared to the baseline experiment, when the CD4 Treg population was abrogated in the simulation and the DCs were allowed to constitutively express the MHC compound, Qa-1 [Williams 2010b]. The elevated population of CD8 Treg can be compared to the baseline level in Figure 2.9 below.

[12] MBP is Myelin Basic Protein, a component of myelin which insulates nerve fibres in the Central Nervous System.
[13] These peptides, also referred to as 'Type 1' peptides, are the digestion products of the CD4 Th1 T-cell receptors once they have been phagocytosed by Antigen Presentingn Cells.

Since there is an elevated population of CD8 Treg effectors, there must also be more priming of CD8 Treg than in the baseline simulation. This priming takes place in the spleen as illustrated in Figure 2.10 below. Similarly, one would expect more CD4 Th1 apoptosis by CD8 Treg - this occurs in the circulation and in the Cervical Lymph Node (CLN) as illustrated in Figure 2.11.

An investigation of the causes of this significant rise in CD8 Treg population under CD4 Treg abrogation forms the basis of the subsequent experimentation described in Chapter 3 of this thesis.

Figure 2.5: The expression of peptides by dendritic cells in the mutually exclusive peptide presentation (MEPP) and baseline experiments. The curves are computed from median data obtained from a set of 1000 runs of the simulator. In the plot, the mutually exclusive peptide presentation (MEPP) experiment frequency curves are shown in solid lines while those for the baseline (Baseline) are shown as dashed lines for reference. MBP is myelin basic protein, a component of the myelin sheath from the CNS. Type 1 refers to peptides derived from type 1 Th cells i.e. type 1 peptides are CDR1/2 and Fr3. This figure was generated from data produced by re-running the mutually exclusive peptide presentation experiment [Williams 2010b].

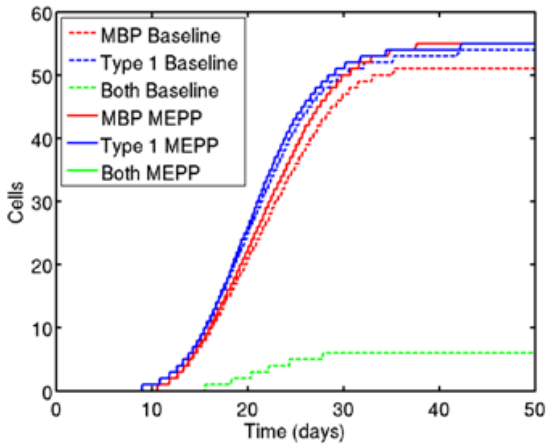

Figure 2.6: Occurrence of CD8 Treg Priming by Compartment in the mutually exclusive peptide presentation and baseline experiments. The curves are computed from median data obtained from a set of 1000 runs of the simulator. In the plot, the mutually exclusive peptide presentation experiment population (MEPP) curves are shown in solid lines while those for the baseline are shown as dashed lines for reference. This figure was generated from data produced by re-running the mutually exclusive peptide presentation experiment [Williams 2010b].

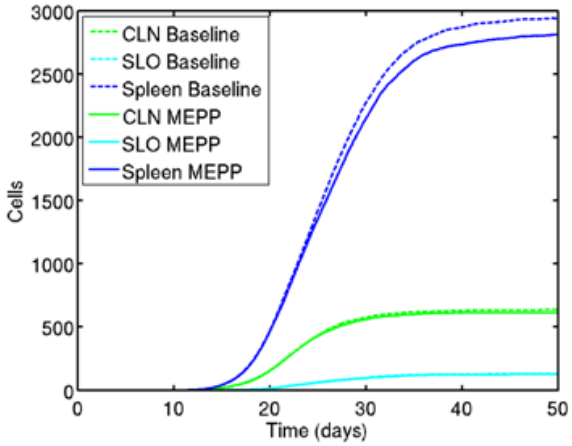

Figure 2.7: Occurrence of CD4 Th1 Priming by Compartment in the mutually exclusive peptide presentation and baseline experiments. The curves are computed from median data obtained from a set of 1000 runs of the simulator. In the plot, the mutually exclusive peptide presentation (MEPP) experiment population curves are shown in solid lines while those for the baseline are shown as dashed lines for reference. This figure was generated from data produced by re-running the mutually exclusive peptide presentation experiment [Williams 2010b].

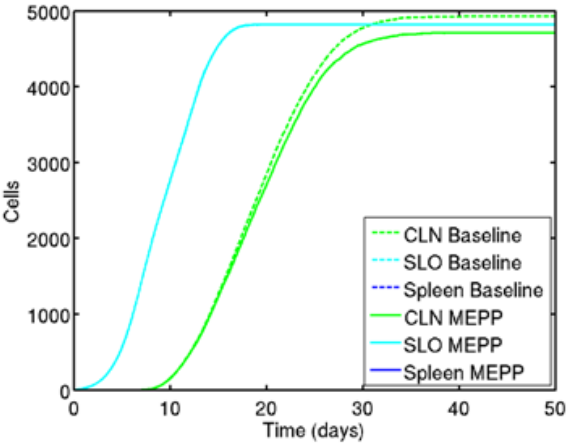

Figure 2.8: Occurrence of CD4 Th1 Apoptosis by Compartment in the mutually exclusive peptide presentation and baseline experiments. The curves are computed from median data obtained from a set of 1000 runs of the simulator. In the plot, the mutually exclusive peptide presentation (MEPP) experiment population curves are shown in solid lines while those for the baseline are shown as dashed lines for reference. This figure was generated from data produced by re-running the mutually exclusive peptide presentation experiment [Williams 2010b].

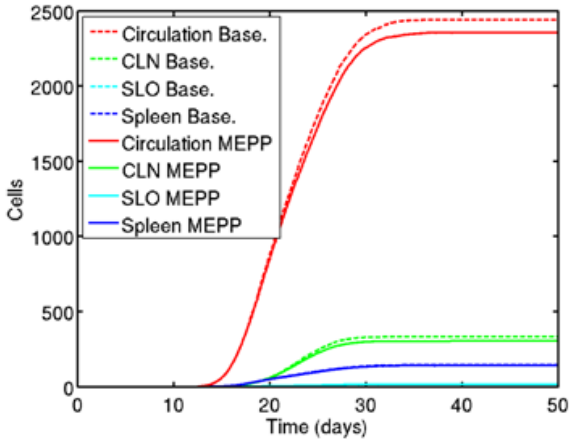

Figure 2.9: System wide effector T-cell populations under the CD4 Treg Abrogation and baseline experiments. The curves are computed from median data obtained from a set of 1000 runs of the simulator. It can be seen that the peak population of CD8 Treg is much higher (~1,000 cells) in the abrogation experiment than in the baseline case described above (~600 cells). The CD4 Treg abrogation (A.) experiment results are shown in solid lines and the baseline (B.) experiments in dashed lines for comparison. There is, of course, no solid blue line as the CD4 Treg population has been removed from this experiment. This figure was generated from data produced by re-running the CD4 Treg abrogation experiment [Williams 2010b].

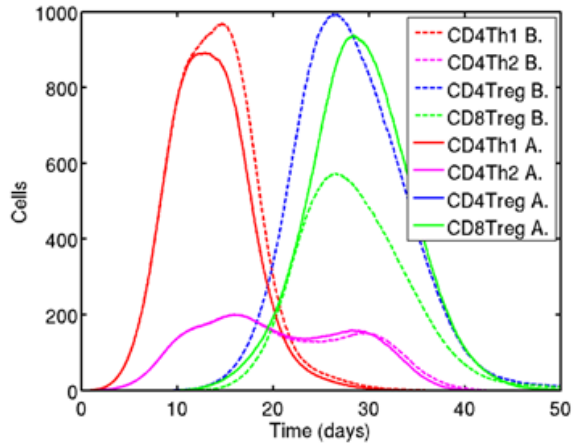

Figure 2.10: Occurrence of CD8 Treg Priming by Compartment in the CD4 Treg abrogation and baseline experiments. The curves are computed from median data obtained from a set of 1000 runs of the simulator. In the plot, the abrogation experiment population curves are shown in solid lines while those for the baseline are shown as dashed lines for reference. This figure was generated from data produced by re-running the CD4 Treg abrogation experiment [Williams 2010b].

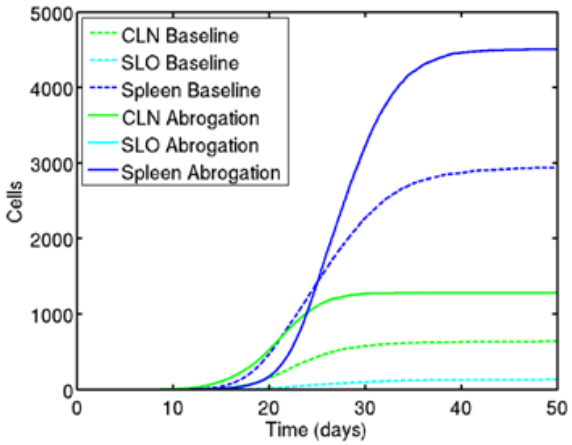

Figure 2.11: Occurrence of CD4 Th1 Killing by CD8 Treg by Compartment in the CD4 Treg abrogation and baseline experiments. The curves are computed from median data obtained from a set of 1000 runs of the simulator. In the plot, the abrogation (Abr.) experiment population curves are shown in solid lines while those for the baseline (Bas.) are shown as dashed lines for reference. This figure was generated from data produced by re-running the CD4 Treg abrogation experiment [Williams 2010b].

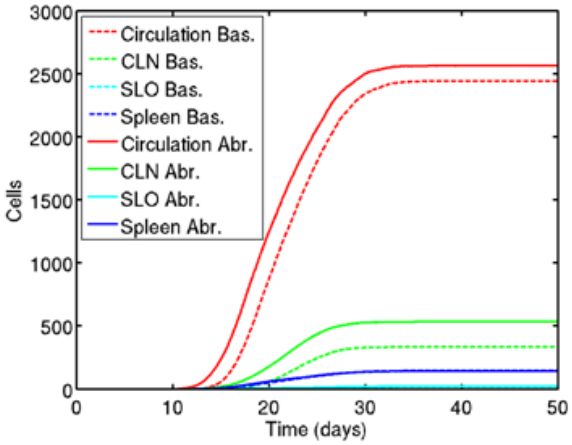

Chapter 3: Explaining the Doubling of CD8 Treg Population on Abrogation of the CD4 Treg Population

3.1 Introduction

The experiments presented seek to explain the observation of elevated CD8 Treg population when we maintain the conditions necessary to CD8 Treg priming, but abrogate the CD4 Treg population from the simulation [Williams 2010b].

To begin, we briefly outline the domain knowledge relevant to the work presented (Section 3.2). Then we develop two hypotheses that explain the observations and state these in a testable form (Section 3.3) before outlining the experimental procedure that we propose to adopt to test the two hypotheses (Section 3.4). Section 3.4 also describes the non-parametric effect size A-Test [Vargha and Delaney 2000], which we use to assess the significance of the effect that a particular parameter change has on simulation behaviour. We then briefly note the implications of what we aim to undertake in terms of the CoSMoS process (Section 3.5). In Section 3.6 we describe the experimentation conducted in the testing of the first hypothesis and in Section 3.7 we present an initial, unsuccessful attempt to address the second hypothesis. The results have, never-the-less, been presented as they give indications of the T-cell population dynamics around DCs during simulation. In Section 3.8 we address the issue of the significance of CD4 / CD8 Treg competition in determining the peak population of CD8 Treg effectors. In Section 3.9 we present data, in the form of A-Test scores, to draw conclusions about the significance of the two hypotheses in determining CD8 Treg population size. Section 3.10 is an experimental postscript and briefly details a further verification of the significance of spatial competition in determining CD8 Treg population. We draw our final conclusions in Section 3.11.

3.2 The Domain

In Chapter 2 we saw how EAE is mediated by MBP-reactive CD4 T-Helper cells [Pender 1995, van den Bark 1985]. These cells become able to access the CNS due to the increased permeability of the BBB [Mostarica-Stojkovic *et al.* 1992, Claudio *et al.* 1990, Baron *et al.* 1993]. Once inside the CNS, they promote the destruction of neurons via the activation of TNF-α secretion by macrophages that have also infiltrated the CNS [Huitinga *et al.* 1990, Pender 1995] and by the resident macrophages of the CNS, the microglia [Carson 2002].

Ultimately the MBP-reactive CD4 T-helper cells in the CNS reach the end of their cell cycle and become apoptotic. They can then migrate out of the CNS where they are phagocytosed by DCs [Pender 1995] which process the dead T-Helper cells and present the Complementarity Determining Region (CDR1/2) and Framework region 3 (Fr3) segments from their T-cell receptor Vβ8.2 chains [Kumar and Sercarz 2001, Tang *et al.* 2005].

The MHC-II-Fr3 complex presented on the DC surface can be recognised and bound by regulatory CD4 Treg [Kumar *et al.* 1996] which then stimulate (or 'license') the DC for production of the MHC Class Ib compound [Cantor *et al.* 1978]. CD8 Treg which recognise Qa-1 – CDR1/2 complex presented on the DC surface, can then bind to DC [Tang *et al.* 2005]. Once bound to the licensed APC, CD8 Treg are activated, with further help from IFN-γ secretion by CD4 Treg, and proliferate [Kumar and Sercarz 2001, Pender 1995].

When CD8 Treg become effector cells, they dissociate from the APC and become mobile. The effectors can recognise the Qa-1 – antigen complex [Kumar and Sercarz 2001] presented on the surface of CD4 T-Helper cells and kill them via the perforin pathway [Beeston *et al.* 2010, Pender 1995]. Consequently, the CD4 Th population falls during the recovery phase of EAE and within 30 days the populations of CD4 T-Helper cells have returned to their normal resting levels. Within ~50 days the full cycle is completed, with Treg populations falling back to their resting levels, and recovery is essentially complete [Pender 1995].

The necessity for CD4 Treg in licensing DC has been demonstrated by Kumar *et al.* 1996 who showed that the inactivation of CD4 Treg cells, via binding of a monoclonal antibody specific for the Vβ8.2 T-cell receptor protein, produced mice that suffered an increase in severity and duration of disease and that did not fully recover from EAE.

Similarly, an experiment that demonstrated the use of the perforin pathway by CD8 Treg to apoptose CD4 Th1 cells, showed that perforin-deficient mice suffered more severe EAE, due to the inability of CD8 Treg to apoptose the CD4 Th1 cells which mediate the disease [Beeston *et al.* 2010].

In summary, CD4 Treg cells are required, *in vivo*, to license DCs to express the MHC Class Ib compound Qa-1. Qa-1 is essential for the presentation of the CDR1/2 antigen from the MBP-reactive CD4 Th1 T-cell receptor. This complex is recognised by CD8 Treg cells and is therefore essential for their binding to DCs and their subsequent activation by them and hence to the role of CD8 Treg in cell-mediated regulation of autoimmunity (illustrated in Figure 3.1 below) [Tang *et al.* 2005].

Figure 3.1: Cartoon representation of the activation of CD8 Treg by dendritic cells that have been licensed for Qa-1 expression by CD4 Treg. In this process, an APC presenting MHC Class II – Fr3 complex can be bound by CD4 Treg whose T-cell receptors recognise it. The CD4 Treg provide the stimulus necessary for the APC to begin expressing the MHC compound Qa-1. The Qa-1 – CDR1/2 complex is recognised and bound by the TCR on CD8 Treg which then become activated. The CD8 Treg receives further stimulatory assistance from the CD4 Treg population via the secretion of IFN-γ. Activated CD8 Treg can then proceed to recognise and bind Qa-1 – CDR1/2 on the surface of the CD4 Th1 that they subsequently apoptose during cell-mediated regulation of autoimmunity. The cup-like projections on the surfaces of the Treg cells depict the T-cell receptors of these cells. 'Cross-Priming' is a process in which a CD8 T-cell response is initiated towards an antigen not synthesised by the APC. Cross-priming is more fully described in Appendix A.3. This figure is adapted from Figure 1 of Tang *et al.* 2005.

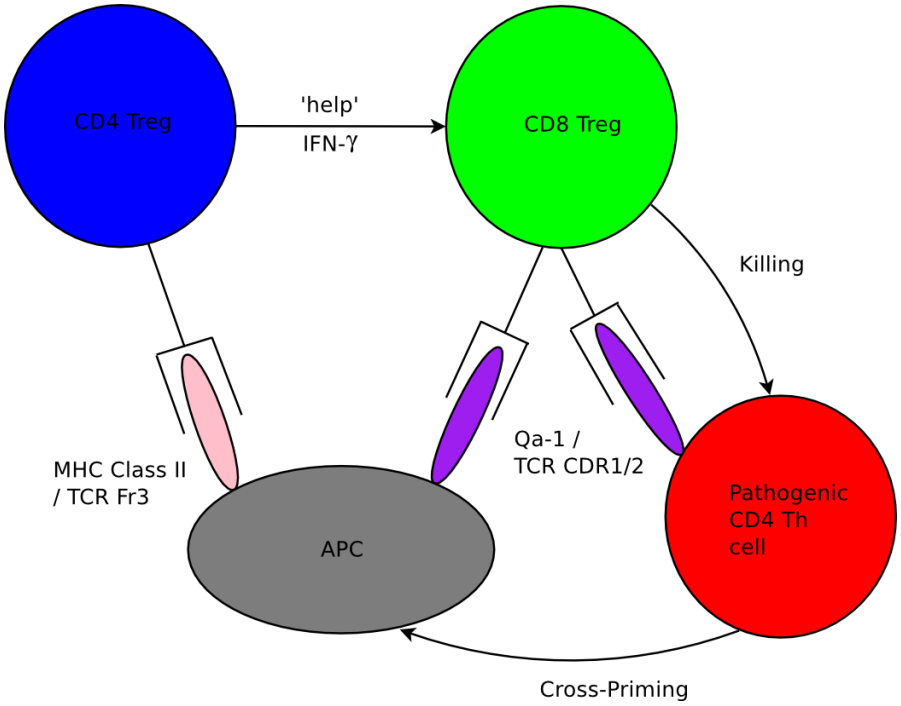

3.3 Motivation for Further Experimentation

As shown in the previous chapter (Section 2.10), earlier work with the EAE Simulator demonstrated that abrogating the CD4 Treg cell population resulted in an approximate doubling of the peak population of CD8 Treg cells provided that DCs were permitted to constitutively express Qa-1 upon maturity [Williams 2010b].

It has been demonstrated that the abrogation of the CD4 Treg population in mice has a negative impact on their recovery from EAE [Kumar *et al.* 1996]. This raises the question as to whether the CD4 Treg population serves to modulate the generation of CD8 Treg and thus the recovery from EAE.

The CD4 Treg abrogation experiment [Williams 2010b] demonstrated the advantage of being able to conduct speculative experiments *in silico* – Williams was able to maintain the conditions needed for CD8 Treg priming whilst removing the CD4 Treg from the system. The work presented here represents a further investigation of this system in an attempt to explain Williams' observations.

The 'age of licensing' hypothesis[14] concerns the constitutive expression of Qa-1 by the DC population. In the baseline simulation it takes ~60 hours (2.5 days) for CD4 Treg cells to mature into effectors and license DC for Qa-1 expression. In Williams' experiment [Williams 2010b], constitutive expression of Qa-1 facilitates earlier priming of CD8 Treg on DC because there is now no delay in waiting for licensing. The earlier priming could, in theory, permit an extra proliferative burst by CD8 Treg as the time that a CD8 Treg must wait to become proliferative is ~19 hours (0.8 days). The hypothesis can be expressed more formally:

> H1: the timing of Qa-1 expression by DC has a significant effect on the size of the CD8 Treg effector population.

The 'spatial saturation' hypothesis concerns the binding space available around the DCs. Both CD4 and CD8 Treg prime on the same DCs and so compete for binding space. Since CD4 Treg may be priming immediately upon DC maturation and CD8 Treg must wait for the DC to be licensed, CD4 Treg would be expected to occupy a greater fraction of the total space around any given DC. If spatial competition is a limiting factor in CD8 Treg population expansion, then removal of the CD4 Treg from the simulation should permit a greater growth in CD8 Treg numbers. Again, the hypothesis can be expressed more formally:

> H2: the competition between CD4 and CD8 Treg has a significant effect on the size of the CD8 Treg effector population.

3.4 Assessing the Hypotheses

Both of the hypotheses described could contribute to the observed rise in CD8 Treg peak population size under CD4 Treg abrogation. We propose two lines of experimentation to assess the relative influence of the hypotheses on increased CD8 Treg population size.

[14] Named because the hypothesis chiefly concerns the timing of Qa-1 expression by dendritic cells

3.4.1 The 'Age of Licensing' Experiment

The aim of this experiment is to gain some insight into how long it takes for DCs to become licensed for Qa-1 expression in a typical baseline simulation. This time approximates the total time taken for CD4 Treg cells to mature into effectors and license DCs. This information would then be used to parameterise an enforced delay in the constitutive expression of Qa-1 by DCs in a CD4 Treg abrogation experiment, thus simulating the DC's wait for a CD4 Treg to license it, without needing to explicitly include the CD4 Treg in the simulation model.

It is a trivial matter to record the times at which key life events occur for each DC in the simulation. We have chosen to record the ages of DCs at the time they become licensed for Qa-1 expression so that we can investigate the spread of values exhibited during a typical set of simulation runs.

We were then able to implement a time delay in the constitutive expression of Qa-1 under abrogation of the CD4 Treg population. In this way we would be able to simulate a system where CD8 Treg have to wait to be licensed but do not have to compete for binding space around the DC. The CD8 Treg population size data collected from this 'Qa-1 delay' experiment could then be compared statistically with that from the original (Qa-1 expression not delayed) CD4 Treg abrogation experiment in order to assess the significance of the timing of Qa-1 expression (i.e. delayed versus immediate) on CD8 Treg population size. The method used to assess the statistical significance of differences between our various experiments is discussed in more detail in Section 3.4.3.

3.4.2 The 'Spatial Saturation' Experiment

The aim of this experiment is to examine how the spatial saturation around the DCs at the time of apoptosis changes between simulator runs performed with immediate constitutive Qa-1 expression and runs performed with delayed constitutive Qa-1 expression. The runs with delayed constitutive expression of Qa-1 were also compared to the baseline runs which had been conducted as a control experiment.

By comparing these experiments using the A-Test, we aim to gain insight into whether removing the CD4 Treg population from the simulation advantages the CD8 Treg in terms of removing their need to compete for DC binding space.

To achieve these goals we implemented methods in the appropriate classes of the simulator to record the neighbours of each DC at its time of apoptosis, the time when DCs are most probably at their maximal spatial saturation, as judged from previous simulation behaviour. Implementation specific details of the changes made to the platform model are presented in Appendix C.1.

We were then able to record the saturation data for DCs in the same baseline runs that we used to derive parameterization of the Qa-1 expression delay. Ultimately we implemented the Qa-1 expression delay and ran more simulator runs under CD4 Treg

abrogation to simulate a system in which DCs still need to wait to be licensed for Qa-1 expression, but in which CD8 Treg do not need to compete for binding space around them.

In calculating the proportion of binding space occupied by a particular cell type one has to be aware of how that space is defined. The simulator is composed of a number of separate simulation compartments (see Appendix B.4). Each of the compartments consists of a lattice of grid points, each of which has the capacity to hold a certain number of cells and each of which has 8 neighbouring cells because the lattice is two-dimensional.

In our current parameterization, one grid point can accommodate 1 DC or 7 T-cells. This means that a soft limit of 56 T-cells may surround any given DC when it is at maximum spatial saturation. Therefore to calculate the proportion of binding space occupied by CD8 Treg we simply counted the number of CD8 Treg neighbouring an apoptotic DC and converted this figure into a percentage of binding space occupied.

3.4.3 Assessing the Relative Importance of the Two Hypotheses

The ultimate aim of the experiments proposed above, is to assess the relative significance of the two factors i.e. timing of Qa-1 expression on DC (and hence, indirectly, of activation of CD8 Treg) and spatial competition with CD4 Treg on the peak population size attained by CD8 Treg.

The simulator readily provides the user with sets of output data detailing the populations of key cell types throughout the simulation. Thus, we need some suitable means to reliably compare the data derived from different experiments.

For each experiment we ran 1000 runs of the simulator and we then calculated median data for the set of 1000 runs (the mean data set was not appropriate as the data did not appear to be normally distributed). We assumed that we were not dealing with normally distributed data and so tests such as Student's t-test were not appropriate and we chose to use Vargha and Delaney's implementation of the A-Test [Vargha and Delaney 2000] first described by Mann and Whitney [Mann and Whitney 1947].

The test is appropriate for working with non-parametric data and for assessing the magnitude of an effect i.e. the impact of changing one parameter value between a pair of experiments whilst holding all other parameter values constant.

The effect sizes corresponding to the various A-Test scores are detailed for the interested reader in Appendix D. For current purposes we have taken A-test scores below 0.29 and above 0.71 to represent scientifically significant effects.

3.5 Implications of the Experiments in Terms of the CoSMoS Process

As we were not intending to modify the logic of the actual simulation itself i.e. we were not introducing any new agent types, permitting changes to state transitions or permitting existing agents to access compartments that they had not previously been permitted to access, we did not need to change the domain model [Read 2011].

However, the simulation platform was modified as we needed to implement additional methods to capture data that already existed within the simulation, for example timings of key cellular events. New data logger and data store objects were implemented for each experiment. These were simply added to the appropriate existing packages in the simulator.

Required code changes were implemented on top of the existing simulator code. Code visualization, editing and compilation were performed using the Eclipse Integrated Development Environment[15].

We have not presented details of the exact changes made to the code. However, these are outlined for the interested reader in Appendix C.1 along with brief comment on the testing process.

3.6 The 'Age of Licensing' Experiment

3.6.1 Collecting Age of Licensing Data from the Baseline Experiment

Before proceeding to collect timing data from the simulator we needed to verify that coding changes had not altered the behaviour of the simulator. 1000 simulator runs were performed in order to verify that the baseline behaviour of the simulator would still be returned when not performing CD4 Treg population abrogation with or without implementing a Qa-1 expression delay.

Plots of the system-wide T-cell effector population levels during the simulation suggested that the baseline behaviour of the simulator had not been disturbed by the coding changes (data presented in Appendix C.3). Having verified that these simulation runs correctly reproduced the baseline, the data produced was used to represent our baseline case data.

3.6.2 Reproducing the CD4 Treg Abrogation Experiment Results

It was then important to verify that CD4 Treg abrogation still functioned properly in the augmented simulator. 1000 simulator runs were performed with CD4 Treg abrogation switched on and the delay in constitutive expression of Qa-1 switched off. The earlier CD4 Treg abrogation results [Williams 2010b] were reproduced by our verification runs (data presented in Appendix C.4). Having verified that these

[15] http://www.eclipse.org

simulation runs correctly reproduced the original CD4 Treg abrogation experiment, the data produced was used to represent our CD4 Treg abrogation case data.

3.6.3 Distribution of DC Ages at Time of Licensing

The baseline runs performed in section 3.6.1 were used to derive parameters for the delay in constitutive Qa-1 expression by DCs. In all, 1000 runs were performed and the spread of DC ages at the time of becoming licensed for Qa-1 expression was analysed.

At first analysis was restricted to a small sample of 5 runs, but the spread of licensing ages across the 5 runs appeared to be too wide to permit the extraction of any meaningful information. It was therefore decided to include data from the full 1000 runs in our investigation. To permit this, analysis and graph plotting was carried out in Matlab[16] via scripts written for the purpose.

First we analysed the median age at licensing across the 1000 simulator runs. The median age of DCs at licensing from these 1000 runs was 3.06 days (73.3750 hours). The inter-quartile range was 2.25 days (54 hours). Here, age statistics have been presented in days and in hours because it is more common in immunology to measure events in days, whereas the simulator parameter files contain cell timing parameters expressed in hours. The distribution of DC ages at licensing for the 1000 baseline runs is shown below in Figure 3.2:

[16] http://www.mathworks.com

Figure 3.2: Histogram showing the distribution of dendritic cell ages at the time of licensing for Qa-1 expression across a set of 1000 baseline simulator runs. The data has been normalised by the total DC count across the 1000 simulator runs in order to facilitate direct comparison with the distributions presented later on in the discussion.

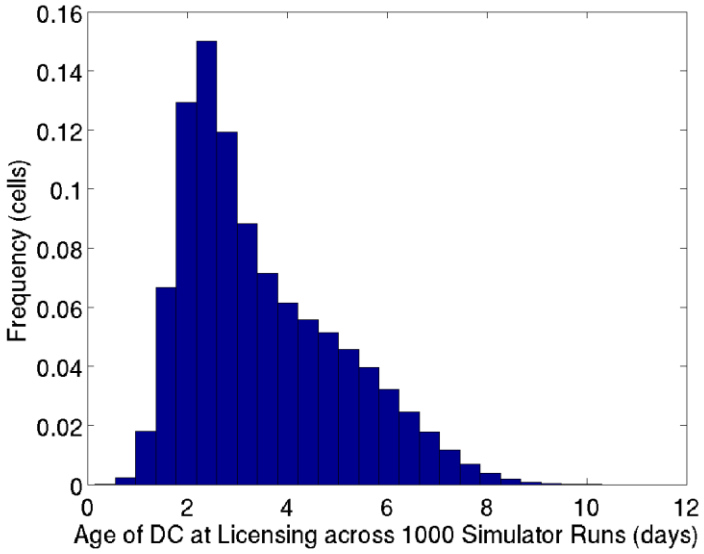

We investigate the possibility that the age of licensing of DCs may be related to the time of creation of those DCs. The simplest means to achieve this is to split the simulation data set for the entire DC population into data from DCs created earlier on in the simulation and data from those created later in the simulation. These two sets of DCs might be expected to exhibit different ages at the time of licensing because early on in the simulation there are fewer CD4 Treg present and therefore the chances of a DC-CD4 Treg encounter productive for licensing are much diminished compared to those later on in the simulation when the CD4 Treg population is reaching its peak.

To investigate this possibility we arbitrarily divided the data for each baseline simulation run into data derived from DCs with a creation time of 20 days (480 hours) or earlier (the 'early' population) and that derived from DCs created after 20 days (the 'late' population). The 20 day cut off time was chosen as it is the point on the baseline simulation system-wide T-cell effector population plot (presented for reference in Figure 3.3) where the CD4 Treg are approximately halfway through their immune response. This is also the point in the simulation when DCs begin to present peptides derived from CD4 Th1 apoptosed by CD8 Treg.

Figure 3.3: System-wide T-cell effector population levels throughout the baseline simulation to illustrate the choice of the 20 day cut-off used for separating the dendritic cell age data into two distinct sets. The approximate mid-point of the CD4 Treg response is indicated by the vertical dashed line at day 20.

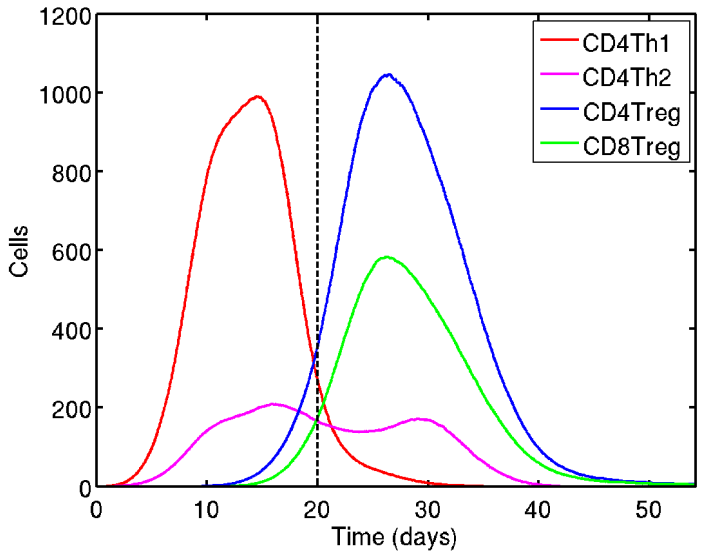

The separated data showed a clear[17] differentiation in medians between the 'early' and 'late' simulation DC. For the population of DCs created before 20 days the median age at licensing was 3.44 days (82.50 hours) with an inter-quartile range of 2.36 days (56.75 hours). For the 'late' population the median age at licensing was 2.48 days (59.50 hours) with an inter-quartile range of 1.45 days (34.875 hours). The distributions of the ages of licensing in the 'early' and 'late' DC populations are presented below in Figures 3.4 and 3.5.

[17] A clear, but not a significant difference. A comparison of the two populations via the A-test gave a score of 0.644 which indicates an effect of only medium size. However, as this arbitrary splitting of the data was performed solely to illustrate the idea that DC age at licensing varies across simulation time, this result has no impact on any further results within the thesis.

Figure 3.4: Histogram showing the distribution of dendritic cell ages at the time of licensing for Qa-1 expression in the 'early' dendritic cell population (i.e. those dendritic cells created prior to 20 days in the baseline simulator experiment) across a set of 1000 baseline simulator runs. The data has been normalised by the total 'early' population DC count across all 1000 simulator runs in order to facilitate direct comparison with the distributions presented later on in the discussion.

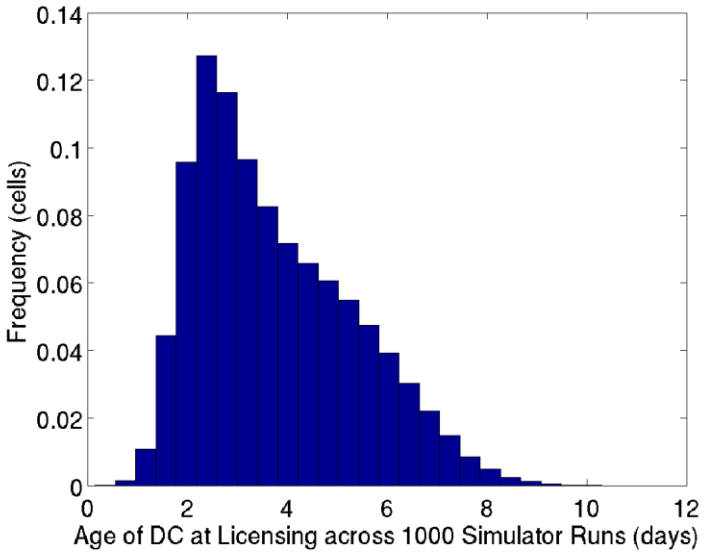

Figure 3.5: Histogram showing the distribution of dendritic cell ages at the time of licensing for Qa-1 expression in the 'late' dendritic cell population (i.e. those dendritic cells created later than 20 days in the baseline simulator experiment) across a set of 1000 baseline simulator runs. The data has been normalised by the total 'late' population DC count across all 1000 simulator runs in order to facilitate direct comparison with the distributions presented later on in the discussion.

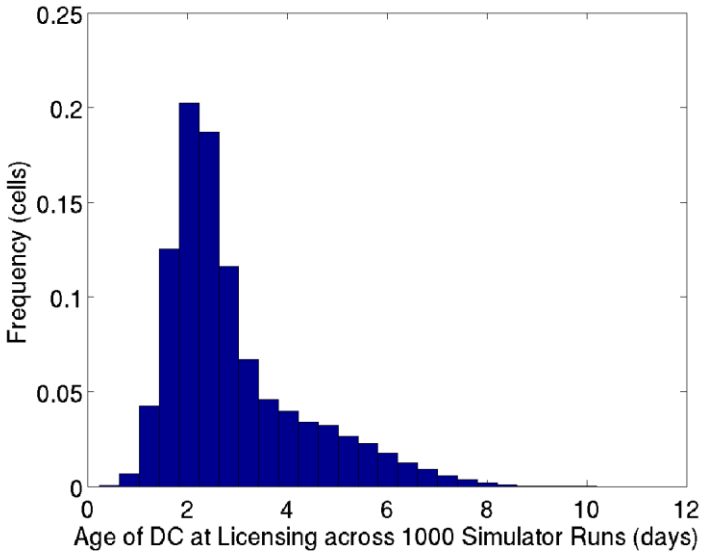

3.6.4 Choice of Parameterization

Since we are mostly interested in the immunology relevant to the late simulation environment (this is when CD8 Treg become activated by licensed DC), a mean delay in Qa-1 expression of 2.48 days (59.50 hours) with two standard deviations being equivalent to 1.83 days (43.803 hours)[18] was chosen to approximate the spread of licensing ages found for the baseline simulation.

It is worth noting at this point that the simulator calculates a future time for an event by drawing a time interval from a normal distribution of time intervals, expressed in hours, centred on a parameterised mean and with a spread determined by the parameterised standard deviation and adding this interval to the current time.

We chose to approximate the non-parametric distribution of DC ages obtained from the baseline simulation with a normal distribution[19] of mean equal to the median age and with an approximate standard deviation given by Equation 1. The distribution of DC ages at the time of licensing was reasonably well approximated by a normal distribution as illustrated in Figure 3.6 below:

Equation 1: 0.5 * Inter-Quartile Range * (68.2/50.0)

[18] These numbers were derived from the distribution of ages of licensing for Qa-1 expression of 'late' population DC in the baseline simulation described in Section 3.6.3.

[19] The comparison of the 'early' population distribution to its normal approximation gave an A-test score of 0.4573 which indicates no effect, whereas the comparison of the 'late' population data to its normal approximation gave a score of 0.5797 i.e. only a small effect. Therefore, neither approximation was significantly different from its corresponding real distribution. Therefore, we might reasonably anticipate that use of this approximation would not affect the results obtained in our simulations. However, to ensure that this is really the case, we should ideally re-run all the experiments in this section with a better approximation to the data.

Figure 3.6: Histogram illustrating the approximation of the experimental distribution (blue bars) of dendritic cell ages at the time of licensing for Qa-1 expression in the ‘late’ population dendritic cells by a normal distribution (red bars) of dendritic cell ages based upon it (the mean is set to the median of the experimental data and the standard deviation is approximated from the semi-inter-quartile range as described in the text). The data has been normalised by the 'late' population DC count across all 1000 simulator runs in order to facilitate direct comparison with the distributions presented later on in the discussion.

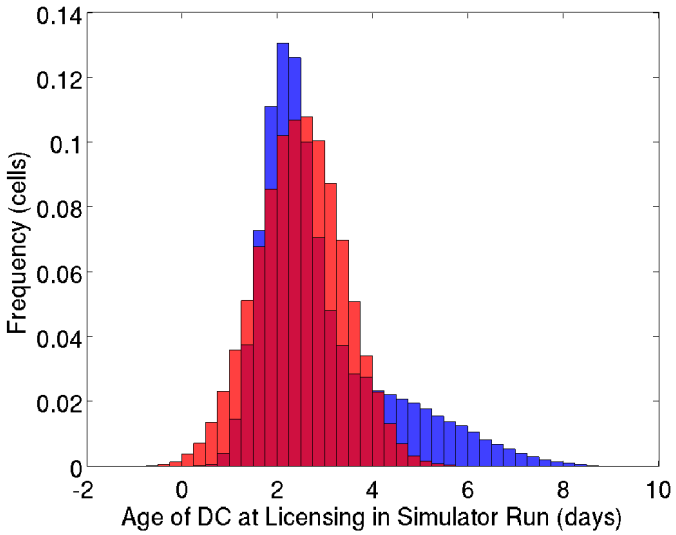

Finally, as a computationally cheap way of assessing the importance of our choice of expression delay parameters on the simulation results in general and the peak CD8 Treg population in particular, we implemented a Qa-1 expression delay with a larger mean of 3.44 days (82.50 hours) with two standard deviations set to 2.97 days (71.278 hours). The parameterisation for this distribution was taken from the distribution of licensing ages in the ‘early’ population DC in the baseline simulation. The comparison of the non-parametric distribution of DC licensing ages with the normally distributed approximation to it is presented in Figure 3.7 below.

Figure 3.7: Histogram illustrating the approximation of the experimental distribution (blue bars) of dendritic cell ages at the time of licensing for Qa-1 expression in the ‘early’ population dendritic cells by a normal distribution (red bars) of dendritic cell ages based upon it (the mean is set to the median of the experimental data and the standard deviation is approximated from the semi-inter-quartile range as described in the text). The data has been normalised by the 'early' population DC count across all 1000 simulator runs in order to facilitate direct comparison with the distributions presented later on in the discussion.

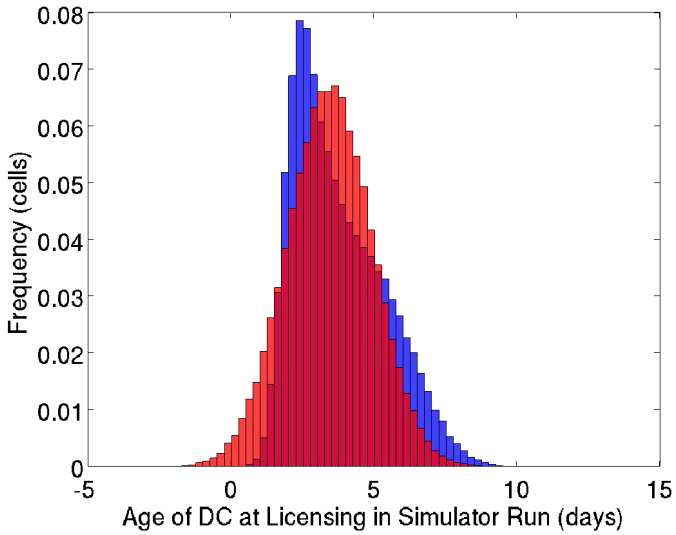

3.6.5 The Results Obtained by Implementing a Delay in Qa-1 Expression Under CD4 Treg Abrogation

3.6.5.1 The 'Qa-1 Delay' Experiment

Having chosen the timing parameters, the 2.48 day Qa-1 expression delay was turned on and 1000 runs of the simulator were performed under CD4 Treg abrogation. The cellular event timing data was briefly analysed to verify that the simulation was reproducing the anticipated range of ages at licensing, and this was the case (data presented in Figure 3.8 below).

Figure 3.8: The distribution of dendritic cell ages at the time of licensing (blue bars) in 1000 runs of the simulator implementing CD4 Treg abrogation with an enforced mean delay of 2.48 days in constitutive expression of Qa-1 (‘Qa-1 delay’) compared to the normal distribution used to generate the times of licensing for the dendritic cells in this simulation (red bars). The data presented has been normalised by the 'late' population DC count across the 1000 simulator runs to facilitate comparison of results in later discussion.

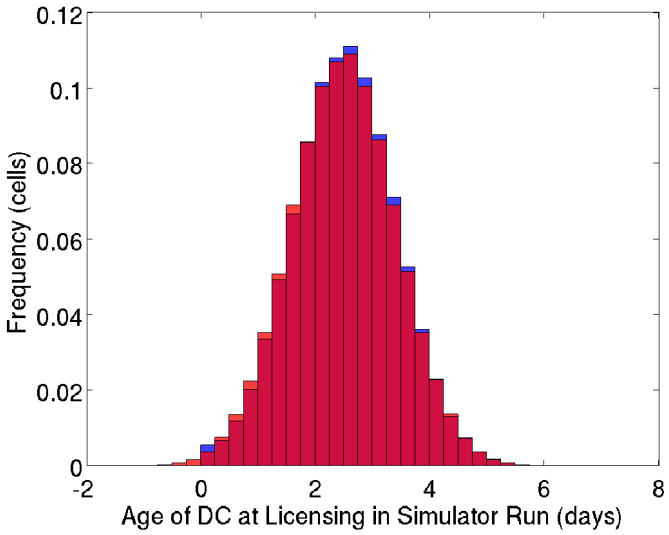

The system-wide effector T-cell population curves were of a similar form to the baseline curves, though peak populations of the different cell types (the CD8 Treg in particular) recorded had shifted from their baseline values as illustrated below in Figure 3.9. As anticipated the peak population of CD8 Treg was reduced from the approximate 1000 cells in the CD4 Treg abrogation experiment to around 750 cells under CD4 Treg abrogation coupled with an enforced delay in constitutive Qa-1 expression by DCs.

Figure 3.9: Median T-cell effector population levels ('responses') from 1000 simulator runs under CD4 Treg abrogation and employing a mean delay of 2.48 days in constitutive expression of Qa-1 by dendritic cells ('Qa-1 delay'). The peak population of effector CD8 Treg has been reduced to ~750 cells from the original CD4 Treg abrogation experiment peak of ~1,000 cells. This is still above the ~600 cell peak population of CD8 Treg effectors found for the baseline experiment.

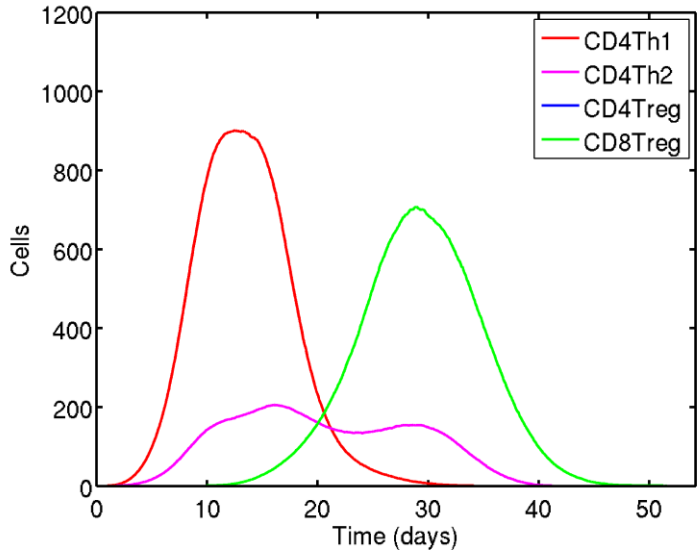

3.6.5.2 The 'Longer Qa-1 Delay' Experiment

When we repeated the simulation using a longer mean delay in constitutive Qa-1 expression of 3.44 days, the system-wide effector T-cell population curves were again of a similar form to the baseline curves (the effector T-cell population sizes are presented in Figure 3.10 below). The peak population of CD8 Treg was reduced from the approximate 1000 cells in the CD4 Treg abrogation experiment to around 400 cells under CD4 Treg abrogation coupled with the longer enforced delay in constitutive Qa-1 expression by DCs.

Figure 3.10: Median T-cell population responses from 1000 simulator runs under CD4 Treg abrogation and employing a mean delay of 3.44 days in constitutive expression of Qa-1 by dendritic cells ('longer Qa-1 delay'). The peak population of effector CD8 Treg has been reduced to ~400 cells from the original CD4 Treg abrogation experiment peak of ~1,000 cells. This is now lower than the ~600 cell peak population of CD8 Treg effectors found for the baseline experiment.

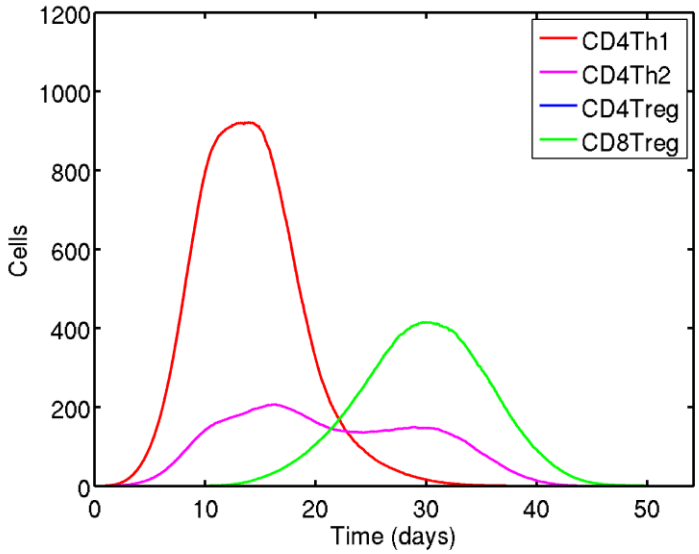

Again we verified that the observed distribution of DC ages at the time of licensing reproduced the normal distribution from which the times of licensing were drawn. This was again the case; the comparison of the two distributions is presented in Figure 3.11 below.

It should be noted that the apparently anomalous datum at age zero arises due to the way that the simulator calculates times for future events. The normal distribution used to parameterise our time for DC licensing actually extends nearly two days into the past and therefore, if these times are selected by the simulator they have to be rounded up to 0 or else the simulator would be trying to place the licensing event in the past.

Figure 3.11: Distribution of dendritic cell ages at the time of licensing (blue bars) in 1000 runs of the simulator implementing CD4 Treg abrogation with an enforced mean delay of 3.44 days in constitutive expression of Qa-1 (‘longer Qa-1 delay’) compared to the normal distribution used to generate the times of licensing for the dendritic cells in this simulation (red bars). The data presented has been normalised by the 'early' population DC count across the 1000 simulator runs to facilitate comparison of results in later discussion.

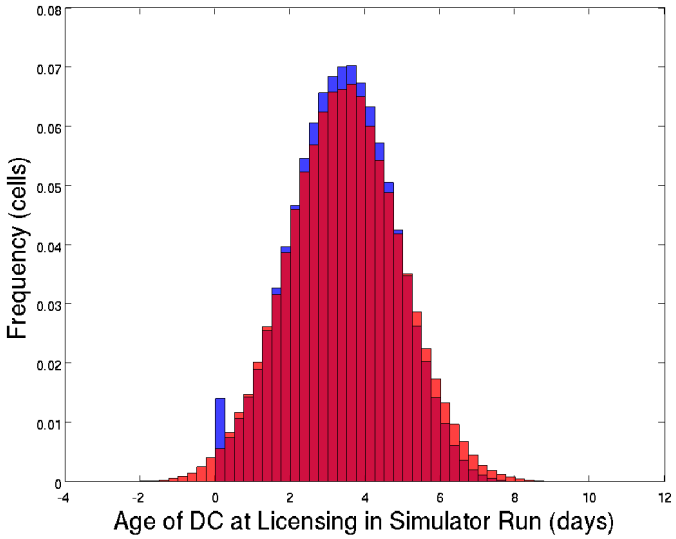

The significance of the changes in T-cell peak populations was assessed using the Mann and Whitney A-test described in section 3.4.3 [Mann and Whitney 1947, Vargha and Delaney 2000]. The CD8 Treg effector population in the ‘Qa-1 delay’ experiment showed a medium-sized increase as compared to the baseline case and a large decrease as compared to the CD4 Treg abrogation experiment. The ‘longer Qa-1 delay’ experiment showed a large decrease compared to both the baseline and CD4 Treg abrogation experiments.

For the CD4 Th1 cells the peak populations showed small changes overall. The number of Th1 apoptosed appeared to be more significantly affected, as expected as significant changes in the size of CD8 Treg population might reasonably be expected to have significant impact on the numbers of CD4 Th1 cells they could apoptose.

The experimental results and their significance are discussed in greater detail in Section 3.9 where we then attempt to assess the relative influence of the two hypotheses on CD8 Treg peak effector population. A short summary of the relevant experiments is presented in Table 3.1 below.

Table 3.1: Summary of the experiments compared during the investigation of the effects of variation in the timing of Qa-1 expression by dendritic cells. The table details the median ages of DCs at licensing in the various experiments along with the peak population attained by the CD8 Treg. Notes are provided to further identify each experiment in the text. All experiments except 'Qa-1 delay' showed a significantly changed CD8 Treg population as compared to the baseline simulation as assessed using the A-test.

Experiment	Peak CD8 Treg effector population	Notes
Baseline Median age of DC at time of licensing = 3.06 days	600	default behaviour of simulator as described in Read *et al.* 2009a.
CD4 Treg abrogation Median age of DC at time of licensing = 0.00 days (constitutive and immediate)	1000	as described in Williams 2010b.
'Qa-1 delay' Median age of DC at time of licensing = 2.48 days	750	parameterised from the distribution of ages of licensing of DCs created later than 20 days into the baseline simulation.
'Longer Qa-1 delay' Median age of DC at time of licensing = 3.44 days	400	parameterised from the distribution of ages of licensing of DCs created prior to 20 days into the baseline simulation.

3.7 The 'Spatial Saturation' Experiment

3.7.1 Conducting the 'Spatial Saturation' Experiment

The aim of this experiment was to provide a means by which we could compare experiments that employed CD4 Treg abrogation with those that did not. This was done in order to assess the impact of the spatial competition of CD8 Treg with CD4 Treg on the peak population of CD8 Treg effectors in the simulation, using the A-Test as described in section 3.4.3.

Our previous experiments concerned the 'age of licensing' hypothesis and contrasted systems where CD4 Treg abrogation had occurred and the expression of Qa-1 by DC had been delayed with those where Qa-1 expression was immediate. This experiment relates to the spatial saturation hypothesis and should allow us to compare systems where there are CD4 Treg present and those where they are not.

However, analysis of the spatial saturation data was problematical as there was great variation in the proportion of the binding space occupied around DCs as well as in the

numbers of DC apoptosed during the different simulator runs. This made it difficult to present coherent and meaningful datasets appropriate for use with the A-Test [Vargha and Delaney 2000]. No other suitable test was identified for demonstrating the significance of differences in the data. The data is, however, presented here because it provides insight into the population dynamics of T-cells during the simulation.

Initially, we attempted to examine the distribution of the spatial saturations of the DC in the 1000 baseline runs and the two sets of 1000 runs with CD4 Treg abrogation coupled with different delays in constitutive Qa-1 expression implemented (the 'Qa-1 Delay' and 'longer Qa-1 Delay' experiments described in Section 3.6.5).

After initial inspection of the resulting data, we investigated whether the distribution of neighbours around the DC changed over simulation time. This investigation is particularly appropriate here as a large proportion of the logged DCs that are created prior to the 20 day cut off express MBP, whereas those created later in the simulation tend to express Th1-derived peptides. Very few DC express both MBP and Th1-derived peptides. This is not unreasonable, in the simulation DCs are created expressing MBP and early on in the simulation have not had time to encounter CD4 Th1 cells and phagocytose them, the phagocytosis being a necessary preliminary step to the expression of Th1-derived peptides.

The data from the baseline and each ‘Qa-1 delay’ run was split into 'early' and 'late' DCs as in the previous experimental analysis and the space occupied around the DCs by each T-cell population calculated. As the data is complicated, only specific emerging trends are presented for the sake of simplicity and brevity.

3.7.2 Results of the ‘Spatial Saturation’ Experiment

Of particular note is the fact that the spatial saturation of the DCs expressing Th1-derived peptide is dominated by CD4 Treg in the baseline experiment in a ratio of roughly 2:1 with CD8 Treg (see Table 3.2 which summarises the histograms presented in Figure 3.12). The median spatial occupancy for CD4 Treg is around 60% and for CD8 Treg around 30%, but is irrelevant to the Qa-1 expression delaying experiments as CD4 Treg abrogation was implemented during these simulations.

The median spatial saturation by CD4 Th cells around MBP expressing DCs is 90% for the baseline runs and appears to be reduced in the experiments where we implemented a Qa-1 expression delay. It is not known why this saturation does not actually reach 100% in the baseline experiment. The two trends are illustrated in Figures 3.12a-f below.

A further trend that emerged from the data was that towards greater spatial saturation by Tregs with time i.e. DCs tend to become more 'crowded' by Treg toward the latter half of the simulation. In contrast, the DC expressing MBP tended to become slightly less saturated by CD4 Th1 as the simulation progressed (this was markedly more significant in the CD4 Treg Abrogation experiment). This could potentially arise

from the fact that later on in the simulation there is a greater population of CD8 Treg and thus a greater possibility of CD4 Th1 apoptosis occurring and in the CD4 Treg abrogation experiment there are significantly more CD8 Treg present. This could be confirmed from simulation by recording CD8 Treg-CD4 Th encounters resulting in apoptosis of CD4 Th. Figures 3.12a-f illustrate the trends in the baseline experiment. These trends were observed for DC expressing Th1-derived peptides in the baseline, the CD4 Treg abrogation and both Qa-1 expression delaying experiments (as illustrated for the ‘Qa-1 Delay’ experiment in Figures 3.13a-d).

Figure 3.12: Histograms illustrating the trend toward lesser spatial saturation of MBP-expressing dendritic cells by CD4 Th1 and toward greater spatial saturation of Th1-derived peptide expressing dendritic cells by Treg with simulation time across 1000 baseline simulator runs. The data presented has been normalised by the DC population count pertaining to the subset of interest (i.e. 'early' or 'late') across the 1000 simulator runs to facilitate comparison of results in later discussion. Panels a) and b) show the fraction of ‘early’ and ‘late’ population MBP-expressing DCs having a given proportion of their binding space occupied by CD4 T-helper cells. Panels c) and d) show the fraction of ‘early’ and ‘late’ population Th1 peptide-expressing DCs having a given proportion of their binding space occupied by CD4 Treg. Panels e) and f) show the fraction of ‘early’ and ‘late’ Th1 peptide-expressing population DCs having a given proportion of their binding space occupied by CD8 Treg.

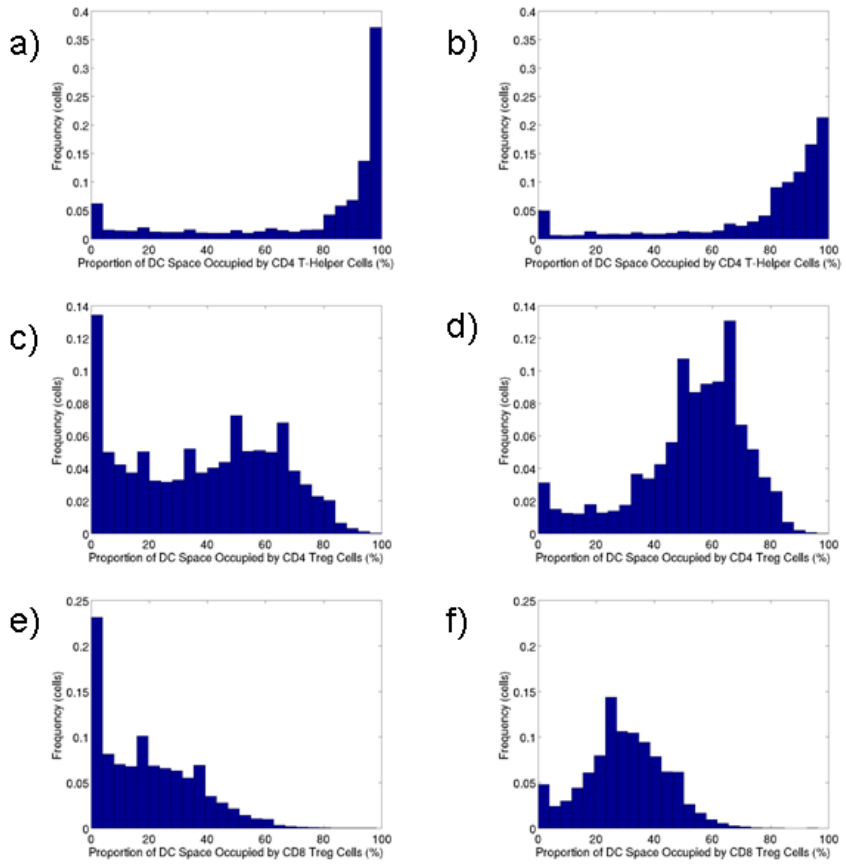

The Qa-1 expression delaying experiments ('Qa-1 Delay' and 'Longer Qa-1 Delay') showed similar trends to those described for the baseline experiment above (the trend toward lesser spatial saturation of DC by CD4 Th1 with time is illustrated in Figures 3.13a and b, while that for greater spatial saturation of DC with time by CD8 Treg is illustrated in Figures 3.13c and d).

Figure 3.13: Histograms illustrating the trends toward lesser spatial saturation of dendritic cells by CD4 Th1 and greater spatial saturation of dendritic cells by CD8 Treg with simulation time across 1000 simulator runs under CD4 Treg abrogation with a mean delay of 59.2 hours implemented in the constitutive expression of Qa-1 by dendritic cells. The data presented has been normalised by the DC population count pertaining to the subset of interest (i.e. 'early' or 'late') across the 1000 simulator runs to facilitate comparison of results in later discussion. Panels a) and b) show the fraction of 'early' and 'late' population MBP-expressing DCs having a given proportion of their binding space occupied by CD4 T-helper cells. Panels c) and d) show the fraction of 'early' and 'late' population Th1 peptide-expressing DCs having a given proportion of their binding space occupied by CD8 Treg.

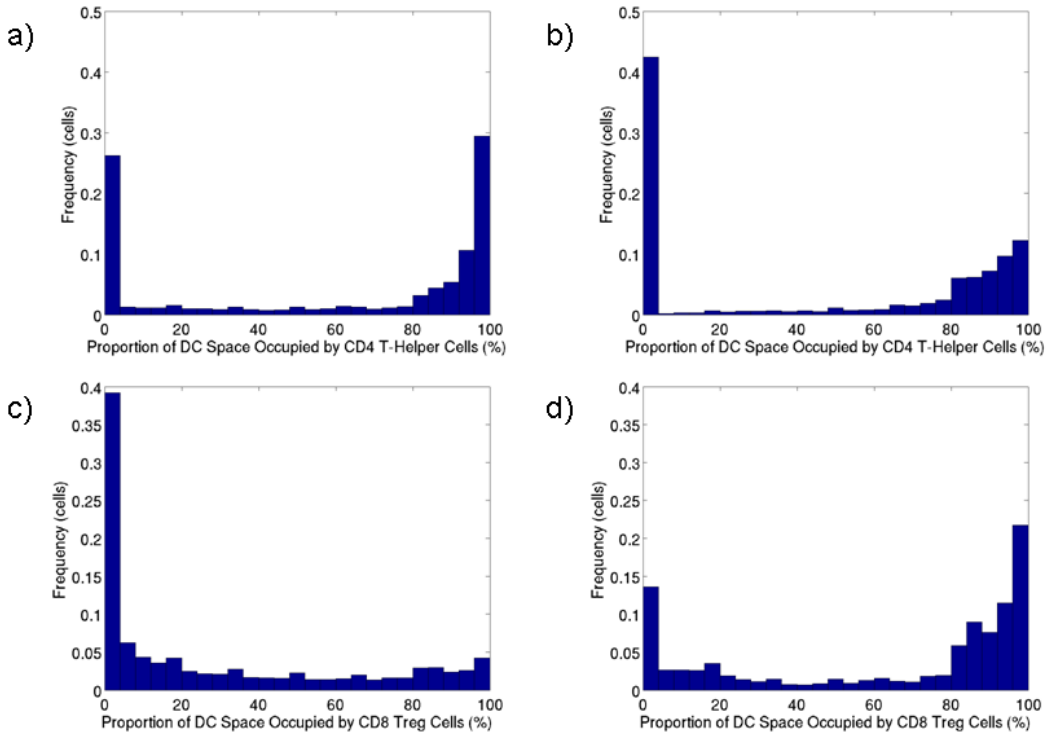

However, there were also some differences between the Qa-1 expression delaying experiments and the baseline. In the 'Qa-1 delay' experiment[20] the maximum proportion of binding space around DCs occupied by CD8 Treg is greater than in the

[20] Parameterised with data from dendritic cells created after 20 days into the baseline experiment

baseline. This should come as no surprise as in the baseline experiment the CD8 Treg have to compete for binding sites around the DC with the CD4 Treg population whereas in the Qa-1 delay experiments the CD4 Treg population has been abrogated. In contrast, the ‘Qa-1 Delay Longer’ experiment exhibited similar median saturation of DCs by CD8 Treg in the latter half of the simulation to that shown in the baseline, despite the fact that the CD4 Treg were abrogated in the Qa-1 delay experiment. A possible reason for this may be simply the lower population of CD8 Treg present in the simulation due to the longer wait for activation.

The median spatial saturation levels for the DC in the baseline, the CD4 Treg abrogation and the two Qa-1 expression delaying experiments are presented as a summary of the above in Table 3.2

Table 3.2: Summary of the median spatial saturations of apoptotic dendritic cells across 1000 runs of the EAE simulator in the baseline, CD4 Treg abrogation and Qa-1 expression delaying experiments. Median data has been presented since the distributions are non-normal meaning that presentation of means would be inappropriate. The median data serves to distinguish the type of cellular neighbourhood of DCs early in the simulation from that later in the simulation.

Experiment	**Cell type**	**Median proportion of space around DC occupied 'early' (%)**	**Median proportion of space around DC occupied 'late' (%)**
Longer Qa-1 delay (median DC age at licensing = 3.44 days)	CD4 Treg	0 (NA)	0 (NA)
	CD8 Treg	3.57	32.14
	CD4 Th	82.14	73.21
Qa-1 delay (median DC age at licensing = 2.48 days)	CD4 Treg	0 (NA)	0 (NA)
	CD8 Treg	8.93	73.21
	CD4 Th	89.29	72.21
CD4 Treg Abrogation (median DC age at licensing = 0.00 days)	CD4 Treg	0 (NA)	0 (NA)
	CD8 Treg	12.50	83.93
	CD4 Th1	83.93	57.14
Baseline (median DC age at licensing = 3.06 days)	CD4 Treg	39.29	57.14
	CD8 Treg	17.86	30.36
	CD4 Th	92.96	87.5

3.8 Demonstrating the Effect of Spatial Restriction about Dendritic Cells on the Peak Population of CD8 Treg

Given that the spatial saturation data described in Section 3.7 was unsuitable for comparison using the A-Test on account of the different sizes of dataset between the experiments, a different method was required to provide a comparison between experiments where CD8 Treg faced spatial competition from CD4 Treg and experiments where they did not. Ideally this method relies on comparing simulator time-series data, particularly the T-cell effector population levels with time, as these are guaranteed to be of the same size i.e. the number of simulation time steps. To this end further experiments are conducted in which the spatial restrictions applicable to CD8 Treg are modified as compared to other T-cell populations.

The simulator baseline behaviour is to allow T-cells to migrate to a grid point only if there are less than 7 T-cells already at that point. We selectively removed this test for saturation of the destination grid point for CD8 Treg migrating to this position. In two separate experiments we allowed CD8 Treg to continue to enter a grid point until there were first 14 and secondly 21 T-cells present in the grid point, the values 14 and 21 representing twice and three times the baseline limit on the number of T-cells occupying any one grid point. This allowed us to assess the behaviour of the system when CD8 Treg face significantly less spatial competition but still need to wait for DC to become licensed before they can be activated.

It is important to emphasise that the restrictions pertaining to T-cell migration are all that is changed between the baseline simulator and these two experiments. Expression of Qa-1 on DC occurs as for the baseline case and CD4 Treg are present in the simulation. The experiments relevant to the comparison of levels of CD8 Treg spatial restriction on peak CD8 Treg population are enumerated in Table 3.3 below. The author has adopted short-hand names for the two additional experiments in order to simplify the text and discussion.

Table 3.3: List of experiments facilitating the comparison of the effects of varying the spatial restrictions imposed on CD8 Treg.

Experiment	**Referred to in text as**
Baseline	baseline
Allow CD8 Treg to enter a grid point until there are 14 T-cells at that location	maximal crowding 14
Allow CD8 Treg to enter a grid point until there are 21 T-cells at that location	maximal crowding 21

As a result of the relaxation of the crowding criteria for CD8 Treg, but not other T-cell populations, the peak population of CD8 Treg was greatly increased compared to the baseline experiment. When allowing CD8 Treg to enter a grid space up to a total of 14 T-cells, the peak response rose to ~2,300 CD8 Treg cells. This figure rose even further, to ~3,000 cells in the 'maximal crowding 21' experiment. This is illustrated in Figure 3.14 a, b and c below and summarised in Table 3.4.

Figure 3.14: Effector T-cell population ('response') curves obtained under the baseline conditions and with modified rules for CD8 Treg migration – firstly, allowing CD8 Tregs to migrate into a grid point if there are fewer than 14 T-cells already present there and secondly, allowing CD8 Tregs to migrate into a grid point if there are fewer than 21 T-cells already present there. Panel a) shows the baseline experiment. Panel b) shows the results of the experiment with the maximal crowding parameter set to 14 and panel c) shows the results of the experiment with the maximal crowding parameter set to 21. All results were obtained from median data from samples of 1000 simulator runs. The maximal crowding parameter determines how many T-cells can migrate into a grid point and is more fully described in the text.

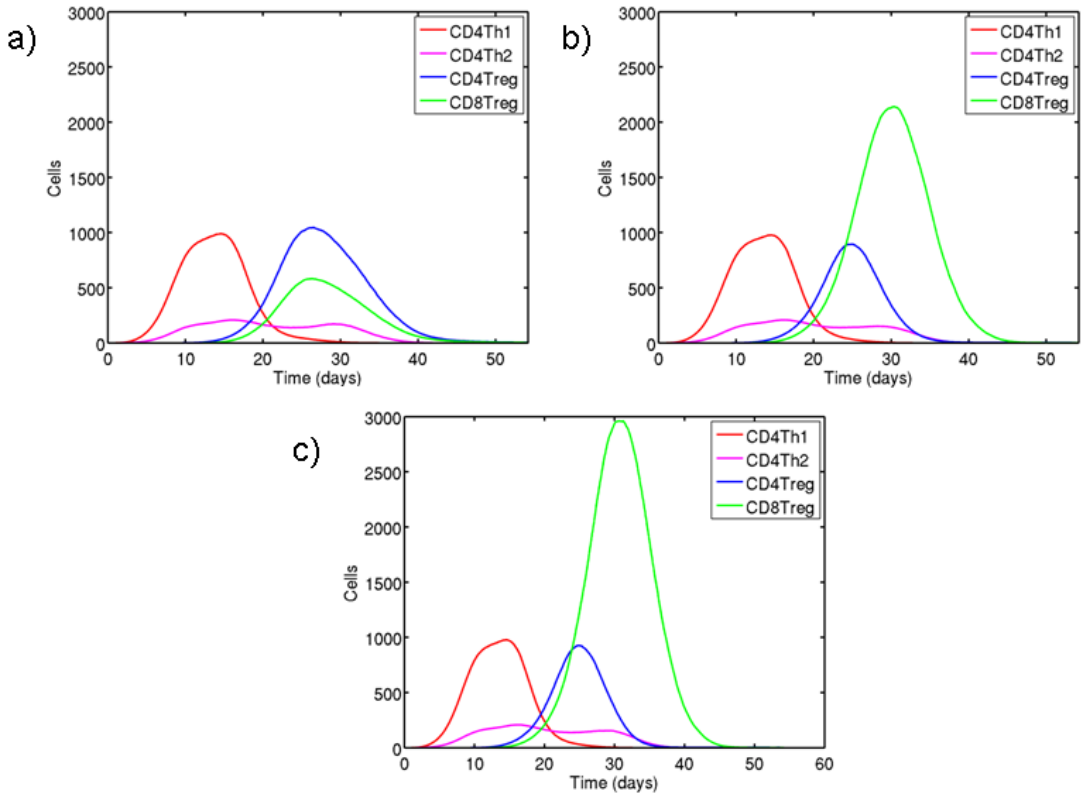

Table 3.4: Summary of the peak population of CD8 Treg effectors attained in the baseline, CD4 Treg abrogation and spatial restriction experiments. The non-linear relationship between the number of T-cells permitted per grid space and the maximum CD8 Treg effector population arises from system stochasticity.

Experiment	**Peak CD8 Treg effector population (cells)**	**Number and nature of T-cells allowed to enter a grid-space**
baseline	600	CD4 Th / CD4 Treg / CD8 Treg enter if < 7 T-cells already present
CD4 Treg abrogation	1000	CD4 Th / CD4 Treg / CD8 Treg enter if < 7 T-cells already present
Maximal crowding 14	2300	CD4 Th / CD4 Treg enter if < 7 T-cells already present, CD8 Treg enter if < 14 T-cells already present
Maximal crowding 21	3000	CD4 Th / CD4 Treg enter if < 7 T-cells already present, CD8 Treg enter if < 21 T-cells already present

An investigation of the cellular neighbourhood of apoptotic DCs in these experiments showed that the DC experienced an increase in spatial occupancy by CD8 Treg, with some showing more than the theoretical maximum occupancy as defined by the baseline experiment i.e. more than 56 T-cells around a DC. The trend toward super-crowded DC was more pronounced towards the end of the simulation for CD8 Treg. CD4 Treg showed the reverse trend, as CD8 Treg numbers expand, the CD4 Treg are apparently 'crowded out' from the DC and the amount of space occupied by CD4 Treg diminishes toward the end of the simulation. This is illustrated in Figure 3.15a-d below.

3.9 The Significance of the Changes in Maximum CD8 Treg Populations Between Experiments

We are now in a position to try to assess the relative influence of the age of licensing and spatial saturation hypotheses on the peak populations of CD8 Tregs in the system. We aim to achieve this using non-parametric A-Test statistics [Vargha and Delaney 2000] as described in Section 3.4.3.

The T-cell population (or 'response') curves from each experiment were analysed to locate the maximum response of each T-cell subtype and the time within the simulation at which this occurred. This data was then compared to that from the baseline and original CD4 Treg abrogation experiments. The comparisons of the peak CD8 Treg populations in the experiments from the 'age of licensing' hypothesis are

presented in Table 3.5 and those for the spatial competition experiments in Table 3.6. The corresponding measures of significance for the changes relative to baseline and CD4 Treg abrogation experiments are presented in Tables 3.7 and 3.8 respectively.

Figure 3.15: Illustrating spatial saturation of DC under the modified CD8 migration rules – in this instance we allowed CD8 Treg to continue to migrate into a grid space up until there were 14 T-cells present (the default simulator behaviour is to permit a maximum of 7 T-cells per grid point). The data are normalised by the 'early' and 'late' DC populations to facilitate comparison. Panels a) and b) show the fraction of 'early' and 'late' Th1 peptide-expressing population DCs having a given proportion of their binding space occupied by CD8 Treg. Panels c) and d) show the fraction of 'early' and 'late' population Th1 peptide-expressing DCs having a given proportion of their binding space occupied by CD4 Treg.

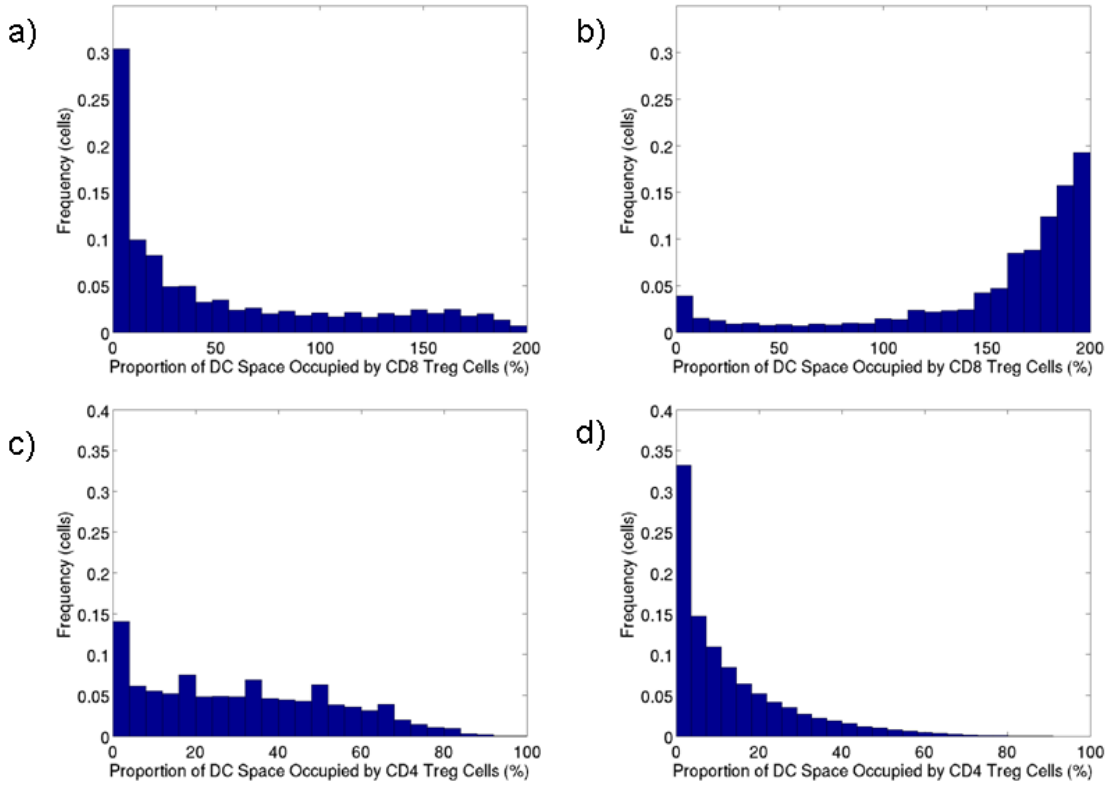

Table 3.5: Comparison of the peak CD8 Treg effector populations in the different experiments performed on the licensing of DC by CD4 Treg.

Experiment	Peak CD8 Treg population (cells)	delay in Qa-1 expression /hours
CD4 Treg abrogation	1000	0.000
'Qa-1 Delay'	750	59.500
Baseline	600	73.375
'Longer Qa-1 Delay'	400	82.500

Table 3.6: Comparison of the effects of the different experiments performed on the spatial competition between CD4 Treg and CD8 Treg around DC.

Experiment	Peak CD8 Treg population (cells)	CD8 Tregs can enter grid point if <X T-cells present: X= (cells)
Baseline	600	7
CD4 Treg abrogation	1000	7
'Maximal crowding 14'	2300	14
'Maximal crowding 21'	3000	21

Table 3.7: Comparison of the effects of different experiments performed on the licensing of DC by CD4 Treg on peak population attained by effector CD8 Treg.

Experiment 1	Experiment 2	A-Test measure	Effect Size
CD4 Treg abrogation	baseline	0.941468	Large
'Qa-1 delay'	baseline	0.704443	Medium
'Longer Qa-1 delay'	baseline	0.163168	Large

Experiment 1	Experiment 2	A-Test measure	Effect Size
'Qa-1 delay'	CD4 Treg Abrogation	0.176274	Large
'Longer Qa-1 delay'	CD4 Treg Abrogation	0.011294	Large

The A-Test statistics suggest that delaying Qa-1 expression by DC during a simulation in which CD4 Treg have been abrogated, has a significant effect on the

peak population attained by the CD8 Treg compared to CD4 Treg abrogation experiments in which Qa-1 expression is constitutive and immediate. The effect sizes of the changes relative to the baseline experiment are less marked as the mean delay times in these experiments are not as far removed from that found in the baseline i.e. 3.06 days as they are from the delay i.e. 0.00 days in the CD4 Treg abrogation experiment.

Conversely, the effect sizes of the peak CD8 Treg population changes caused by the changes of migration rules for CD8 Tregs ('maximal crowding 14' and 'maximal crowding 21' experiments) were both large compared to both the baseline and the CD4 Treg abrogation experiments as shown in Table 3.8.

Table 3.8: Comparison of the effects on the peak population of CD8 Treg effectors of different experiments performed on the spatial competition between CD4 Treg and CD8 Treg around DC.

Experiment 1	**Experiment 2**	**A-Test Score**	**Effect Size**
Maximal crowding 14	baseline	0.999985	Large
Maximal crowding 21	baseline	1.000000	Large

Experiment 1	**Experiment 2**	**A-Test Score**	**Effect Size**
Maximal crowding 14	CD4 Treg Abrogation	0.994426	Large
Maximal crowding 21	CD4 Treg Abrogation	0.999787	Large

From Table 3.7 it can be seen that the timing of licensing is profoundly important in determining the size of the population of CD8 Treg during immune response (CD4 Treg abrogation with Qa-1 delay implemented versus CD4 Treg abrogation with no delay in Qa-1 expression). Equally, the effect of reducing the spatial restrictions applied to CD8 Treg migrating into the vicinity of a DC, are seen to have a large effect on the peak population of CD8 Treg compared to the baseline and CD4 Treg abrogation experiments.

Therefore we conclude that the lifting of spatial competition in the original CD4 Treg abrogation experiment was at least as important as the constitutive expression of Qa-1 by DC in promoting a much inflated population of CD8 Treg since we cannot realistically say any more based on the A-Test results derived in the course of this work.

3.10 Verification of the Significance of the Effect of Spatial Competition on Peak CD8 Treg Effector Population – an Experimental Postscript

In Section 3.8 we presented results that demonstrated the significance of the impact that CD4 Treg – CD8 Treg spatial competition has on the peak population of CD8 Treg Effectors in the simulation. During the writing process and following discussion of the experimental design, it was decided that the experiments of Section 3.8 were not the fairest comparison of the baseline and CD4 Treg abrogation experiments possible.

To gain as fair a comparison as possible of these two experiments in order to assess the effect of Treg spatial competition, we need to be able to allow the same number of CD8 Tregs into a baseline simulation grid space as can occupy an abrogation experiment grid space whilst keeping all other factors unchanged.

From Table 3.2 which details the median spatial occupancy around DC at their time of apoptosis, we see that the median fraction of DC binding space occupied by CD8 Treg is ~84%, corresponding to 6 CD8 Treg around each DC. In comparison, the median fraction of DC binding space occupied by CD8 Treg in the baseline is ~30%.

A further experiment was conducted in which we again modified the rules by which we assess whether T-cell migration into a given grid point are allowed. In this particular experiment we continued to test CD4 Treg and CD4 Th migration as in the default simulation i.e. a CD4 Treg or CD4 Th can only enter a grid point if there are fewer than 7 T-cells already present there. The change was implemented in how we tested for CD8 Treg migration. When attempting to move a CD8 Treg into a grid point we only counted the CD8 Treg that were already present at that point – and if there were fewer than 7, the CD8 Treg could migrate to the grid point.

Once again, the change in the way we allow CD8 Treg to migrate permitted the CD8 Treg to overcome the spatial competition from CD4 Treg and greatly enhanced the peak population of CD8 Treg Effectors. The baseline population of ~600 cells was increased to ~1400 when CD8 Treg migration rules were relaxed. The non-parametric effect size A-Test scored the effect of this experiment as large when comparing against either the baseline or the CD4 Treg abrogation experiment (0.762 for peak CD8 Treg population and 0.001 for peak CD4 Treg population when comparing to the baseline simulator. The corresponding scores for comparison to the CD4 Treg abrogation experiment are 0.074 for CD8 Treg effector population and 0.000 for the CD4 Treg population).

This experiment serves to corroborate the experimentation described in Section 3.8 and further strengthens the idea that CD4 Treg – CD8 Treg spatial competition exerts a powerful effect on the eventual population size of CD8 Treg during an immune response.

3.11 Conclusions

The distribution of median ages of DCs at time of licensing for Qa-1 suggest that the age of DCs at licensing is quite diverse but tends to lie between 3 and 4 days (70-80 hours) as calculated from the entire simulation population.

If DC data is split into two populations, those DC created before and after 20 days, we see a clear distinction of the medians of the two data sets. Intuitively this makes sense as early on in the simulation there are fewer CD4 Tregs in the simulation to license the DCs for Qa-1 expression and we observe a longer median age of 3.48 days (82.50 hours). For the DCs created later in the simulation we observe a median age of approximately 2.48 days (59.50 hours). At this stage in the simulation there are a greater number of CD4 Treg and so a greater chance of encounter between them and DC effectors so making the possibility of productive encounters greater.

The distribution of median spatial saturations shows that for DC expressing Th1-derived peptide, binding space is dominated by CD4 Treg in a ratio 2:1 reminiscent of the ratio found for CD4 to CD8 T-cells produced in Efroni *et al.*'s model of thymocyte development in the thymus [Efroni *et al.* 2007]. The proportion of binding space occupied is roughly CD4 Treg: 60% and CD8 Treg: 30%.

DCs expressing MBP are saturated by CD4 Th cells with approximately 90% of available binding space occupied by these cells in the baseline runs. One possible explanation for the fact that this figure is not 100% lies in the stochasticity of the system. When a T-cell attempts to bind to an APC in the simulation, it does not check every single grid space but randomly samples several spaces prior to moving away and it is this less than full sampling of DC binding space that can result in less than 100% occupancy.

Within these simulation runs there were very few DC expressing MBP and Th1-derived peptides together and it is difficult to draw meaningful generalizations from the data. There were no CD8 Tregs observed around these cells and very few CD4 Tregs. Several simulation runs exhibited DCs with high proportions of occupied binding space – up to ~75%, the neighbouring cells being almost exclusively CD4 Th1.

Generally spatial saturation of DCs expressing Th1-derived peptides tends to increase as the simulation progresses, whereas the spatial saturation around MBP expressing DCs falls slightly.

Abrogation of the CD4 Treg population within the simulator appears to permit the spatial saturation of Th1-derived peptide expressing DC by CD8 Treg to increase to levels significantly above those observed for the baseline runs.

The further experiments, in which we relaxed the rules governing CD8 Treg migration into a grid space, showed that permitting CD8 Treg to benefit from a more lenient migration condition had a very significant impact on the peak response for

these cells. It is entirely possible that there may be some limit on just how many T-cells can be allowed to occupy any one binding site around a DC. Eventually if CD8 Treg proliferate to the extent that they then completely prevent CD4 Treg from accessing DC and then licensing them, no more CD8 Treg become activated or proliferative and the population plateaus. It remains to be determined just where this population plateau lies in relation to the baseline restrictions on T-cell migration. The significance of the effect of relaxing spatial competition was further demonstrated by an additional experiment in which we counted only CD8 Treg in the destination grid cell if we were attempting to move a CD8 Treg into that grid cell.

Overall, we observe that the simulator is producing an elevated peak population of CD8 Tregs under CD4 Treg abrogation via a combination of Hypothesis 1 and Hypothesis 2. That is to say that both reduced spatial restriction due to the missing competition from CD4 Treg and the fact that constitutive Qa-1 expression was permitted on DCs have allowed the CD8 Treg to become activated and proliferate more freely than they would have done in the baseline simulation.

One could be tempted to speculate that the spatial competition between CD4 Treg and CD8 Treg for binding space around DC coupled with the requirement for DC to be licensed by CD4 Treg before CD8 Treg can be activated, serves as a brake on potentially runaway responses by CD8 Treg.

An attempt has been made to relate the observations drawn from the presented experimentation to the literature. This has proved difficult since the work conducted has been speculative and would have been problematic to conduct in a wet lab. However, the observation that CD4 Treg outnumber CD8 Treg around DC by 2:1 mirrors the result of CD4 T cells outnumbering CD8 T-cells 2:1 obtained by Efroni *et al.* 2007 in their simulation of T-cell development in the thymus. To date we have been unable to obtain further corroboration of our results from the literature.

Chapter 4: Investigating the Effect of Adding the CD200-CD200R Regulatory Pathway to the Domain Model

4.1 Introduction

We are principally interested in aspects of CD8 Treg involvement in EAE regulation. The expression of CD200 by CD8 Treg effectors potentially allows these cells to signal down-regulation of expression of costimulatory molecules, Qa-1 and MHC-II molecules by DCs. This in turn impacts the ability of the DC to prime all T-cell sub-populations, so not only regulating the disease but also CD8 Treg mediated killing of auto-reactive CD4 T-helper cells.

We seek to implement a limited model of the CD200-CD200R regulatory axis within our existing domain model to permit us to address greater model complexity within the context of the simulation and to investigate the effect that this regulation axis has on CD8 Treg behaviour and numbers within our EAE disease model.

We have chosen to implement a simplified model of the axis to permit us to understand simulation behaviour in a less complex model prior to adding layers of greater complexity at a later stage. This greater complexity consists of extending the effects of CD200 and CD200R expression into the CNS.

We present the revisions that we have made to the domain model in Section 4.3 and discuss this model and its implications in Section 4.4. The augmented model introduces two new parameters – the probability that a CD200-CD200R interaction (‘negative signal’) causes a DC to down-regulate costimulatory molecule ('CoStim') expression and the probability that a negative signal causes a DC to down regulate Qa-1 and MHC-II molecule ('MHC') expression.

The work presented here attempts to locate values for these two parameters that return baseline simulation behaviour when this additional regulatory axis is implemented. We aim to achieve this via systematic mapping of simulation behaviour at an array of combinations of the parameter values and assessing which lie close to baseline behaviour (a 'factorial analysis'). Ultimately, our aim is to rebalance the CD200-CD200R regulatory mechanism with the CD8 Treg mediated cell killing mechanism so as to restore baseline simulation behaviour.

The preliminary factorial analysis is described in Section 4.5 and a brief investigation of the immunological effects of varying the probability parameters is presented in Section 4.6. A further factorial analysis at smaller values of the two parameters was required, and is described in Section 4.7. In Section 4.8 we present a brief experiment as a preliminary attempt to assess the extent of the interaction between the additional disease regulation axis and the CD8 Treg mediated killing mechanism that existed in the baseline simulation. Our conclusions about the outcomes and efficacy of this work are presented in Section 4.9.

4.2 The Domain

4.2.1 CD200 and CD200R

CD200, also called OX2 [Minas and Liversidge 2006, Feuer 2007], is a small glycoprotein [Hoek *et al.* 2000, Wright *et al.* 2003, Meuth *et al.* 2008] expressed on the surface of a wide variety of cell types [Hoek *et al.* 2000, Wright *et al.* 2003]. Different authors report CD200 expression on different cell types but there appears to be consensus that CD200 is expressed by neurons [Hoek *et al.* 2000, Minas and Liversidge 2006, Copland *et al.* 2007, Meuth *et al.* 2008, Liu *et al.* 2010] and by T-Cells [Minas and Liversidge 2006, Copland *et al.* 2007, Liu *et al.* 2010]. Other authors report expression by B-cells [Hoek *et al.* 2000, Minas and Liversidge 2006] and DCs [Liu *et al.* 2010].

CD200 interacts with the CD200R family of receptors [Gorczynski *et al*. 2004a] which is also a cell surface expressed glycoprotein. CD200R is generally expressed in cells of myeloid origin [Wright *et al.* 2003, Gorczynski *et al.* 2004a, Meuth *et al.* 2008] such as DCs [Copland *et al.* 2007, Liu *et al.* 2010], macrophages and microglia [Liu *et al.* 2010]. Minas and Liversidge 2006 report CD200R expression on APCs generally.

The interaction of CD200 with its receptor has important consequences for the behaviour of the cell receiving the receptor signal. CD200 has been shown to mediate inhibitory signals via CD200R [Meuth *et al.* 2008] which play a role in the regulation of macrophages [Hoek *et al.* 2000, Wright *et al.* 2003, Minas and Liversidge 2006]. Through this signalling, CD200 helps to regulate the adaptive immune system through lymphocyte-APC interactions [Gorczynski *et al.* 2004a].

The CD200-CD200R axis is potentially a very complex regulatory network as there is some evidence that CD200R is expressed on certain T-cell subsets whilst, conversely APCs may also express CD200 [Minas and Liversidge 2006, Liu *et al.* 2010]. This additional complexity may serve as a form of 'fine control' of immune response [Minas and Liversidge 2006].

The evidence for, and the implications of this role for CD200 are examined in the following Section (4.2.2).

4.2.2 Experimental Evidence for the Role of the CD200-CD200R Regulatory Axis

A number of authors report the increased activation of macrophages and severity of autoimmune disorders in mice lacking the ability to express CD200 ($CD200^{-/-}$ mice) [Hoek *et al.* 2000, Copland *et al.* 2005, Melchior *et al*. 2006]. Hoek *et al.* 2000 also demonstrated that disruption of the CD200-CD200R interaction precipitates susceptibility to Collagen-Induced Arthritis (CIA). Copland *et al.* 2005 found greater infiltration of the retina by macrophages during Experimental Autoimmune Uveoretinitis (EAU). However, the authors also found that administration of a

CD200R agonist was capable of suppressing EAU and produced an earlier resolution of the disease despite maintenance of T-cell proliferation. Subsequently, Melchior *et al.* 2006 found that microglia in the CNS of $CD200^{-/-}$ mice showed an activated phenotype even in the absence of pathogenic stimuli.

In contrast, a number of authors have demonstrated that stimulation of CD200R can lead to suppressed macrophage activity and reduced severity of inflammation [Lue *et al.* 2010] or autoimmune disease [Feuer 2007, Liu *et al.* 2010]. Liu *et al.* 2010 found that the decreased severity of EAE in their study was linked to reduced axonal damage and demyelination, a fact that they attributed to suppression of macrophage and microglial accumulation in the CNS.

On the other hand, blockade of CD200R such that CD200 cannot bind leads to mice that exhibit an aggravated form of EAE [Meuth *et al.* 2008]. This study also observed an enhanced infiltrate of T-cells and macrophages in the spinal cord lesions of the affected mice, an observation that is explained in terms of CD200-CD200R interaction permitting a form of spatially restricted signalling between neurons and infiltrating macrophages [Meuth *et al.* 2008].

Finally, a study of Wld^{s}[21] mice showed that they suffer an attenuated form of EAE because these mice possess a gene phenotype that confers axon protection [Chitnis *et al.* 2007]. The authors also found decreased macrophage accumulation in the CNS which they associated with the constitutive expression of CD200 on neurons. However, administration of anti-CD200 antibody to Wld^{s} mice abrogated the axonal protection that they benefitted from and the mice then suffered increased CNS inflammation and neuro-degeneration. The authors concluded that CD200-CD200R plays a critical role in the attenuation of EAE and in reducing inflammation-mediated damage in the CNS [Chitnis *et al.* 2007].

Similarly, the myxoma virus can express a cell surface protein that possesses significant similarity to CD200 and which is implicated in myeloid lineage cell activation. The M141R protein is found to be required for full development of lethal myxoma virus infection. Compared to wild type M141R infected rabbits, M141R 'knock-out' virus ($M141R^{-/-}$) infected rabbits showed higher activation levels of macrophages and lymphocytes suggesting that M141R can inhibit macrophage activation and ability to prime lymphocytes, allowing the virus to bypass the host's immune system [Cameron *et al.* 2007].

Taken together, the evidence suggests that the CD200-CD200R axis serves, at a very minimum, as a form of immune response regulation, acting via modulation of DC or macrophage ability to activate T-cells.

[21] A specific genetic strain of mouse used in genetic and immunological experiments.

4.3 Implications in Terms of the CoSMoS Process

4.3.1 The Existing Domain Model

The purpose of a domain model is to fully describe the domain expert's current understanding of the domain of interest so that a model may be implemented from it. The EAE simulator domain model was briefly introduced in Section 2.10. The model comprises 25 UML diagrams which between them fully describe the baseline simulator behaviour [Read *et al.* 2009a, Read 2011].

The domain model is fully described in [Read *et al.* 2009a, Read 2011]. To give the reader some appreciation of what has been changed in the simulator in order to carry out the work presented in this chapter, a short description of the domain model diagrams is given below.

The largest set of diagrams in the model represent the lowest level of abstraction i.e. the most detailed descriptions, and are the state machine diagrams for the various cell types and molecules in the simulator. These outline what states a particular kind of cell may exist in and how the cell's behaviour changes according to what state it is in e.g. DCs can be in a non-mature or a mature state.

Similarly there are state diagrams for the other cell types of the simulator [Read 2011] – CD4Treg, CD4THelper, CD8Treg, neuron, microglia (CNSmacrophage) and a special diagram for generic T cell behaviour within Lymphoid Organs. The model also details the possible states and behaviours of the important chemical entities in the system – MBP, TNF-α and Types 1 and 2 cytokines.

The domain model [Read 2011] contains several UML class diagrams whose purpose is to provide descriptions of how the various classes interact to produce certain systemic behaviours such as propagation and regulation of EAE. The diagrams detail the pattern of inheritance among classes, and in what numbers the classes interact with each other ('aggregation') and in what way they interact ('association'). The Simulator Domain Model includes four class diagrams which describe how the various cell types interact to produce the EAE disease cycle and the regulation mediated recovery cycle (both of which are informally depicted in Figure 2.3 in Chapter 2 of this thesis).There is also a diagram describing how cellular interactions give rise to the 'type 2 deviation' which is the tipping of the CD4 Th1 to Th2 ratio in favour of Th2 as recovery progresses and Th1 cells are killed by CD8 Treg.

The model [Read 2011] also contains four activity diagrams that correspond roughly to the class diagrams described above. Activity diagrams detail complex processes by documenting sequences of actions performed by the various system agents. The Domain Model includes diagrams illustrating how cellular level events give rise to instigation and perpetuation of disease as well as to recovery and type 2 deviation.

Finally there are diagrams which illustrate the conventions used in describing cellular relationships such as interrupted interactions and spawning behaviour. There are also

diagrams that depict the system at higher levels of abstraction – informally illustrating the disease-recovery cycle and the expected behaviours of the system which lead to the directly observable results.

4.3.2 A Revised Domain Model

We are introducing two new chemical entities, CD200 and CD200R which certain cells are allowed to express under specific conditions. The two proteins will be allowed to interact, with the interaction producing specific results in the interacting cells.

For this reason we need to modify the existing domain model diagrams [Read 2011] to incorporate the expression and effects of the new protein entities. The relevant diagrams contain some elements that are not part of standard UML. Read provides a full description of the domain model and all of the symbols employed within it [Read 2011]. The symbols coined by Read that appear in the diagrams relevant to the changes effected here are described in Table 4.1 along with those notations added by the author.

Table 4.1: The non-standard UML notations employed in the UML domain model diagrams presented in the Sections 4.3.2.1 to 4.3.2.5. These notations have been devised in an attempt to express concepts that are not readily conveyed in standard UML.

Symbol	Explanation
λ	on a state transition indicates that a transition between states occurs after some discrete period of time over which a specified event occurs – a temporal guard.
δ	on a state transition forms a probabilistic term in a guard
pair of parallel dashed lines	indicates that an event or transition is probabilistically interrupted (by analogy with Read's interrupt relationship [Read 2011]).
dashed arrows	indicate event transitions that occur probabilistically (by analogy to the probabilistic interrupt described above).

Modifications have been made to six of the domain model diagrams that are relevant to DCs, CD8 Treg and their roles in our model of the CD200-CD200R regulatory axis. We present the relevant portions of the model and outline the changes made to the existing EAE model in Sections 4.3.2.1 to 4.3.2.5.

4.3.2.1 The State Machine Diagram for CD8Treg

We restrict CD200 expression to CD8 Treg Effector cells. CD200 expression is constitutive and immediate upon a CD8 Treg becoming an effector, local activation of the CD8 Treg not being required prior to CD200 expression. We have maintained the current assumption that CD8 Treg cannot enter the CNS, despite evidence to the contrary [Zozulya *et al.* 2009], in order to maintain the simplicity of the model so that we can effectively explore, parameterise and understand the system more fully before incorporating further levels of complexity. By allowing CD8 Treg to migrate into the CNS we would also then have to allow microglia to express CD200R and neurons to express CD200. Before we reach this level of complexity we need to be able to reliably assess the impact of the regulatory axis on the DCs. The state diagram needs to be changed to reflect this altered situation as illustrated in Figure 4.1.

Figure 4.1: Revised state diagram for the CD8Treg class. The modifications made by the author are highlighted in red and show that expression of CD200 commences immediately upon a CD8 Treg maturing into an effector cell and without the prior need for local activation. Otherwise, the state behaviour of the CD8Treg class is substantially unaltered. Adapted from the original CD8Treg state diagram in [Read 2011].

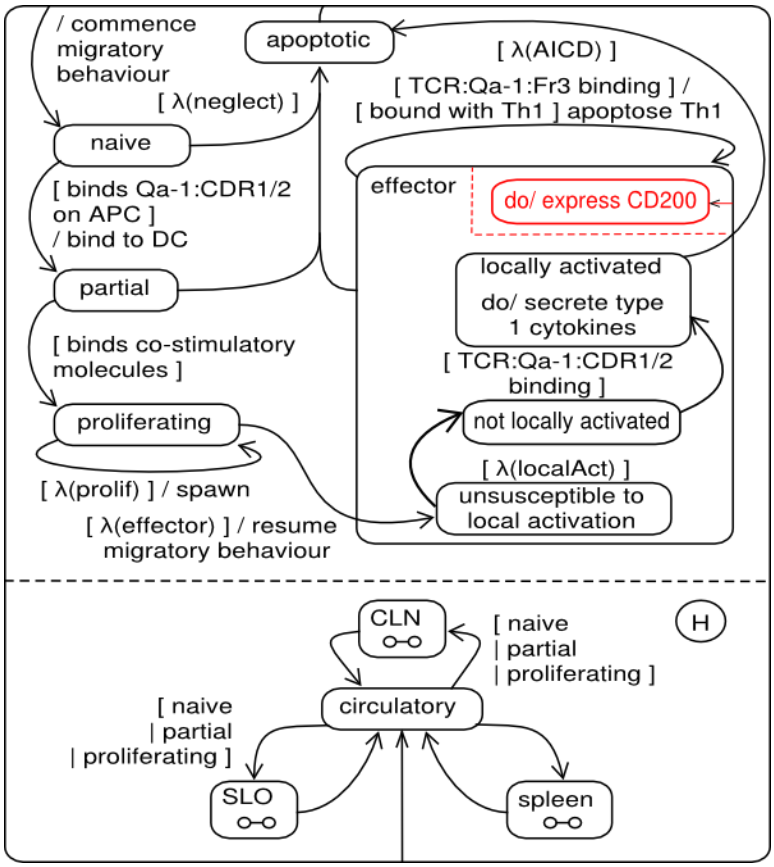

4.3.2.2 The State Machine Diagram for DendriticCell

We enable mature DCs to express CD200R immediately upon maturity. CD200-CD200R interaction is implicit when a CD8 Treg and a DC occupy neighbouring grid spaces and no further signals or activation are required. The 'negative signal' generated by CD200-CD200R interaction is to be allowed two possible probabilistic effects a) to down-regulate expression of MHC compounds ('MHC') i.e. MHC Class II and Qa-1 completely and b) to down-regulate expression of costimulatory molecules ('CoStim') completely. The potential for re-induction of CoStim expression is not explicitly ruled out, nor is that of re-licensing of DC to express Qa-1, though in practice this cannot occur because MHC-II down-regulation is permanent once it occurs. The changes made to the model are illustrated in Figure 4.2.

4.3.2.3 The EAE Instigation, EAE Perpetuation and Type 2 Deviation Activity Diagrams

Three of the four activity diagrams also had to be modified to incorporate changes in behaviour associated with the CD200-CD200R negative signal. These changes reflect the potential for the changed ability of DCs to prime T-cells once they have received the negative signal. In the EAE instigation diagram the possibility of disrupting the priming of CD4 Th in the SLO has been added and in the perpetuation diagram we have added the possibility of disruption of CD4 Th priming in the CLN. Finally, modifications to the type 2 deviation diagram illustrate the disruption to Th and Treg priming owing to negatively signalled DC. The changes are presented in Figures 4.3, 4.4 and 4.5.

Figure 4.2: Revised state diagram for the DendriticCell class. The modifications made by the author are highlighted in red and show that expression of CD200R commences immediately once a DC reaches maturity. The diagram also incorporates the anticipated effects of the CD200-CD200R signal on DC behaviour. For instance we have added the suppression of Qa-1 expression on a probabilistic basis upon receipt of the negative signal. Likewise, MHC-II and costimulatory molecule expression is curtailed. However, it should be noted that while MHC-II down-regulation is permanent in our new model, Qa-1 and CoStim can at least in theory be switched back on if the appropriate conditions apply (though, of course, for Qa-1 this does not occur as CD4 T-cells require MHC-II expression before they can bind to and license DC). Adapted from the original DC state diagram in Read 2011 (in this diagram some non-standard notation has been employed: δ denotes a probabilistic process whereas λ denotes a process that occurs after some undefined delay).

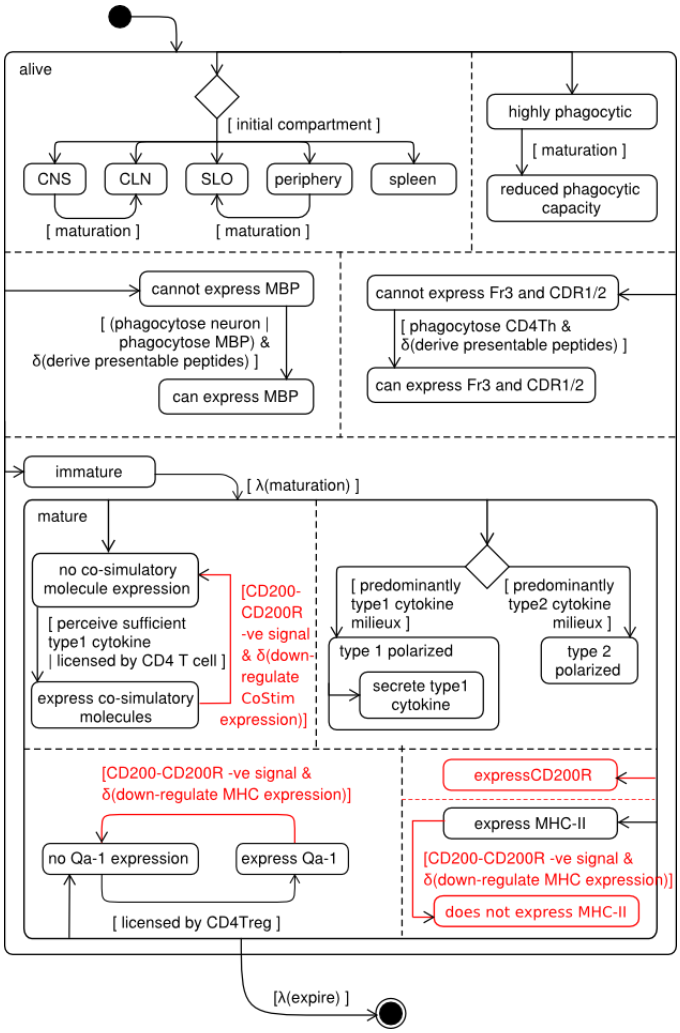

Figure 4.3: Illustrating the modifications made to the EAE instigation activity diagram. The only major change made to this diagram (shown in red) is that the mature DC entering the SLO are now capable of expressing CD200R and thus of being negatively signalled by CD8 Treg effectors expressing CD200. The CD4 Th in the SLO are not able to bind to the signalled DC if these have ceased to express MHC-II compounds and so CD4 Th are not able to receive 'signal 1' necessary for activation. The diagram contains some non-standard UML symbols, the oblique dashed parallel lines are here meant to denote that the cessation of MHC-II expression is probabilistic and therefore the interruption of CD4 Th1 priming is likewise probabilistic. Adapted from the original EAE instigation activity diagram in Read 2011.

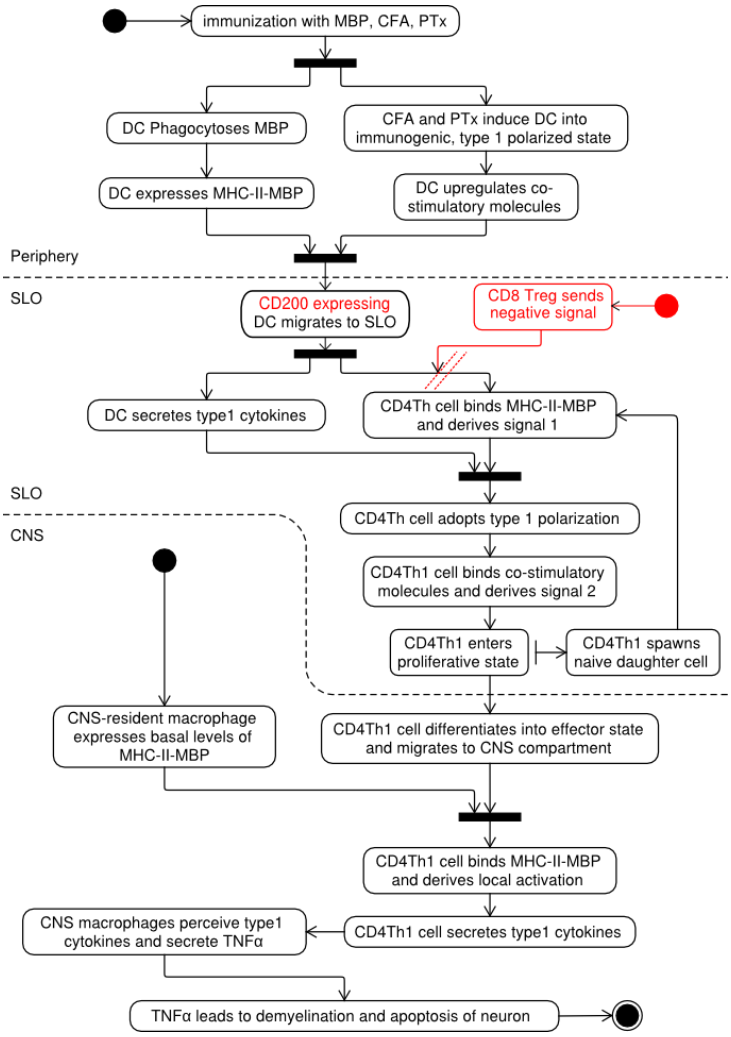

Figure 4.4: Illustrating the modifications made to the EAE perpetuation activity diagram. The only major change made to this diagram (shown in red) is that the mature DC entering the CLN are now capable of expressing CD200R and thus of being negatively signalled by CD8 Treg effectors expressing CD200. The CD4 Th in the CLN arel not able to bind to the signalled DC if these have ceased to express MHC-II compounds and so CD4 Th are not able to receive 'signal 1' necessary for activation. The diagram contains some non-standard UML symbols, the oblique dashed parallel lines are here meant to denote that the cessation of MHC-II expression is probabilistic and therefore the interruption of CD4 Th1 priming is likewise probabilistic. Adapted from the original EAE perpetuation activity diagram in Read 2011.

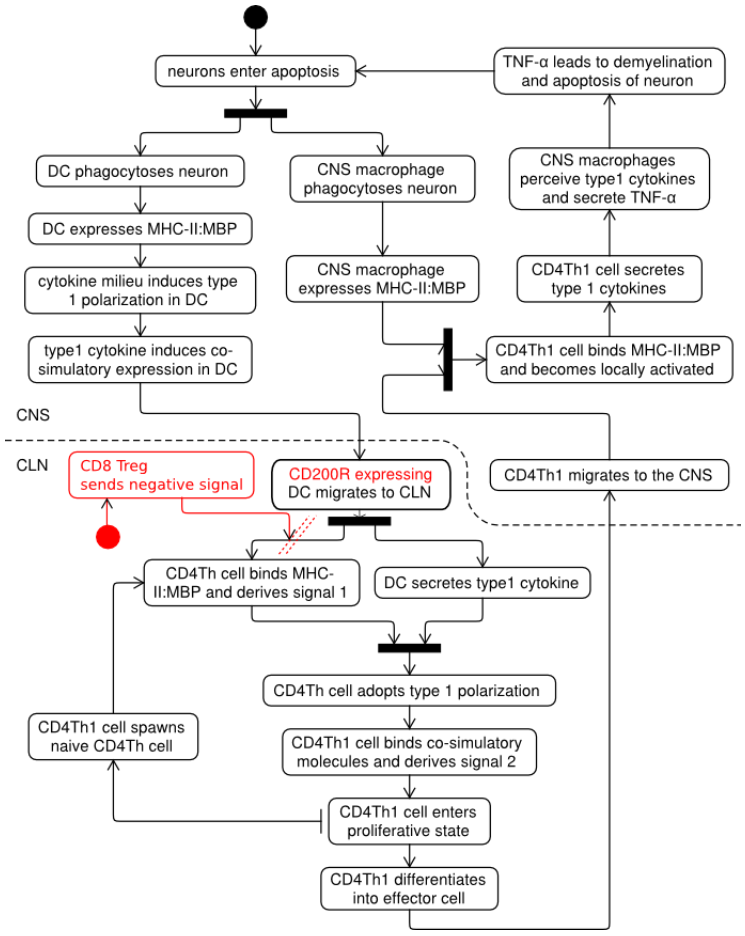

Figure 4.5: Illustrating the modifications made to the type 2 deviation activity diagram. The changes made to this diagram (shown in red) indicate that immunogenic DC are now capable of expressing CD200R and thus of being negatively signalled by CD8 Treg effectors expressing CD200 and that their priming of CD4 Th1 and Th2 populations can be disrupted. Similarly, the possibility of disrupting Treg priming by negatively signalled DC is included here. The diagram contains some non-standard UML symbols, the oblique dashed parallel lines are here meant to denote that the cessation of MHC-II, CoStim and Qa-1 expression is probabilistic and therefore the interruption of T-cell priming is likewise probabilistic. Adapted from the original type 2 deviation activity diagram in Read 2011.

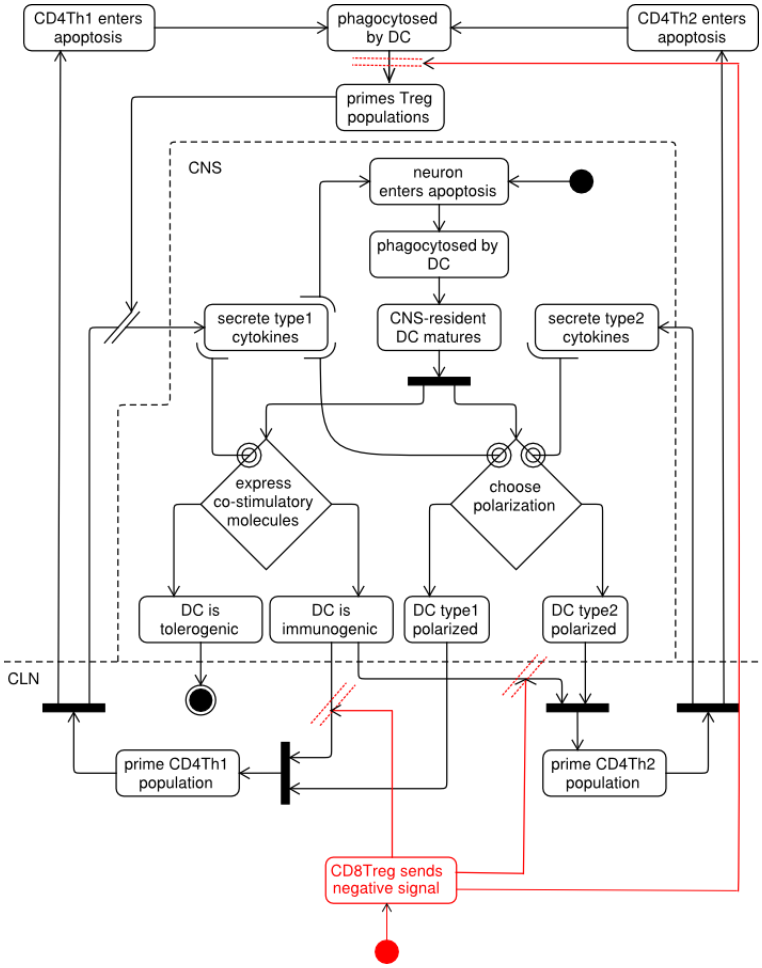

4.3.2.4 The Expected Behaviour Diagram

Finally, we needed to modify the expected behaviour diagram which details the system at the highest level of abstraction. The modified diagram (Figure 4.6) incorporates the potential for down-regulation of the expression of CoStim, MHC-II and Qa-1 on DC by negative signalling from effector CD8 Treg.

Figure 4.6: The Expected Behaviour Diagram which illustrates the macroscale, observable effects of the microscale interactions of the individual cells. The diagram has been modified simply by the addition of an arrow indicating the negative signalling of DCs by CD8 Treg (the change is high-lighted in red). As before the figure is adapted from the original as presented in Read 2011.

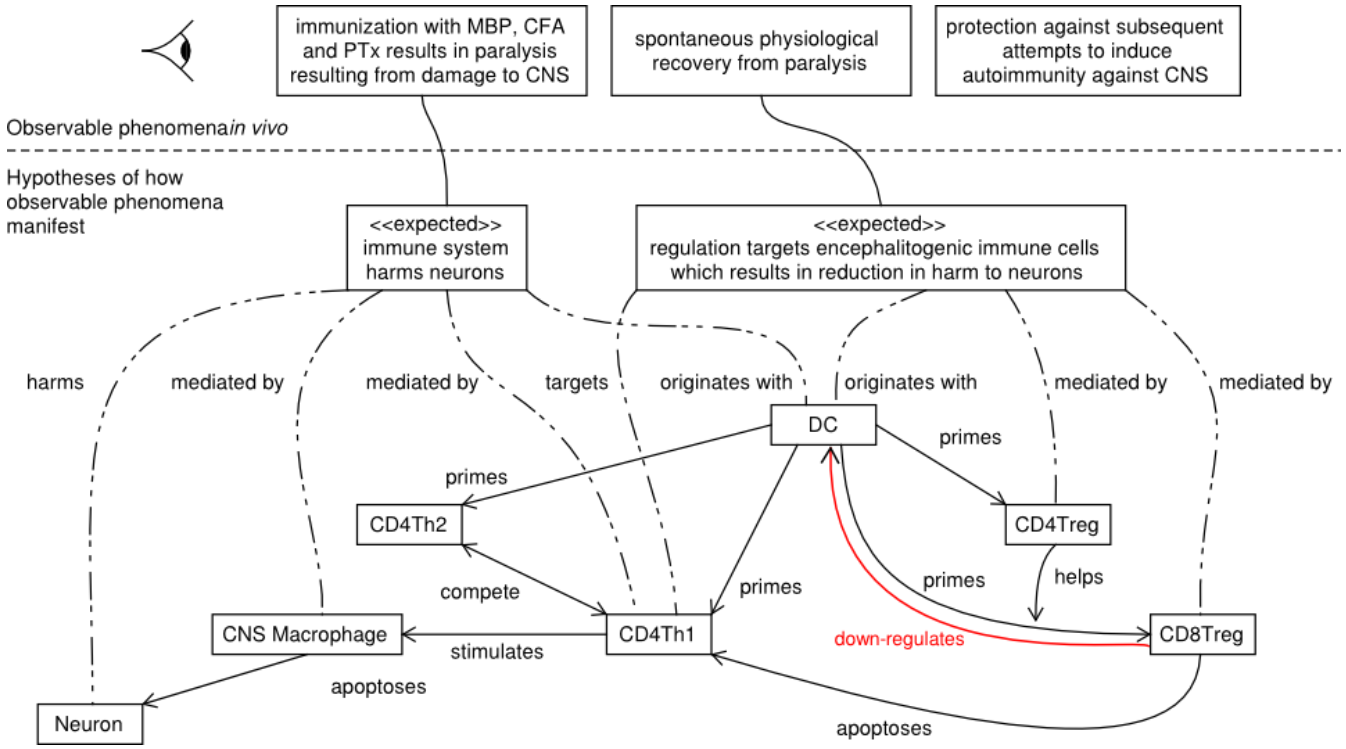

4.3.2.5 Creation of an EAE CD200-CD200R Regulation Axis Activity Diagram

Additionally, we attempted to create an activity diagram that expresses the sequence of events set in motion by negative signalling of DC. A very basic CD200 Axis regulation activity diagram is presented in Figure 4.7.

Figure 4.7: EAE CD200 Axis regulation activity diagram. Non-standard notation has been employed to express ideas that are difficult to express within standard UML. The dotted directed lines are meant to denote that the following action is performed probabilistically – that is down-regulation of Qa-1, MHC-II and CoStim are all probabilistic. No attempt is made in this diagram to connect this regulation axis into the original CD8 Treg mediated regulation. This diagram was not part of the original domain model and has been added by the author.

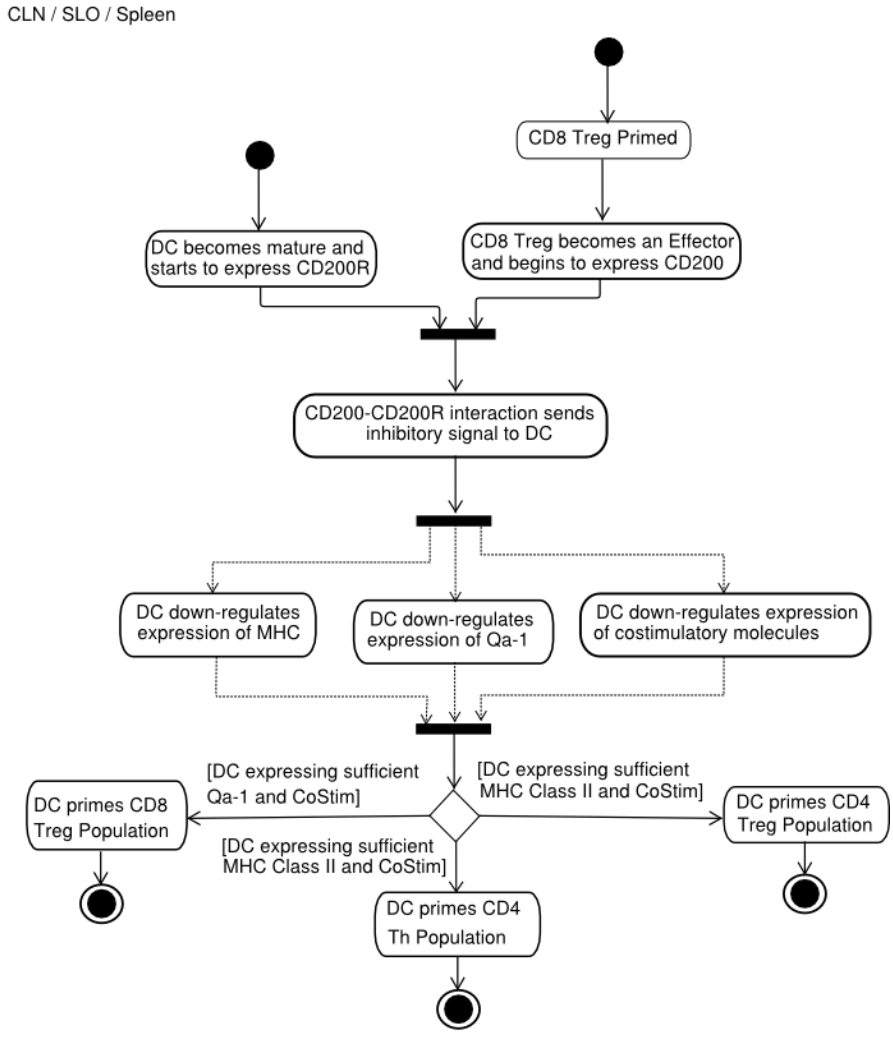

4.4 The Model Implemented

Modifications are made to the EAE Simulator to facilitate incorporation of a simple model of the CD200-CD200R regulatory axis. The altered model contains two new parameters which we wish to assign values to in such a way that simulator baseline behaviour can be reproduced even when the new regulatory axis is included in our implementation.

CD8 Treg effectors are made capable of CD200 expression without local activation, with immediate effect upon becoming an effector. Mature DCs can constitutively express CD200R, again immediately upon maturity.

At the moment CD8 Treg remain excluded from the CNS, neurons are not allowed to express CD200 and microglia are not permitted to express CD200R. This allows us to explore the new model at a lower level of complexity compared to if we had included all the cell types in our model immediately, facilitating our understanding of novel regulatory axis behaviour in the simulation prior to adding further layers of complexity at a later date.

In our model, interaction of CD8 Treg with a DC implies CD200-CD200R interaction and DC receives a 'negative signal' via this interaction. Currently, one negative signal can completely switch off CoStim and / or MHC (MHC-II and Qa-1) expression on a probabilistic basis.

This change potentially radically alters the simulation behaviour and therefore it is necessary to parameterise the probabilities for CoStim and MHC down-regulation upon receipt of the CD200-CD200R negative signal by DC.

There are two probability parameters to explore – the probability that a negatively signalled DC down-regulates MHC and the probability that a negatively signalled DC down-regulates CoStim e.g. one negative signal may have a 10% chance of down-regulating MHC expression and at the same time a 50% chance of down-regulating CoStim expression by DC.

The work presented here attempts to locate values for these two parameters that return baseline simulator behaviour when this additional regulatory axis is implemented in the simulator. We achieve this via systematic mapping of simulation behaviour at an array of combinations of the parameter values and assessing which lie close to baseline behaviour (a 'factorial analysis'). This experiment is described in Section 4.5.

4.5 A Factorial Analysis of the Two Probability Parameters

To gain insight into how we should parameterise such a model, we explore a 2D parameter space over the full range of values for the two probability parameters. We varied each parameter from 0% to 100% in increments of 10% giving 121 simulator experiments with unique combinations of parameter values. We calculated T-cell (Th1, Th2, CD4 Treg and CD8 Treg) populations and EAE Severity scores [Read 2011] (which we refer to as 'responses') for 500 runs of the simulator using each pairing of parameter values. We analysed the maximum population size for each T-cell sub-type along with the time of its occurrence in order to gain insight into the size of the immune reaction and the extent of its regulation. The EAE Severity Scores are an attempt to calculate a severity for the EAE suffered by the simulated mouse at each time-step of the simulation. These scores are meant to mirror the clinical measures used to assess the severity of disease in laboratory animals.

The effect of each pair of parameter settings on peak effector cell population for each T-cell subtype is presented in the figures below. We also present a deeper discussion of the immunological changes occurring alongside the changes in our chosen parameters (probabilities of a negative signal down-regulating MHC or down-regulating CoStim) in Section 4.6.

4.5.1 Effect on Peak CD4 Th1 Effector Population

Figure 4.8: 3-dimensional plot illustrating the peak CD4 Th1 effector population at all of the possible pairings of probability parameter values. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

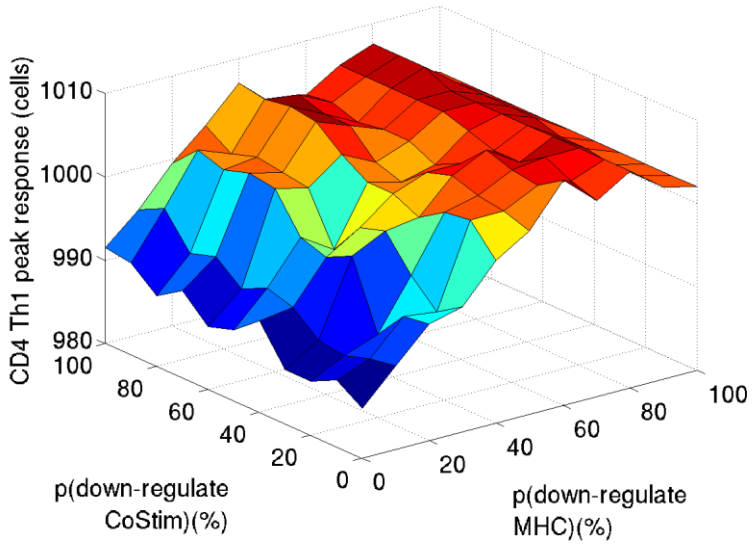

Variation in the value of the CoStim down-regulation probability has relatively little impact on the CD4 Th1 population, whereas gradual increase in the MHC down-regulation parameter leads to a concomitant rise in the peak population of CD4 Th1 observed in the simulation. One possible explanation for this lies in the behaviour of the CD8 Treg population under the influence of the two parameters (see Section 4.5.7). There is a rapid decline in CD8 Treg effector numbers upon increasing the probability of a DC down-regulating MHC expression upon receipt of a CD200-CD200R negative signal.

4.5.2 Effect on Time Taken to Reach Peak CD4 Th1 Effector Population

Unlike with the actual peak population of CD4 Th1, the time taken to reach this peak population within the simulation does not appear to be related to the two probability parameters in a simple manner. The variation of this response is reasonably small (about 1 day in a simulation covering ~50 days) and the A-Test results (presented in Appendix E) comparing runs with varying p(down-regulate CoStim) and varying p(down-regulate MHC) show that variation in either of these probability parameters has no significant effect on the response and therefore the variation is most likely due to the inherent stochasticity of the system.

Figure 4.9: 3-dimensional plot illustrating the time taken to reach peak CD4 Th1 effector population at all of the possible pairings of probability parameter values. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

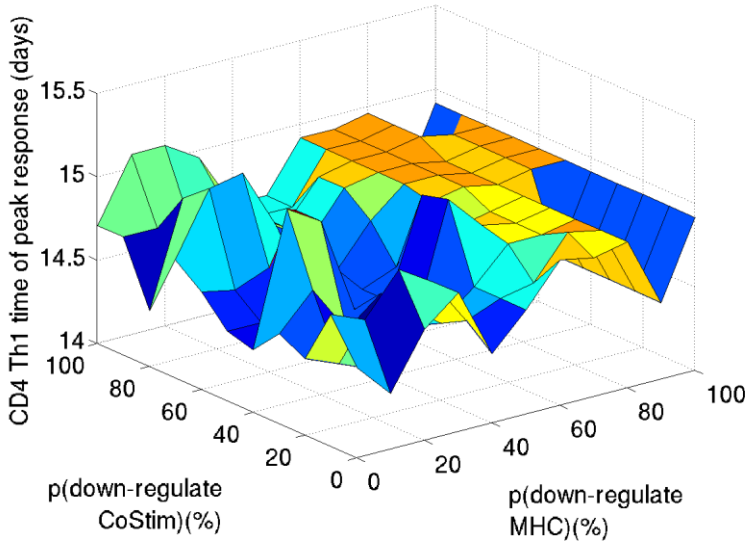

4.5.3 Effect on Peak CD4 Th2 Effector Population

Unlike for the CD4 Th1 population, the peak effector CD4 Th2 population does not appear to bear any straight-forward relationship to the values of the two probability parameters. Here the change in numbers of cells is much smaller (a range of just 3 cells compared to a range of ~20 cells for the CD4 Th1). Again, the relevant A-Test scores (presented in Appendix E.2) suggest that the effect size of varying the two probability parameters is small and so the variation observed is just system stochasticity.

Figure 4.10: 3-dimensional plot illustrating the peak CD4 Th2 effector population at all of the possible pairings of probability parameters. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

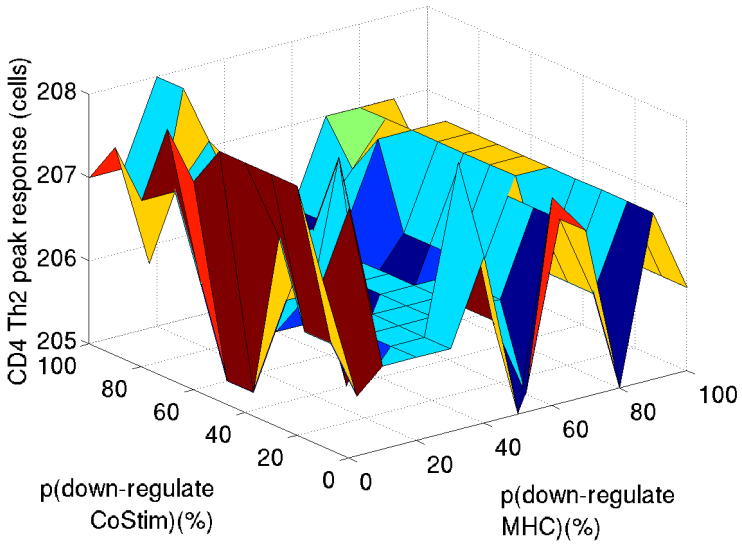

4.5.4 Effect on Time Taken to Reach Peak CD4 Th2 Effector Population

As with the time taken to attain peak CD4 Th1 effector population, the time taken to achieve peak CD4 Th2 effector population in the simulation is not straight-forwardly related to the values of the probability parameters. As with the time of CD4 Th1 peak population variation in the times of peak population here show very little real variation (~1 day). Although the A-Test scores are indicative of small to medium effect due to variation of p(down-regulate MHC) and small effect only for varying p(down-regulate CoStim), the variation observed is probably due to noise.

Figure 4.11: 3-dimensional plot illustrating the time taken to reach peak CD4 Th2 effector population at all of the possible pairings of probability parameter values. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

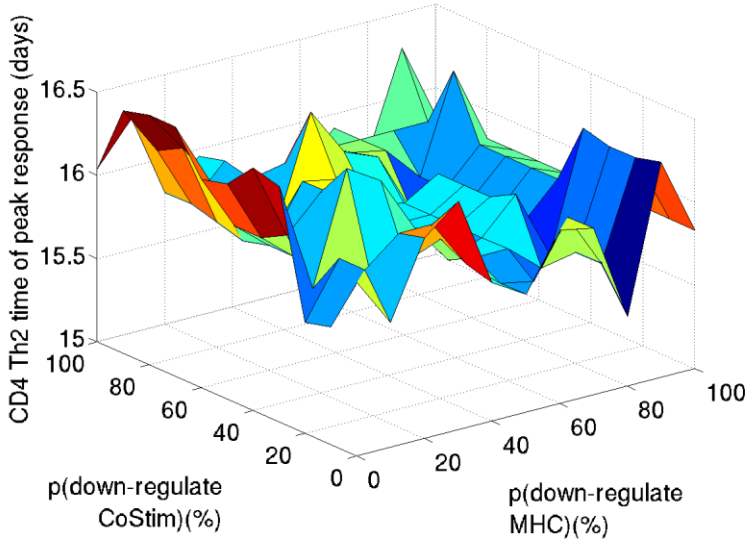

4.5.5 Effect on Peak CD4 Treg Effector Population

The peak CD4 Treg effector population is rapidly diminished once we start allowing MHC expression to be switched off. Compared to the effect of varying the probability that negative signalling switches off MHC expression by DC, the effect of varying the probability of CoStim down-regulation is small, but does exist. The fall in peak CD4 Treg population in going from p(down-regulate CoStim) = 0% to p(down-regulate CoStim) = 10% is from 1053 cells to 1000 cells and stands at 846 cells with p(down-regulate CoStim) at 100%.

Figure 4.12: 3-dimensional plot illustrating the peak CD4 Treg effector population at all of the possible pairings of probability parameters. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

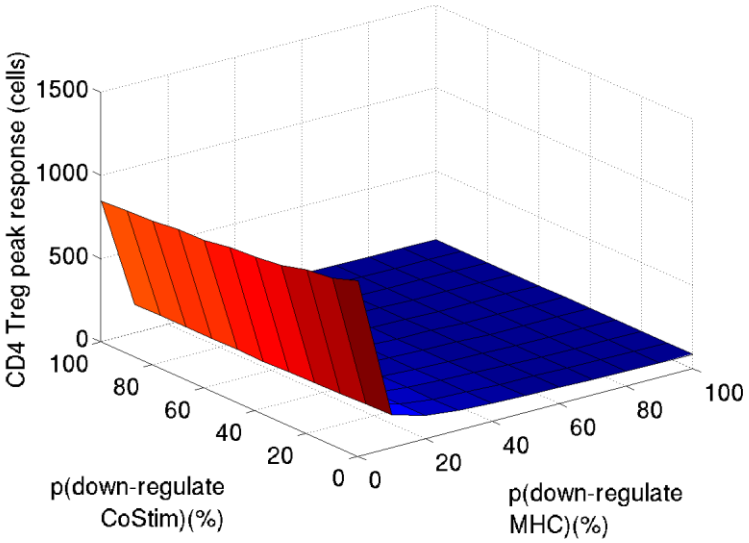

4.5.6 Effect on Time Taken to Reach Peak CD4 Treg Effector Population

As we start to increase the probability of negative signalling down-regulating the expression of MHC, the time taken to reach the peak CD4 Treg population diminishes rapidly. This reflects the much reduced size of the overall population of CD4 Treg in the simulation at these higher parameter values.

Figure 4.13: 3-dimensional plot illustrating the time taken to reach peak CD4 Treg effector population at all of the possible pairings of probability parameter values. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

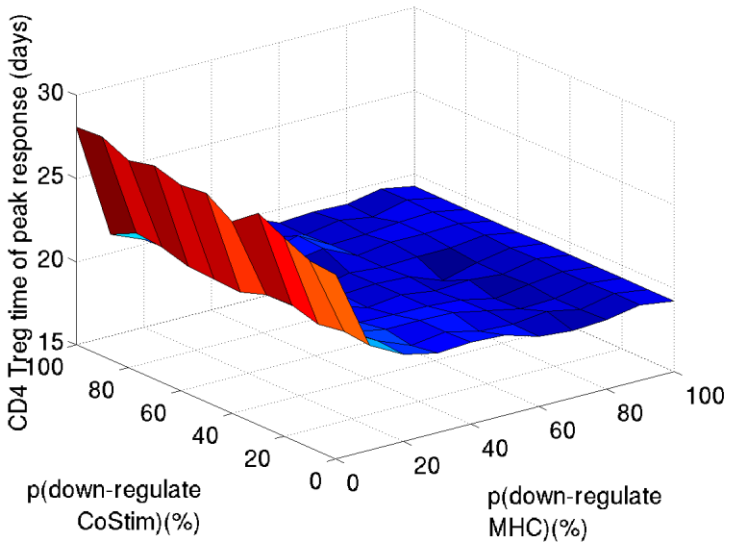

4.5.7 Effect on Peak CD8 Treg Effector Population

As with the CD4 Treg effector population, peak numbers of CD8 Treg effectors are highly sensitive to any fall in the levels of MHC expression. This may well have impacted on the peak population size observed for the effector CD4 Th1 cells in Section 4.5.1.

Figure 4.14: 3-dimensional plot illustrating the peak CD8 Treg effector population at all of the possible pairings of probability parameters. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

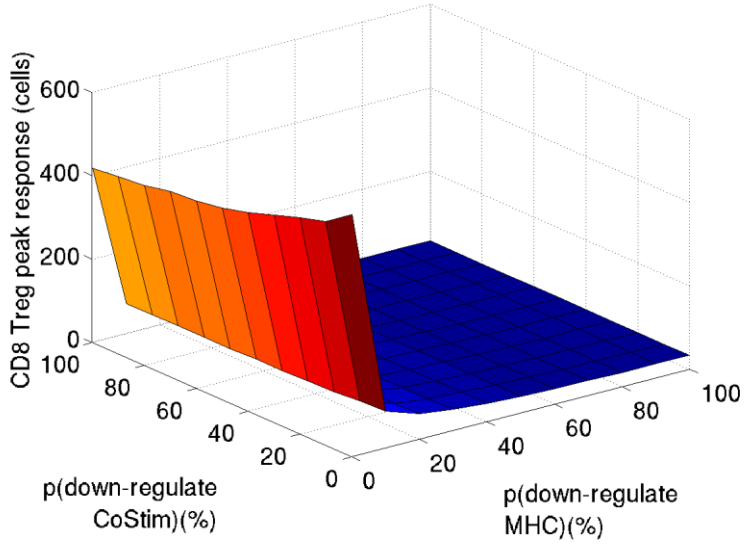

4.5.8 Effect on Time Taken to Reach Peak CD8 Treg Effector Population

As we permit a greater possibility that a DC completely down-regulates MHC expression upon receiving the negative signal, the time taken to reach peak CD8 Treg effector population size falls.

Figure 4.15: 3-dimensional plot illustrating the time taken to reach the peak CD8 Treg effector population at all of the possible pairings of probability parameters. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

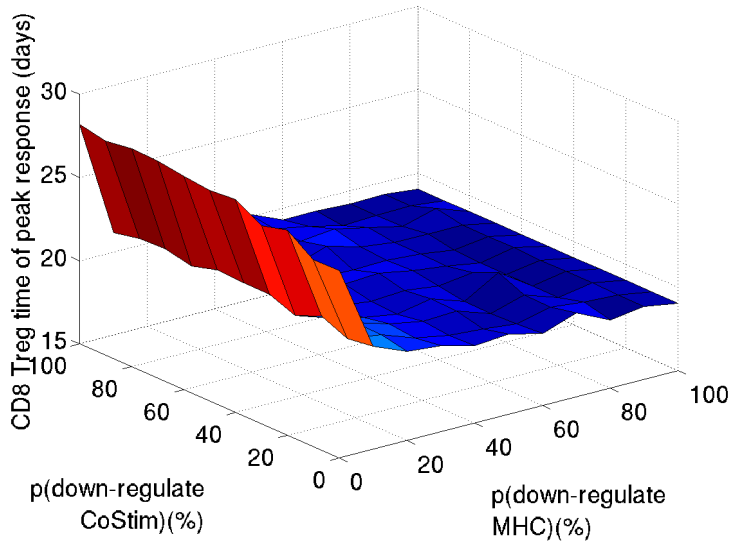

4.5.9 Effect on CD4 Th1 Effector Response at 40 Days

As we allow MHC expression by DC to be curtailed by negative signalling, the size of the CD4 Th1 effector population persisting to day 40 of the simulation starts to rise. As the CD4 Th1 effector population is tending to persist coupled with the fact that both the CD4 and CD8 Treg effector populations are much reduced from their baseline levels, one might anticipate that EAE severity becomes greater as we increase the probability of MHC down-regulation following negative signalling.

Figure 4.16: 3-dimensional plot illustrating the size of the effector CD4 Th1 population on day 40 at all of the possible pairings of probability parameters. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

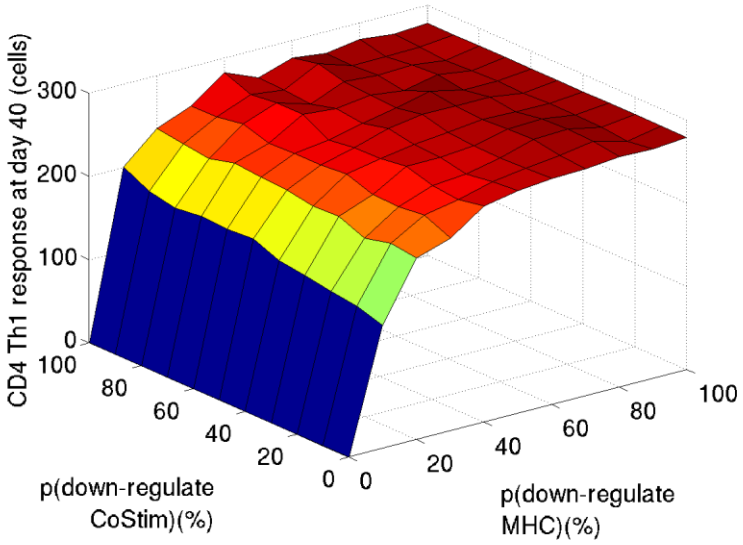

4.5.10 Effect on Mean EAE Severity

The expectation of increased EAE severity at higher probability parameter values is largely borne out by the simulation data, presumably due to the elevation of the peak population observed for effector CD4 Th1 and the much reduced peak populations of the Treg effectors.

Figure 4.17: 3-dimensional plot illustrating the mean EAE severity score at all of the possible pairings of probability parameter values. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

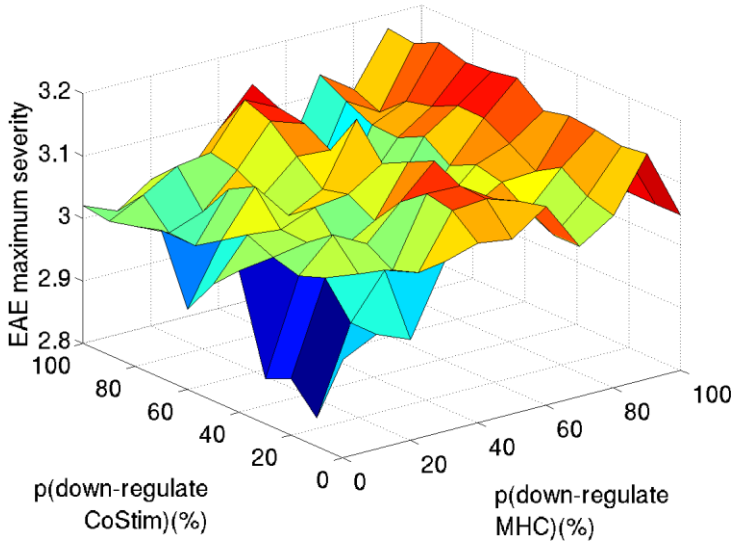

4.5.11 Effect on EAE Severity at 40 Days

For EAE severity at day 40 it is hard to discern any significant trend. However, it does appear that stronger negative signalling i.e. greater probability of MHC and / or CoStim being down-regulated raises the EAE Severity at day 40.

Figure 4.18: 3-dimensional plot illustrating the mean EAE severity score on day 40 of simulation at all of the possible pairings of probability parameter values. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour.

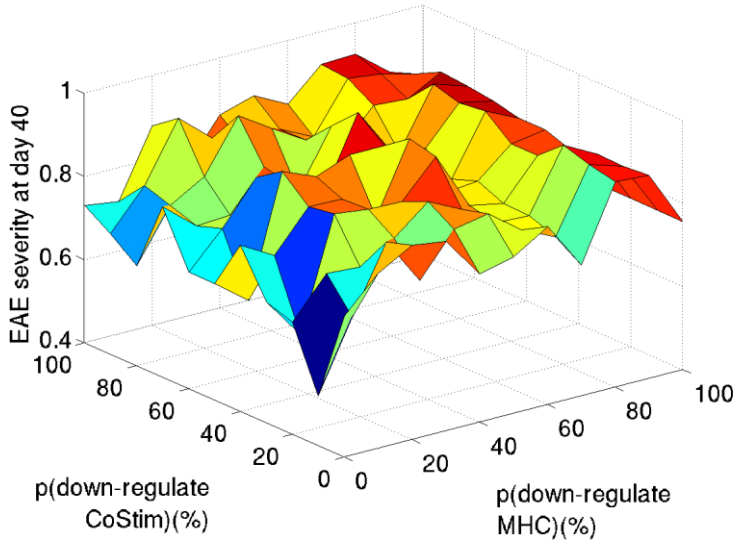

4.6 Summary of Findings

A common theme of all the response plots is that the value of the CoStim down-regulation probability (p(down-regulate CoStim)) appeared to have relatively little influence on the peak responses. The probability of MHC down-regulation (p(down-regulate MHC)) in contrast, had a dramatic impact on the peak responses, with peak Treg effector populations falling to less than half their baseline values when it was set to 10% from 0%. This is perhaps not so surprising; for every CD8 Treg that a DC primes there is a CD200-CD200R interaction (and thus negative signalling) since expression is immediate. In this manner, small changes to the MHC probability parameter has potentially dramatic effects on T-cell populations, whereas changes to the CoStim parameter are not as influential. We propose that this is the case because in the current model, CoStim expression can be switched back on once down-regulated, but MHC cannot (see Section 4.3.2.2).

4.6.1 Other Immunological Effects of the Parameter Values Changes

Comparing the response curves of runs with differing p(down-regulate CoStim) but constant p(down-regulate MHC) we see that most of the responses are essentially unchanged between runs (data not presented). This is borne out by A-Test (introduced in Section 3.4.3) statistics (presented in Appendix E.2) which allow us to compare data sets from two experiments differing only in one parameter setting.

The chief effect of allowing the down-regulation of CoStim expression on DC following negative signalling is that we now see one or two (literally) partially activated T-cells (Th1/CD4 Treg/CD8 Treg) in the resulting responses. This reflects the fact that T-cells in the simulator can only access the proliferating state directly from the naïve state if they receive both signal 1 and signal 2 i.e. recognise and bind cognate antigen **and** receive costimulation. This is surprising as one might reasonably have expected a good many more partially activated T-cells owing to the lower levels of CoStim expression. However, it is possible that, as the down-regulation of MHC and CoStim upon receipt of the negative signal is probabilistic, CoStim may be down-regulated while MHC is not. Under these circumstances, T-cells can still bind to APCs and if the APCs perceive sufficient Type 1 cytokine in the environment then they may be re-induced to express CoStim as no explicit bar to this was added to our model. It remains to be verified if this is indeed what has happened, this being relatively easily achieved by logging how many times CoStim expression has been up- or down-regulated on a particular DC.

There is also evidence of a slightly elevated rate of neuron killing after 30 days as compared to the baseline.

Comparing experiments which differ in p(down-regulate MHC) value and having the same p(down-regulate CoStim) the differences are much more marked. The peak population of Treg effectors is much reduced compared to the baseline (data not presented). There is much less Treg priming occurring in non-zero p(down-regulate MHC) simulations – CD4 Treg priming in the spleen is radically reduced (5000 in baseline and just 400 with p(down-regulate MHC) = 10%, p(down-regulate CoStim) = 0%) (data presented in Figure 4.19).

There are similar numbers of CD4 Th1 effectors in the CNS, but these persist after day 20 if p(down-regulate MHC) is greater than zero. Allied to this, there is neuron killing beyond 30 days in simulations with p(down-regulate MHC) > 0% and the cumulative count of neurons killed is much greater (50,000 as opposed to 30,000 in the baseline)(data not presented).

Coupled with this the CD4 Th1 population in general falls off much more slowly when we allow MHC down-regulation on DC. APC behaviour is also modified by the parameter change. There are more immature APC in the spleen after day 20 (data presented in Figure 4.20), though there are fewer immature APC in the CNS (but an elevated number of immunogenic APC) when p(down-regulate MHC) > 0% (data presented in Figure 4.21).

Figure 4.19: the effect of increasing p(down-regulate MHC) from 0% to 10% on CD4 Treg priming in the spleen, CLN and the SLO in the modified simulation. In the plot below, the baseline experimental results are shown by the dashed lines and the experiment in which MHC expression has a 10% probability of being switched off by negative signalling ('Stop MHC') is shown by solid lines. The different coloured lines represent different simulator compartments as indicated in the plot legend.

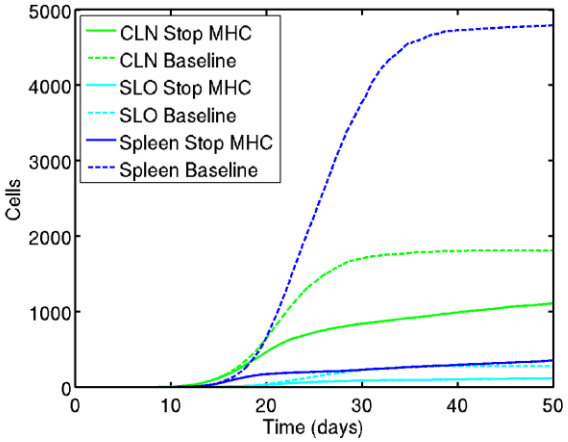

Figure 4.20: the effect of increasing p(down-regulate MHC) from 0% to 10% on the number of immature APC in the spleen in the modified simulation. In the plot below, the baseline ('Base') experimental results are shown by the dashed lines and the experiment in which MHC expression has a 10% probability of being switched off by negative signalling ('MHC') is shown by solid lines. The different coloured lines represent the different APC states within the simulator (Immature (Imm.), Tolerogenic (Tolero.) and Immunogenic (Immuno.).

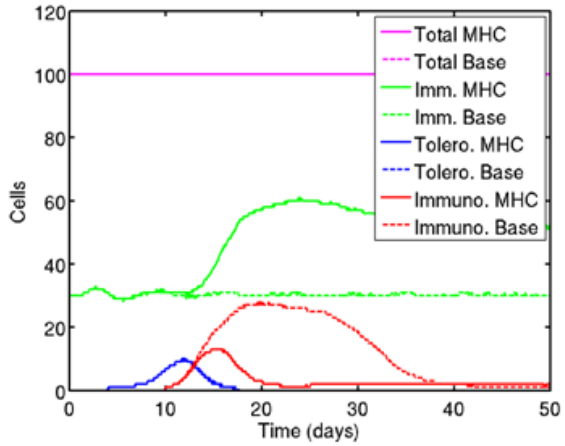

Figure 4.21: the effect of increasing p(down-regulate MHC) from 0% to 10% on the numbers of APCs in the CNS in the modified simulation. In the plot below, the baseline ('Base') experimental results are shown by the dashed lines and the experiment in which MHC expression has a 10% probability of being switched off by negative signalling ('MHC') is shown by solid lines. The different coloured lines represent the different APC states within the simulator (Immature (Imm.), Tolerogenic (Tolero.) and Immunogenic (Immuno.).

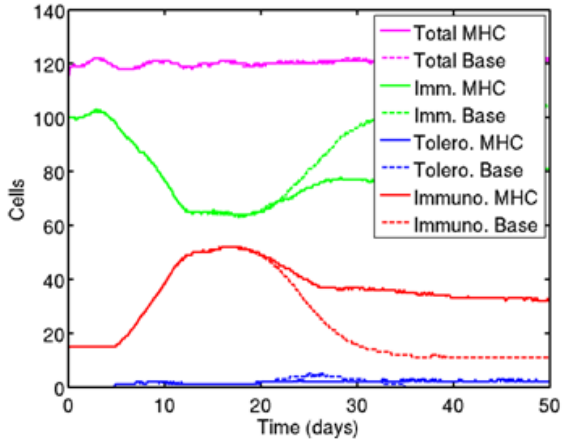

4.7 A Further Factorial Analysis

As even our smallest non-baseline parameter pairing (i.e. 10%, 10%) gave Treg and Th1 effector populations quite significantly changed from the baseline values, we decided to explore this region of parameter space in more detail.

Further simulator runs with progressively smaller values of p(down-regulate MHC) indicated that values between 0% and 1% might return responses reasonably close to those in the baseline experiment.

Therefore, we decided that a second parameter analysis to explore the parameter space for p(down-regulate MHC) and p(down-regulate CoStim) between 0% and 1% in steps of 0.1% would be appropriate. The effect of these parameter settings on peak effector responses is presented below.

4.7.1 Effect on Peak CD4 Th1 Population

Even at these low probabilities of down-regulating MHC expression on DC receiving the CD200-CD200R negative signal, we observe some shift away from baseline figures in terms of the peak CD4 Th1 effector population, though this is only 1-2 cells at p(down-regulate MHC) = 0.1%. The probability of down-regulating MHC appears to have a less significant effect on peak CD4 Th1 population size than the probability of down-regulating CoStim at low values of the parameters. However, the data is probably noise based on the values of the A-Test scores for variations of these parameters between experiments (see Appendix E.2).

Figure 4.22: 3-dimensional plot illustrating the peak CD4 Th1 effector population at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

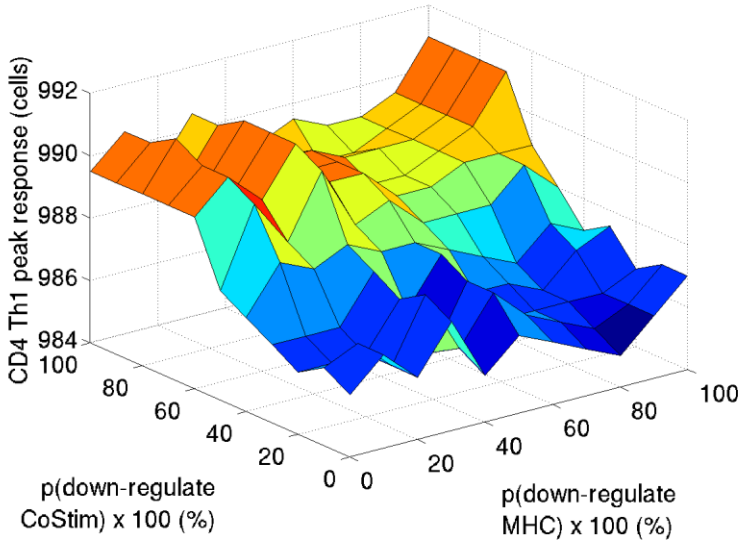

4.7.2 Effect on Time Taken to Reach Peak CD4 Th1 Effector Population

The time taken to attain peak CD4 Th1 effector population is not shifted too far away from the baseline level even at p(down-regulate MHC) = 1%. The 3D plot for the parameter mapping of this response is not included in the main flow of the document since the data is probably noise, a conclusion borne out by A-Test scores for varying each of the probability parameters separately. Instead the plot is presented in Appendix E.1 and the A-Test data in Appendix E.3 for reference.

4.7.3 Effect on Peak CD4 Th2 Effector Population

There does not appear to be any clear relationship between the two down-regulation probabilities and the CD4 Th2 effector peak response. The A-Test scores in Appendix E.3 suggest that again, the data is noise and so the plot has been presented in Appendix E.1.

4.7.4 Effect on Time Taken to Reach Peak CD4 Th2 Effector Population

As with the peak CD4 Th2 response, there does not seem to be any straightforward relationship between the down-regulation probabilities and the time of peak CD4 Th2 response. The A-Test scores (presented in Appendix E.3) suggest that the data is noise, so the plot is presented for completeness in Appendix E.1.

4.7.5 Effect on Peak CD4 Treg Effector Population

Even at p(down-regulate MHC) = 1% the peak population of CD4 Treg effectors has fallen from ~1000 to less than 600. The peak response is only reasonably close to the baseline value at p(down-regulate MHC) < 0.1%. The value of p(down-regulate CoStim) has relatively little impact compared to the probability of down-regulating MHC.

Figure 4.23: 3-dimensional plot illustrating the peak CD4 Treg effector population at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

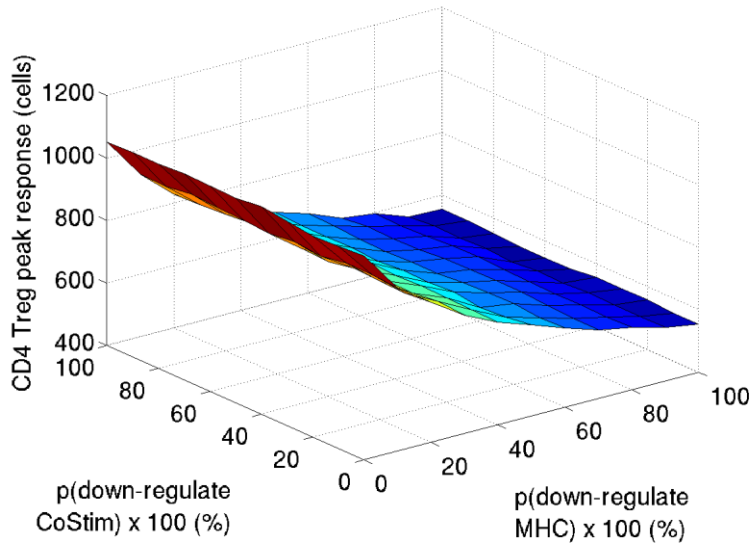

4.7.6 Effect on Time Taken to Reach Peak CD4 Treg Population

The effect of the parameters on the time taken to reach the peak CD4 Treg effector response mirrors their effect on the size of the response itself.

Figure 4.24: 3-dimensional plot illustrating the time taken to reach the peak CD4 Treg effector population at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

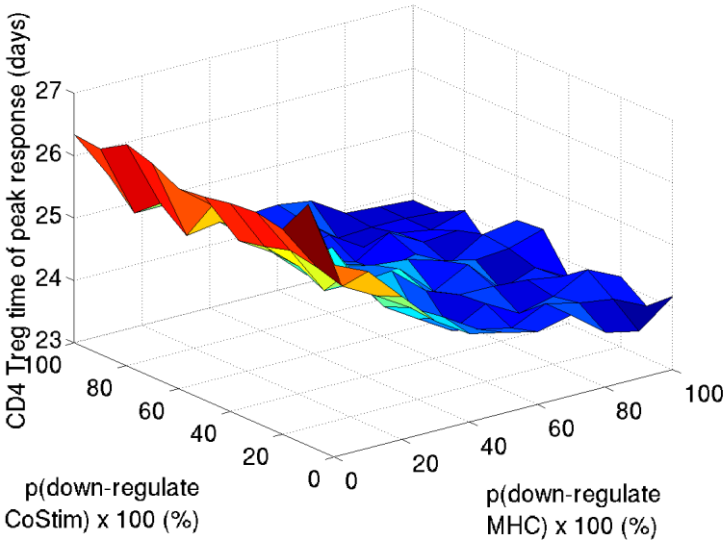

4.7.7 Effect on Peak CD8 Treg Effector Population

As with the CD4 Treg effector response, the CD8 Treg effector response is very sensitive to changes in p(down-regulate MHC) and falls from ~600 cells to ~500 cells in changing from 0% to 0.1%. Beyond this parameter value the response deviates increasingly from the baseline value.

Figure 4.25: 3-dimensional plot illustrating the peak CD8 Treg effector population at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

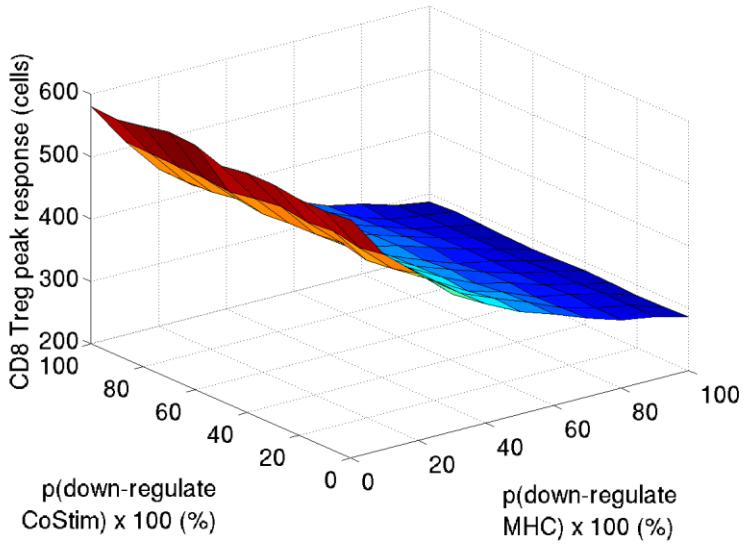

4.7.8 Effect on Time Taken to Reach Peak CD8 Treg Effector Population

The time taken to reach peak CD8 Treg effector response roughly echoes the shape of the peak response plot.

Figure 4.26: 3-dimensional plot illustrating the time taken to reach the peak CD8 Treg effector population at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

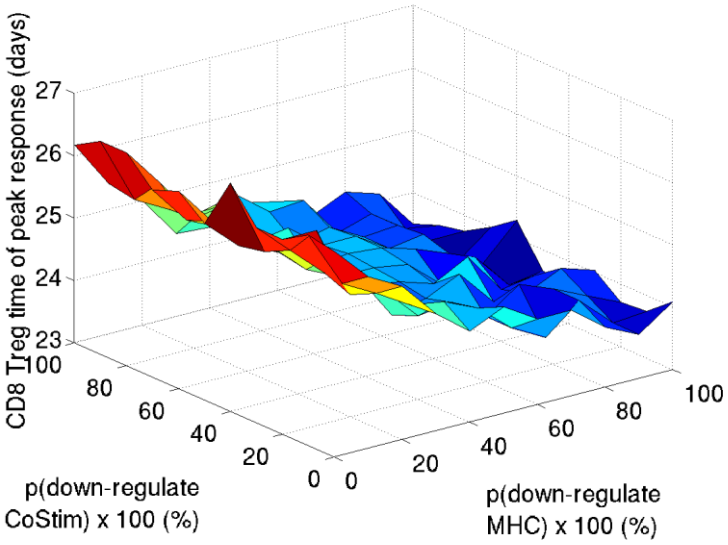

4.7.9 Effect on CD4 Th1 Effector Response at 40 Days

The population size of CD4 Th1 effectors at 40 days appears to be less sensitive to the value of p(down-regulate MHC) than the other responses. The baseline behaviour is maintained until p(down-regulate MHC) reaches 0.5%.

Figure 4.27: 3-dimensional plot illustrating the CD4 Th1 effector population persisting at day 40 at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

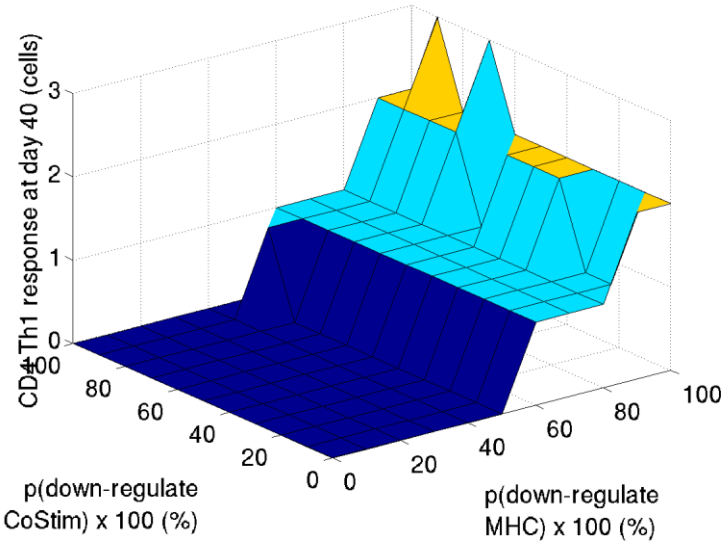

4.7.10 Effect on Mean EAE Severity

Mean severity of EAE appears to be less susceptible to changes in the two down-regulation probabilities at these low values.

Figure 4.28: 3-dimensional plot illustrating the mean EAE severity at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

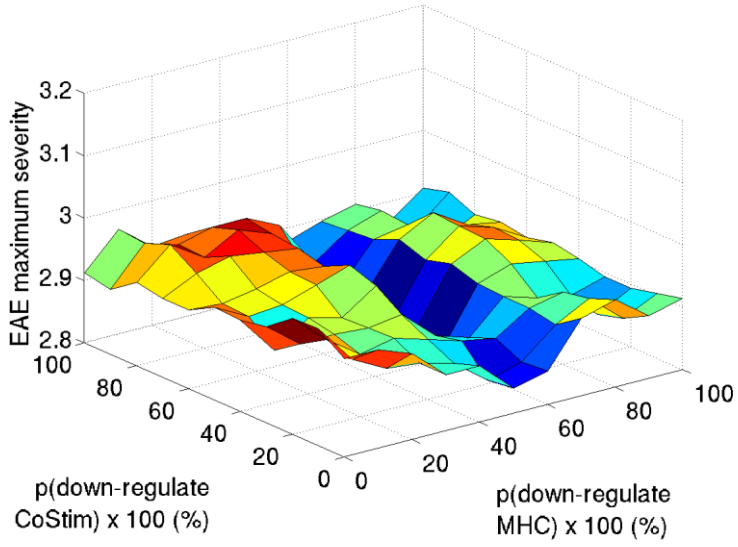

4.7.11 Effect on EAE Severity at 40 Days

The EAE Severity at 40 days plot appears to behave very like the mean EAE Severity plot.

Figure 4.29: 3-dimensional plot illustrating the mean EAE severity at 40 days at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

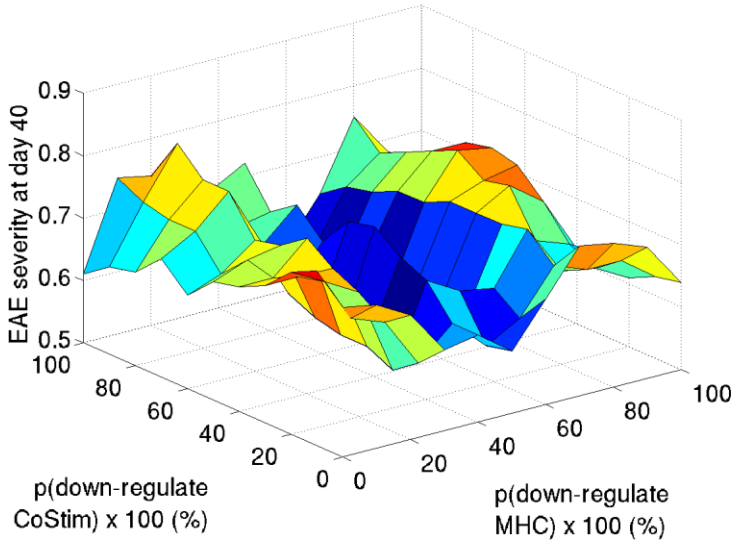

4.8 A Brief Investigation of the Extent of Interaction Between the two Disease Regulation Mechanisms

In the augmented domain model presented in Section 4.3 we have added an additional regulatory mechanism to the existing CD8 Treg mediated killing mechanism. Since the CD200-CD200R axis affects the ability of DC to prime T cell populations, the two regulatory mechanisms exhibit some degree of interaction. This interaction depends on the extent to which the T cell priming capacity of DCs is affected by the added axis and by the efficiency of CD4 Th1 killing by CD8 Treg. In the work presented in this thesis we have assumed that CD8 Treg killing of CD4 Th1 is 100% efficient. However, there is some evidence that this is not the case [Tang *et al.* 2006]. Altering the efficiency of CD4 Th1 killing in our model offers us a means to assess how much the CD8 Treg mediated killing mechanism of disease regulation is interacting with the CD200-CD200R axis mediated down-regulation of CD4 priming by DCs and hence to re-balance the two mechanisms within the simulation to return a more baseline-like behaviour.

To investigate the extent of such regulatory mechanism interaction we conducted a further experiment in which the efficiency of CD4 Th1 killing by CD8 Treg was adjusted from its baseline level of 100% to a lower level of 30% whilst implementing the negative signalling with both probability parameters set at 1%. The results of this experiment would serve as a guide as to whether re-balancing was needed

However, an investigation into the robustness of the system described above to changes in the killing efficiency of CD4 Th1 by CD8 Treg (a ‘robustness analysis’ [Read *et al.* 2011]) showed that the only response that changed significantly relative to the simulation using killing efficiency of 100% was the population of CD4 Th1 remaining at day 40 of the simulation. The number of CD4 Th1 remaining at 40 days was 26 cells compared to just 2 in the experiment using 100% killing efficiency and 0 in the baseline experiment. All other responses show no effect after reduction of the killing efficiency to 30%. This suggests that the model of the CD200-CD200R mediated regulation is probably too severe and is completely dominating the regulation of T-cell populations within the simulation.

4.9 Overall Conclusions

Overall, the results from the two factorial analyses suggest that the simulator behaviour is particularly sensitive to the ability of DC to carry on expressing MHC (this being particularly true of the CD4 Treg, CD8 Treg and CD4 Th1 effector responses). The system deviates quite significantly from the baseline T-cell effector responses once the probability of the negative signal down-regulating MHC expression on DC rises to 1%. It remains to decide whether these deviations lie within tolerable bounds and whether we are able to satisfactorily parameterise the effects of the CD200-CD200R negative signal on DC expression of MHC and CoStim.

It must be borne in mind that only the parameter pairing (0%, 0%) returns **exactly** the baseline behaviour as this represents the situation implemented in the baseline simulator. We should also bear in mind that it is entirely possible that the 'all-or-nothing' model of negative signalling adopted currently employed is probably too simplistic and that a 'phased' model in which negative signalling can reduce MHC and / or CoStim expression by some fraction less than 100% (and so subsequent signals can therefore compound to further reduce MHC or CoStim expression) might be more appropriate and would almost certainly more accurately reflect the biological situation. This 'phased' down-regulation would allow for a much lower level of response to negative signalling than the current model permits.

The current analysis was carried out assuming that CD8 Treg are 100% efficient at killing CD4 Th1. Thorough investigation of the effects of lower killing efficiency in tandem with the implementation of the CD200-CD200R axis is needed and requires a more extensive re-parameterization to properly recalibrate the simulator. Our initial experiment with lowered killing efficiency suggests that this parameter has no significant effect on simulation behaviour when using our current model of the CD200-CD200R axis even at low values of the two down-regulation probability parameters. This suggests that the model of the CD200-CD200R regulation axis is too simplistic and too severe, making it difficult to relate the results of this preliminary exploration back to the real domain.

Chapter 5: Conclusions and Further Work

5.1 Contribution

We conducted experimentation to assess the significance of two hypotheses in explaining the observed CD8 Treg effector population size in a CD4 Treg abrogation experiment. The hypotheses investigated were:

H1: timing of Qa-1 expression on DCs has a significant impact on CD8 Treg population.

H2: competition between CD4 and CD8 Treg for binding space around DC has a significant impact on CD8 Treg population.

Hypothesis significance was assessed by means of the non-parametric effect size A-test [Vargha and Delaney 2000].

The domain expert subsequently indicated an interest in investigating a different mechanism for the regulation of disease not currently implemented in the simulation. In a separate experiment we augmented the simulation to incorporate a simplified model of this additional, but poorly understood, immune regulatory axis. The implemented model incorporated two new simulation parameters which represented probabilities that expression of certain molecules by DCs would be down-regulated. We sought to assign values to these by exploring simulation behaviour over all combinations of values of the two parameters. We then attempted to identify which values returned behaviour close to that of the baseline simulation (where the two parameters are both effectively set to 0%).

5.2 Conclusions

5.2.1 Conclusions Drawn from the Experimentation on CD8 Treg Activation by Dendritic Cells

In chapter 3 of this thesis we assessed the significance of two hypotheses in determining the peak population of CD8 Treg effectors in an experiment where the CD4 Treg population had been abrogated [Williams 2010b], clearly demonstrating how *in silico* experimentation allows us to examine hypothetical circumstances that would have been impossible to test in the same way in a wet laboratory. This body of experimentation required no changes to be made to the existing domain model [Read 2011].

Both hypotheses were found to be significant in explaining the observed CD8 Treg population levels, illustrating that both the timing of Qa-1 expression by DCs and the spatial competition of CD4 with CD8 Treg for binding space around DC have a profound impact on the scale of the cytotoxic response following the onset of disease. This implies that CD4 Treg are capable of exerting a dual influence on the size of the CD8 Treg population attainable.

Although these results are grounded specifically in the EAE system, the results are potentially applicable to any disease system where CD4 Th and CD8 T-cells are involved in the immune response and have significance for attempts to create therapies based on boosting the CD8 Treg population [Beeston *et al.* 2010], for example in inflammatory bowel disease where CD8 Treg have been shown to cure colitis in mice [Endharti *et al.* 2010].

5.2.2 Conclusions Drawn from the Experimentation on Adding the CD200-CD200R Regulatory Axis to the Simulation

In chapter 4 of this thesis we employed the CoSMoS process to implement additional functionality in our model of EAE. The implementation of the additional cellular functions required changes to the domain model [Read 2011], which were discussed with the domain expert. The CoSMoS process facilitated structured thinking about what changes to the logic of the simulation were actually required and how best to implement them.

Our additions to the basic domain model [Read 2011] were simple in that we required expression of a new cell surface protein on CD8 Treg and of another on DC. The two new proteins were allowed to interact, with the interaction having predefined effects on the behaviour of the DC.

The interaction ('negative signal') was defined in the model, at the suggestion of the domain expert, as probabilistically switching off the expression of MHC compounds and / or costimulatory molecules. This introduced two new parameters to our model; the probability that negative signalling switches off MHC expression and the probability that it switches off CoStim expression. We restricted our model to a simple switching off of MHC and / or CoStim expression by DC as we have no prior knowledge as to the effects of the pathway on the system and we wished to explore the results of the augmentation on the simulation. This representation was considered to be more realistic than a model in which all negative signals caused complete down-regulation of DC activity and easier to implement than a model in which each negative signal produced partial down-regulation of MHC and / or CoStim by DC.

We implemented the new regulatory pathway on the advice of the domain expert, who believes this pathway to be important. However, we were unable to obtain any guidance from the literature regarding the possible values of the two parameters. In an attempt to locate values for the two parameters that would replicate the baseline simulation, we carried out two factorial analyses and found that even a value of 1% for the probability of turning off MHC expression yielded significant impact on the peak populations of certain key cell populations (CD4 Th1, CD4 Treg and CD8 Treg).

It is unlikely that the mechanism abstracted in our model is a fair representation of what happens *in vivo* as the T-cell populations are radically reduced in our simulation. For this reason we need to explore other representations of the system which exhibit a better balance between CD8 Treg killing of CD4 Th1 and CD8 down-regulation of DC. Alternative abstractions are considered in the section on further work (Section 5.3.3). In this light it is quite possible that our initial model is too simplistic and severe. This line of experimentation would certainly benefit from further *in vivo* data concerning the regulatory pathway.

5.3 Further Work

In the following sections, we consider further strands of investigation that could extend and strengthen the work presented in the thesis.

5.3.1 General Considerations about the EAE Simulator

At a general level there are two considerations to be aware of that may impact on the reliability of the results of any simulation: the limited size of the model compared to that of the real system and the amount of detail that is abstracted away in the model.

The immune system consists of some 10^{12} cells [Seiden and Celada 1992], with 10^{9} lymphocytes in circulation [Janeway *et al.* 2008]. The EAE Simulator employed in the experiments presented in this thesis uses cell populations significantly lower than these numbers.

This may be a concern in that we may need to account for size-related artefacts in the simulation [Seiden and Celada 1992]. However, it could be validly argued that there is very little purpose in creating a simulation that can handle truly biological sized populations of cells (along with the concomitant demands on computing power that this would entail) if the results of smaller, tractable simulations adequately address the questions we ask of them.

Initial cell populations are simulation parameters and as such their influence on simulation behaviour can be examined via a sensitivity analysis as described in [Read *et al.* 2011]. If the cell populations are not found to be influential then they are not anticipated to bring about any scale-related effects in the simulation results. However, .if one population is found to be influential when all others are held constant, then there could be population size related effects. Generally one would change all population numbers simultaneously in such a way as to preserve population ratios. In this way EAE scores and other simulation results of interest would remain constant. Generally, if there are qualitative changes following changes in initial cell populations, one must be cautious in the interpretation of simulation results. The course of remedial action chosen would then be problem specific.

Our second consideration is the level of detail included in our model. There are two features of the immune system that are potentially relevant to EAE, which are abstracted out of our current domain model.

The first of these concerns a sub-population of T-helper cells, the Th17 cells. These cells are generated in a different cytokine milieu (generally the cytokine environment of an early immune response [Janeway *et al.* 2008]) to the Th1 and Th2 cells and have been implicated in the early phase of EAE induction [Reboldi *et al.* 2009]. Our domain model does not include these cells, and indeed abstracts away the period of innate immune response into a special immunization method and by the setting up of an initial population of mature and immature DC (see Appendix B.7).

Again, one might validly argue that the further complexity of the model entailed in explicitly including this T cell subpopulation in our model serves no useful purpose given that our current model gives reasonable and encouraging results.

In Chapter 4 of this thesis we deal with adding complexity to our domain model by allowing CD8 Treg effectors to express CD200 and mature DC to express CD200R and by allowing the interaction of CD200 with CD200R to bring about inhibitory effects on the DC population. Among the effects included in our model was the probabilistic down-regulation of costimulatory molecule expression by DC.

In actual fact, our model abstracts away considerable detail about the signals exchanged by DC and T cells. In particular the simulation uses the concept of 'costimulatory molecule' to cover several distinct entities which are expressed on DC. These molecules are ligands to receptors which can either stimulate or inhibit T cell responses to antigens [Keir and Sharpe 2005]. These molecular interactions serve to fine tune immune response and consist of many different receptors and ligands [Crawford and Wherry 2009, Orabona *et al.* 2004].

It is more difficult to explore abstraction space than parameter space, as it is difficult to be sure that one has explored every possible abstraction of a system, possibly owing to incomplete domain knowledge. In this instance the only possible procedure is careful iterative calibration in collaboration with a domain expert as described in [*Read et al.* 2011]. This remains a problem for modelling that is not fully resolved.

It is important that we remain aware of the detail that has been abstracted away in the domain model and as we have implemented an, albeit very simplified model of the CD200-CD200R axis, it is not inconceivable that a more explicit description of DC to T-cell signalling may one day also be included in the simulation. This would become a particularly relevant issue if, in future, some aspect of the simulation concerned with mechanisms of costimulation were not to produce results in agreement with what is observed *in vivo*. However, we have no evidence as yet that such a discrepancy would be likely and the potential benefits of adding this level of detail to the model would have to be carefully assessed in light of our current work with the CD200-CD200R axis.

5.3.2 Further Work Arising from the Chapter 3 Experimentation

A general concern that has emerged from this work has been that when comparing experiments using the A-Test, we need the comparison to be as fair as we can make it. With valid, but 'unfair' comparisons such as those cited in Section 3.9 we can only validly assess which of our two hypotheses are significant.

Ideally, we should continue to attempt to develop a completely fair comparison on which to use the A-Test. This would allow us to assess the relative significance of the two hypotheses by directly comparing A-Test scores.

5.3.2.1 Further Work Arising from the Analysis of Dendritic Cell Age at Time of Licensing

There remain immunologically interesting data to be obtained from our simulations. The simulator makes much of this relatively straightforward to obtain, requiring only alterations to the code for storage and output of the required quantities.

For example, we could use the simulation data to investigate the extent to which mature DC phagocytose other cells (immature DC have greater phagocytic ability than mature DC). This would provide us with a more detailed picture of DC behaviour within the simulation, which would be potentially useful in helping to explain or interpret T-cell population dynamics as DC play a pivotal role in priming T-cells.

Also of interest is whether DC acquire their peptide presentation capabilities when they are immature or only after maturation. This is interesting because DC are more phagocytic when immature, but they spend more of their life span in the mature state under current simulation parameterization (48 hours in the immature state versus 110 hours in the mature state).

We could also record data from the simulation that would allow us to assess how quickly T-cells interact with DCs within the simulation. At the moment we have an approximation to how quickly CD4 Treg can reach and bind DCs in our age at licensing data, but this measures time from the perspective of the DC and not the CD4 Treg.

These ideas would provide us with greater detail on how the simulation is behaving at a cellular level as opposed to the system level. There is potential for collating data about cell behaviours that may be verifiable by wet-lab experimentation and this would help to increase trust in the simulation. Such data on the behaviour of specific cell types could also potentially be relevant to other disease systems, particular those mediated by CD4 Th and regulated by CD8 Treg.

More directly related to the work presented in the thesis, it has been suggested that we might investigate other representations of how DC age at licensing is related to DC time of creation. In the 'age of licensing' experiment in Section 3.6.3 we have

employed a very simple method of demonstrating the trend toward quicker licensing of DC later on in the simulation i.e. we have divided the population of DC into two sub-populations. However, we could have provided a more detailed picture. e.g. a three-dimensional plot of DC time of creation versus DC age at time of licensing versus frequency of DCs in the data bin. This would enable us to gain insight into the lifecycle of DC within the simulator, potentially suggesting wet lab experiments that would serve to corroborate the *in silico* results.

5.3.2.2 Further Work Arising from the Spatial Saturation Experiments

Whilst conducting investigation into the significance of CD8 / CD4 Treg spatial competition, we omitted to thoroughly investigate how the spatial saturation of DCs varied over simulation time for each peptide presentation profile (in a manner analogous to the 'Age of Licensing' analysis in Section 5.3.2.1). This would enable us to gain insight into the population dynamics of T-cells and DC and into their patterns of interaction over the course of an immune response.

We also need to investigate further the idea of increasing the maximal crowding parameter introduced in Section 3.8. We would like to know if there is a point beyond which increases in this parameter cease to increase peak CD8 Treg effector population and if this point is reached before the simulation is rendered intractable by the CD8 Treg population explosion. This would provide an estimate of how crowded a DC could theoretically become before CD8 Treg were no longer able to reach them and become activated. It would also serve to model the effect on CD8 Treg effector population size, were this to occur.

These ideas have wider immunological relevance in that they give insight into T-cell population dynamics and into T-cell interactions with DC during the course of an immune response. As argued previously, this information could be particularly pertinent to the development of therapies for autoimmunity which aim to boost CD8 Treg population sizes and / or activity.

5.3.3 Further Work Arising from the Implementation of the CD200-CD200R Regulatory Pathway

The results of our experimentation on adding a representation of the CD200-CD200R pathway to our simulation, led us to conclude that our representation of that pathway was probably too severe. Therefore, in this section we propose further potential abstractions of the pathway that could be investigated further.

The work presented in Chapter 4 introduced the idea of rescue from negative signalling. This form of rescue can in theory occur because the possibility that DC can be re-stimulated into expressing costimulatory molecules or re-licensed to express Qa-1 after a negative signalling event has down-regulated them is not explicitly ruled out in the altered domain model presented in Section 4.3. However, DCs can only up-regulate MHC-II expression at the time of maturity in our current domain model. Janeway *et al.* 2008 state that DC can act as APCs i.e. express MHC

compounds, upon maturity but do not mention any other circumstances under which MHC-II expression can be up-regulated. In our model, this means that once negative signalling has down-regulated MHC-II expression, it remains down-regulated. This in turn means that DC cannot be re-licensed to express Qa-1 (CD4 T-cells cannot bind to a DC that is not expressing MHC-II). This means that our domain model is possibly over proscriptive in that it allows negative signalling to cause DC to down-regulate both MHC-II and Qa-1 expression, thus effectively ruling out the possibility of rescue by CD4 T-cells. We propose that the next phase of investigation should focus on a less severe domain model in which the negative signal triggers Qa-1 but not MHC-II down-regulation on DC.

The simulations carried out in Chapter 4 assumed that killing of CD4 Th by CD8 Treg is 100% efficient. However, it seems unrealistic to assume that killing is 100% efficient [Tang *et al.* 2006] and CD4 Th1 populations may not be as severely affected by the additional regulatory axis as our current model suggests. This would mean that CD4 Th populations would not be reduced as much by CD8 Treg mediated killing as our baseline simulation currently predicts. Without the additional regulatory mechanism provided by the CD200-CD200R axis, EAE Severity scores would also be higher than we currently observe. However, with this axis in place the two regulatory pathways should co-operate in keeping populations of self-reactive CD4 Th in check. The question for our simulation then lies in just how the two mechanisms are balanced.

We have briefly investigated the effect of reducing the killing efficiency parameter to 30% at low values of the new model probability parameters and found the change to show no significant effect (as judged by A-Test scores). To fully investigate the interplay of the two disease regulation mechanisms i.e. direct CD8 Treg mediated killing of CD4 Th and CD200-CD200R axis mediated attenuation of CD4 Th priming by DC, we would need to conduct a factorial analysis for the two probability parameters plus the killing efficiency parameter. Ideally a full sensitivity analysis would be performed over all simulation parameters including the two new probability parameters, permitting a rebalancing of the two regulatory mechanisms within the simulation.

We mentioned in Section 5.2.2 that our model of negative signalling may be too simplistic. In fact, the domain expert has suggested that negative signal has a significant effect on MHC and / or CoStim expression by DC but that it is not 'all or nothing' (which is what our current model assumed for the sake of simplicity). In light of this a better model may be to allow each negative signal to probabilistically reduce (but not turn off completely) MHC and / or CoStim expression by DC, thus requiring several negative signals to be received by a DC before it stops expressing these cell-surface proteins. For example, the first negative signal might reduce MHC and / or CoStim expression by say 10%, each subsequent negative signal reducing expression by some preset decrement over and above the existing extent of down-regulation until MHC / CoStim expression is totally down-regulated.

Another simplification made for the sake of our model which may impact significantly on the scale of the simulation response to negative signalling is the fact that we did not require the CD8 Treg to receive local stimulation prior to expressing CD200. This in effect means that all CD8 Treg effectors express CD200 capable of signalling DC via CD200R. It would therefore be instructive to implement a model in which the CD8 Treg do require local stimulation prior to expression of CD200. This would be anticipated to reduce the numbers of CD8 Treg capable of sending a negative signal to DC and thus moderating the extent of MHC and CoStim down-regulation occurring throughout the simulation.

Although we were aware that CD8 Treg can enter the CNS [Zozulya *et al.* 2009], the domain expert advised us not to investigate this in our preliminary investigations of the regulatory axis. Allowing CD8 Treg to enter the CNS would require that we also permit neurons to express CD200 and microglia to express CD200R, thus introducing a great deal more complexity into our model before we even fully understand our current model. It is also important that we rebalance the regulatory mechanisms within the simulation so that both interact in a reasonable manner (at the moment the CD200-CD200R negative signalling completely dominates the simulation). Ultimately, however, we would like to extend the domain model to incorporate the full complexity of the CD200-CD200R regulatory axis.

5.4 Summary Statement

In Chapter 1 we stated that the goal of the thesis was two-fold: firstly to identify relevant questions that we may direct to our domain expert to address and secondly to gain further insight into the operation of the simulation in hypothetical situations within the EAE system which are not easily implemented *in vivo* or *in vitro*. We now address how far we feel the thesis has attained its aims.

We have utilised *in silico* experimentation to assess the significance of two hypotheses in explaining the results of a previous experiment that would have been impossible to conduct in a wet lab. The results of this hypothesis testing, though useful in explaining a previous observation, could also potentially have consequences for the design of therapies based on boosting CD8 Treg populations.

We have also implemented a limited model of the CD200-CD200R immune regulatory pathway in the simulation. Although this is tempered by the knowledge that our initial model is probably too strict, the experimentation has suggested a number of novel avenues of investigation. It has also prompted a need for specific data concerning the behaviour of DCs upon receiving a CD200-CD200R 'negative signal'.

In conclusion, we have demonstrated the utility of *in silico* experimentation in answering questions that could not easily be investigated in a wet lab, and also the potential relevance of such work to the wider practice of immunology.

Appendix A: A Brief Glossary of Immunological Terms Used Within the Thesis

This appendix briefly presents definitions and descriptions of certain immunological terms and concepts that are used in the text of this thesis. Some are not strictly relevant to the design of the domain model of the simulator, though they are important in the understanding of the results obtained and the limitations of the model utilised for our simulations.

A.1 Major Histocompatibility Complex (MHC) Molecules

These molecules are antigen-presenting marker molecules which the immune system uses to distinguish cells that are harbouring invading pathogens from normal body cells (MHC Class I) and immune cells from other cells (MHC Class II) [Voet and Voet 2004].

All nucleated cells can express MHC Class I compounds, whereas only DC, macrophages and B-cells can express MHC Class II compounds [Kindt *et al.* 2007]. These are the cells that are activated by or activate CD4 T-cells. This allows CTLs to attack only infected host cells, while permitting T helper cells to interact only with immune system cells. APCs display both MHC Classes and can stimulate development of both Tc and Th cells.

The MHC compounds are represented in the simulator as generic class I and class II molecule objects. Methods also exist to test for expression of the MHC class 1b compound Qa-1 on DCs.

A.2 Priming

The activation of a T-cell is sometimes called 'priming'. A primed T-cell can enter the cell cycle and continue its development into an effector cell. Priming occurs when a T-cell binds its cognate MHC-antigen complex (MHC II in case of CD4 Th and MHC I for CD8 Tc) [Kindt *et al.* 2007, Janeway *et al.* 2008].

A.3 Cross Priming

Normally endogenous antigens are proteolysed and displayed on MHC Class I compounds. Similarly, exogenous antigens are presented on MHC Class II compounds. Cross-priming is the initiation of CD8 T-cell response to antigens which are not synthesised by APCs.

In cross-presentation exogenous antigen is displayed on MHC Class I compounds. [Kurts *et al.* 2010, Brode and Macary 2004]. Cross-presentation of exogenous antigens on MHC Class I molecules occurs through several distinct cellular pathways the precise description of which lies outside of the scope of this thesis. Interested readers are referred to [Brode and Macary 2004].

Cross-presentation offers an important mechanism in the activation of CTLs which defend the body against pathogens and tumours [Kurts *et al.* 2010]. To avoid the destruction of uninfected cells that have endocytosed microbial debris; the endocytosed material does not normally enter the MHC class I loading machinery.

Cross-presentation and cross-priming are essential to the CD8 T-cell response to viruses that do not infect APC [Kurts *et al.* 2010].

A.4 Licensing

A textbook explanation of licensing is that inflammation following infection generally activates DCs as part of the innate immune response to the causative pathogen. When inflammation has not occurred then the DCs do not possess the necessary level of activation to stimulate a complete response from CD8 T-cells. CD4 T-cells then serve to stimulate DCs to an activation level sufficient that they can stimulate CD8 T-cell response. This action is termed 'licensing' – CD4 T-cells being said to 'license' DCs [Janeway *et al.* 2008]. However, the role of inflammation is contentious and there are several important unknown factors involved in the licensing process. For instance, is the licensing signal sent via the APC itself or is the process cytokine mediated? Are the CD4 T-cells bound to the APC at the same time as the CD8 Treg?

The requirement for DC licensing in the absence of inflammation is a mechanism for protecting against autoimmunity because it is in effect a requirement for an antigen to be recognised both by CD4 T-cells and CD8 T-cells before the CD8 T-cells can be stimulated to proliferate and differentiate into potentially deleterious CTLs.

The issue of DC licensing by CD4 Treg cells is particularly important in the experiments in Chapter 3 of this thesis and in the simulator licensing is understood to be simply a signal from CD4 Th (or CD4 Treg) to an APC in order to allow the APC to prime CD8 T-cells.

A.5 Apoptosis

Apoptosis is the name given to 'programmed cell death' i.e. the natural death of a cell. Apoptosis serves the function of removing unnecessary cells from the system, returning cell populations to their appropriate levels following clearance of infection and removal of potentially auto-reactive T-cells during selection

Induction of apoptosis involves different signals (e.g. caspases, FasL-Fas) depending on the cell type involved.

Appendix B: Implementation Specific Details of the Simulator

B.1 Simulator Packages

The purpose of a package diagram is to show how the software is structured with regard to the groupings of agents. Agents are grouped by intended function into packages of agents sharing similarly functions e.g. I/O, data logging etc.

Generally, in Java, packages correspond to the location of the different agents in the Java source code directory.

Figure B.1: A package diagram for the EAE Simulator. Since the package structure is implementation specific, the diagram presented here is mainly for the interested reader. The diagram is adapted from [Williams 2010b].

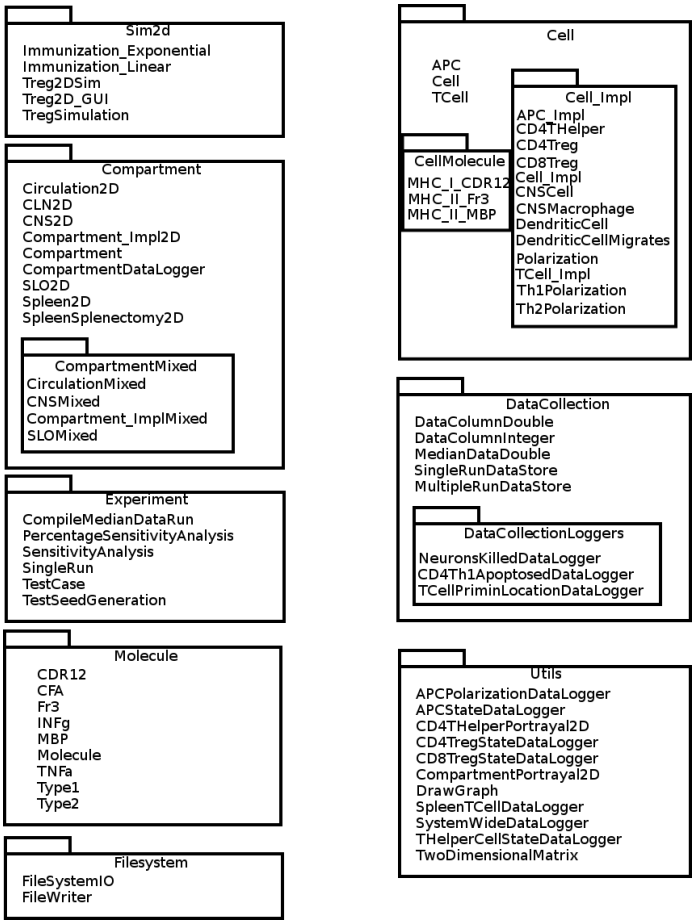

B.2 Interactions Between the Simulator Classes

The purpose of an association class diagram is to give an overview of how the varying classes within a project interact with each other. For biological complex systems such as the immune system in EAE, this can lead to highly connected and potentially unclear diagrams.

Figure B.2: An association class diagram for the EAE Simulator. The diagram aims to give an overview of the various types of interaction (or 'associations') that occur between the different agent types (the Java classes) within the simulator. Again, since the association of classes is implementation specific, this diagram is presented merely for the interested reader. Adapted from [Williams 2010b].

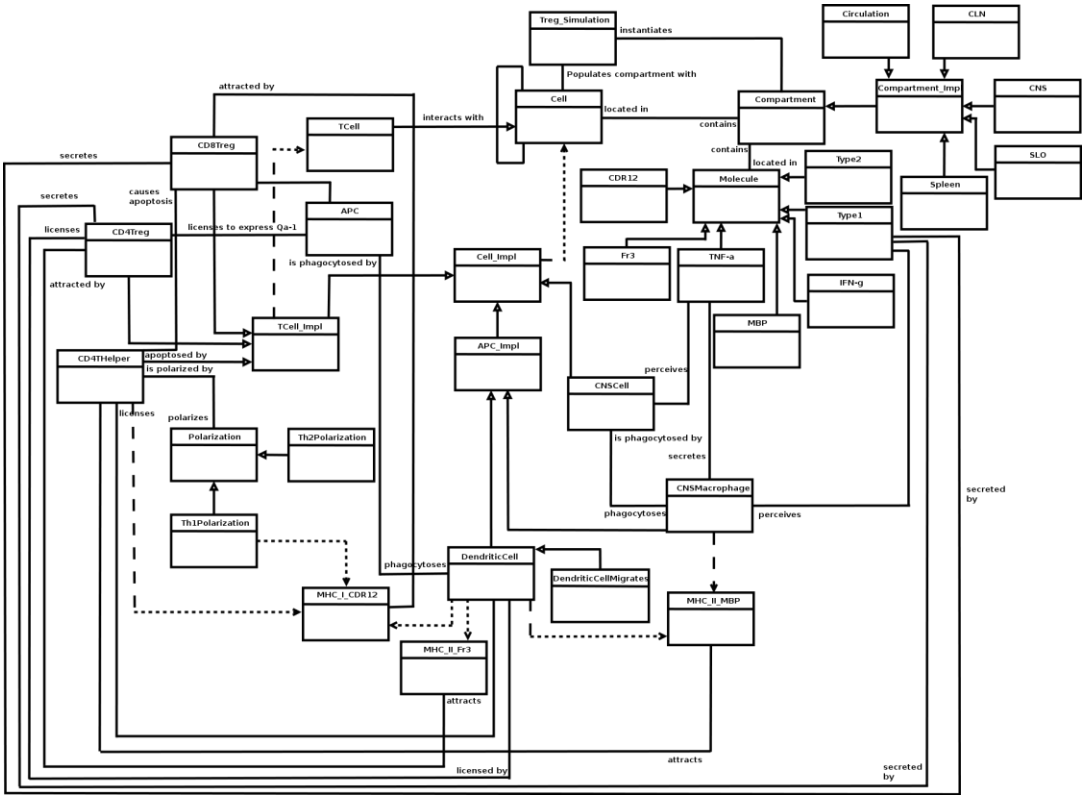

B.3 Cell Types Incorporated into the Domain Model

The following cell types are incorporated into the Domain Model for EAE employed in the simulator [Read *et al.* 2009a, Read 2011].

Figure B.3: The different cell populations defined within the EAE Simulator [Read 2011]. The figure also shows how the different cell types are related to each other i.e. their pattern of inheritance.

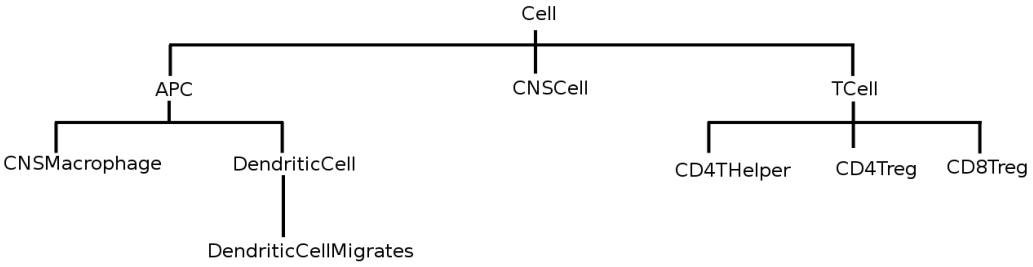

B.4 Simulator Compartments – Nature, Dimensions and Inter-communication

The EAE Simulator defines four different compartments which each represent a distinct location within the body where immune cells may reside during immune response. The compartments are the SLO which represents generic lymphatic tissue, the CLN which is the major lymph node in the neck and serves as a key access point for immune cells into the CNS. Lastly the simulator defines a spleen compartment, which represents the spleen which is an important lymph system organ. These four compartments are notionally connected by the fifth, or circulation compartment. Most cell types can enter the circulation, but strict rules pertain as to which compartments they can migrate to from there.

Table B.1: Enumeration and dimensions of the compartments defined within the simulator. There are five main compartments defined within the simulator, the periphery is not explicitly modelled and is notionally included for the purposes of implementing immunization in the simulation.

Compartment	Length / grid points	Width / grid points
Circulation	62	40
SLO	50	50
CLN	50	50
Spleen	62	40
CNS	50	50

Figure B.4: Illustration of the inter-communication between the different simulator compartments. Each compartment represents a distinct location within the body. Annotations on the communicating arrows define which cell types are allowed to migrate between the specified compartments in the baseline simulator. This figure is adapted from [Read 2011].

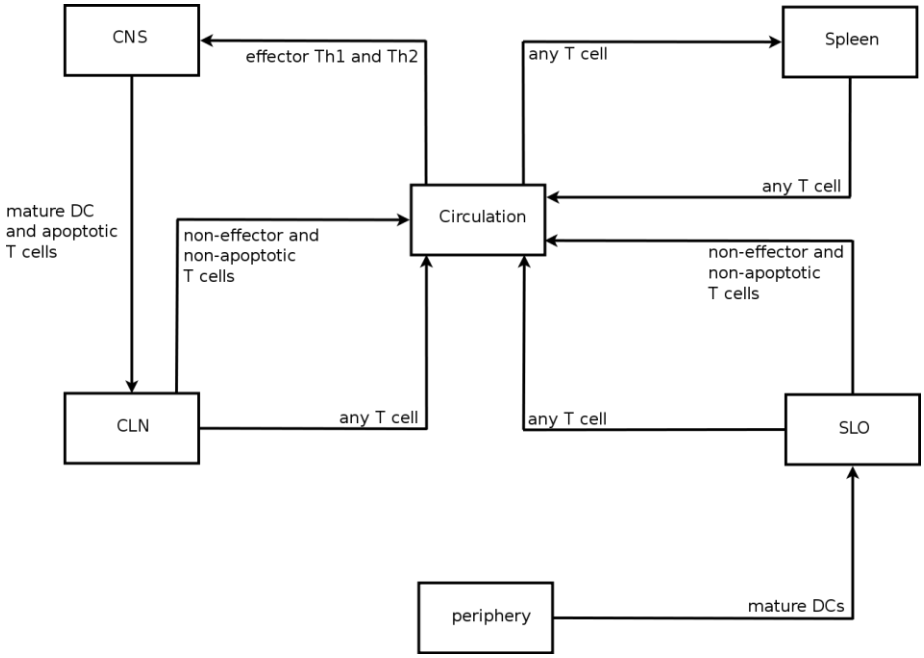

Cell migration occurs on the grid defined for each compartment. A cell can randomly choose a direction in which to move and, if there is space in the chosen destination cell for it to do so, the cell migrates. If a cell reaches the boundary of its current compartment, then the cell can choose whether or not to migrate into the circulation, or if already there, to one of the other compartments to which it has access.

B.5 The States Available to Dendritic Cells

DCs can be either Immature or Mature and a Mature DC can be Tolerogenic or Immunogenic depending on whether the cell can express costimulatory molecules or not. These states are summarised in Table B.2 below.

Table B.2: Description of the states that dendritic cells can adopt within the simulator. The states are defined in terms of whether the cell is expressing MHC compounds (MHC) or costimulatory molecules (CoStim) and whether the cell is Type 1 or Type 2 polarised.

State	Express MHC	Express CoStim	Polarization set
Immature	No	No	No
Mature	Yes	Yes	Yes
Tolerogenic	Yes	No	Yes
Immunogenic	Yes	Yes	Yes

B.6 The States Available to T-cells

Naïve cells are spawned bound to the same APC that their parent was bound to. The naïve cell has the same specificity as its parent.

When the naïve T-cell receives 'signal 1' that is it recognises its cognate antigen, the cell becomes partially activated and can move into a proliferative state if it receives 'signal 2' – costimulation.

The proliferating cell remains bound to the APC and a time is set at which the T-cell can become an effector.

Once the T-cell has become an effector times are set for cell death via Activation Induced Cell Death and for 'death by neglect'. The T-cell becomes unbound from the APC and is free to migrate. Effector function is not active until the T-cell encounters its cognate antigen.

Once the T-cell reaches a preset age, it becomes apoptotic, that is, it undergoes natural programmed cell death.

B.7 The Immunization and Initial Populations of Dendritic Cells

The simulator is designed to appear as if it has been running permanently. To this end there are two special types of DendriticCell agent within the simulator. There are 252 DC in all created at a simulation time designated -1.0. 120 of these (36 immature and 84 mature) are divided between the spleen, CLN and SLO and accounts for all the DC which are permanently resident in these compartments. 40 immature, migrating DC are placed in the CNS and 92 mature, migratory DC in the CLN.

Additionally, DC are introduced to the simulation as 'immunization DCs'. Initially 14 immunization DC are added to the SLO at time 1.0. Further mature DC are created over time according to the immunization protocol.

The above information was derived from simulator parameters and via personal communication with Dr Mark Read.

B.8 Data Logging, Storing and Output

Within the Simulator software there are packages specifically designed for logging simulation data, storing it and ultimately writing it to the file system.

The dataLoggers package contains classes for logging simulation data. Typically this is primarily cell population counts and the counting is performed automatically at every simulation time-step. However, in our experiments investigating the effects of Qa-1 expression timing and of competition between CD4 and CD8 Treg, we needed to log data on apoptotic DCs and so it made more sense to log the data from these cells as they became apoptotic.

The data logged by the dataLogger is stored in a dataStore object during the simulation and is then tabulated and output to the file system. The dataStore object has one record for each time step and one field within each record for each of the quantities that the simulator keeps track of.

Again, for the work in Chapter 3 of the thesis we implemented our own bespoke dataStore objects for recording the age of licensing of DCs and also the proportions of space around each DC occupied by the different T Cell sub-populations.

Appendix C: Specific Details of the Changes Implemented in the Simulator

C.1 Changes in Implementation for the Work in Chapter 3

The new data logger and data store objects were coded and added to the appropriate packages. The data loggers were instantiated as static pointers on the `DendriticCell` class, the instantiation taking place in the `TregSimulation` class. The data stores were instantiated and accessed in the `SingleRun` class in a manner analogous to the existing `SingleRunDataStore` object which holds the data that the simulator writes to the file system.

The stop clock data logger (`DCApoptosedEventsDataLogger`) was implemented so that it logged timing data only from DCs that were expressing Th1-derived peptides (namely the CDR1/2 and Fr3 fragments of the T-cell receptor Vβ8.2 chain). The spatial saturation data logger (`DCApoptosedNeighboursDataLogger`) recorded the proportion of the available space around DCs occupied by different types of T-cells. Specifically data were recorded on the proportion of binding space occupied by CD4 and CD8 Treg around DC expressing Th1-derived peptides, the proportion of space occupied by CD4 Th cells around DC expressing MBP and the proportion of space occupied by all three T-cell types around DC expressing MBP and Th1-derived peptides together.

All other coding changes took place within the `DendriticCell` class. These principally consisted of new methods used to record the cellular event timings on the DCs and relevant accessor ('getter') methods to permit the data loggers to access the private fields on the `DendriticCell`. Of paramount importance to the two experiments were the modifications made to the `becomeApoptotic()` method on `DendriticCell` as these enabled the `DendriticCell` objects to submit themselves to the data loggers when they became apoptotic.

A slight modification was made in the `DendriticCellMigrates` class to allow for full inheritance of methods implemented in `DendriticCell` class. This related to the `becomeApoptotic()` method from which the apoptotic `DendriticCell` objects submit themselves to the two data loggers.

The code was tested after major revisions. Testing consisted principally of tracing flow of control via diagnostic use of the `System.out.println()` method. This method was also used to report on values passed between methods on different objects and of key variables at decision points in the code.

C.2 Unusual Data in the Age of Licensing Distribution

While collecting the age at licensing data in Section 3.6.3, the simulation presented some unusual timing data. The age of licensing data logger was found to have logged DCs which had creation times of -1.0 hours, a time for Qa-1 licensing and no recorded time of maturity in several runs out of the 1000. These data were particularly unusual in that the age of licensing data logger during each simulator run typically logged between 120 and 150 Th1-derived peptide expressing DCs and this kind of cell had been logged only between one and three times per run. It was therefore important to verify that these data were legitimate observations and not an artefact of a software bug.

The simulator has been implemented in such a way as to make it appear as if it has been running forever. This is achieved by creating a ready established population of DCs across the simulator's compartments at time -1.0 hours. Some of these DCs are immature and some are mature, in which case they never invoke the `becomeNonImmature()` method which is where the time of maturity is recorded for each DC. All start the simulation able to express MBP but not Th1-derived peptides. As the mature DC created at time -1.0 hour then function as ordinary DCs it is perfectly feasible for these cells to have become licensed for Qa-1 expression and to be expressing CDR1/2 and Fr3 at the time of cell death. Hence the data logger can legitimately record the timing data pertaining to these DCs.

Similarly, several of the initial population of immature DCs were legitimately logged throughout the 1000 baseline runs. These cells had a creation time of -1.0 hours, but had a time of maturity and a time of licensing recorded for them.

C.3 Results from the Verification of Baseline Behaviour Conducted in Section 3.6.1

The baseline simulation carried out served to provide a test that the alterations to simulation logic had not caused any disturbances in baseline behaviour. It also served to provide us with a set of baseline data against which we could compare subsequent experiments by means of the A-Test.

A brief way to assess that we are returning baseline behaviour is to examine the system-wide T-cell effector population curves. The CD4 Th1 should peak at around day 15 (~1,000 cells), CD4 Th2 should peak around day 12 (~200cells). CD4 Treg peak around 30 days at slightly more than 1,000 cells and CD8 Treg peak around the same time with a lower population of ~600 cells. We find that this is indeed the case as is illustrated in Figure C1.

Figure C.1: The median system-wide T-cell effector populations across 1,000 runs of the simulation using our augmented simulation to verify that we had not disturbed baseline behaviour. The peaks in effector populations occurred at the anticipated times and at the correct population levels.

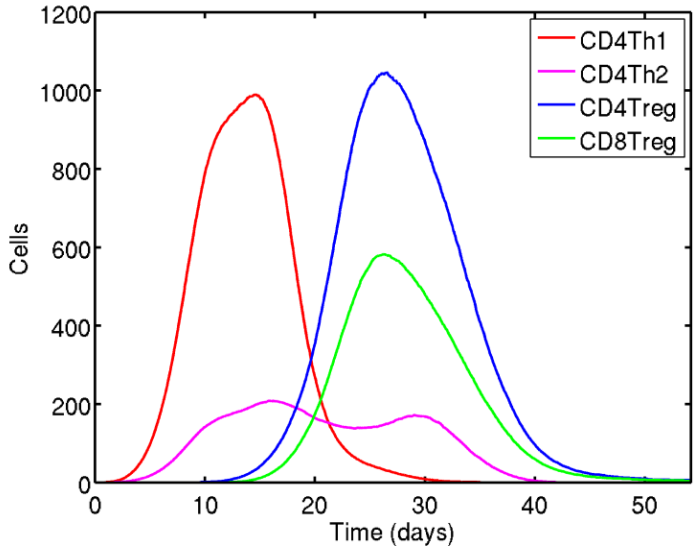

C.4 Results from the Verification of CD4 Treg Abrogation Behaviour Conducted in Section 3.6.2

The CD4 Treg abrogation simulation carried out served to provide a test that the alterations to simulation logic had not caused any disturbances in previously recorded behaviour [Williams 2010b]. It also served to provide us with a set of CD4 Treg abrogation simulation data against which we could compare subsequent experiments by means of the A-Test.

A brief way to assess that we are returning the proper behaviour is to examine the system-wide T-cell effector population curves. The CD4 Th1 should peak at around day 12 (~900 cells), CD4 Th2 should peak around day 12 (~200cells). CD8 Treg should peak at around 30 days ~1000 cells. We find that this is indeed the case as is illustrated in Figure C2.

Figure C.2: The median system-wide T-cell effector populations across 1,000 runs of the simulation using our augmented simulation to verify that we had not disturbed baseline behaviour. The peaks in effector populations occurred at the anticipated times and at the correct population levels.

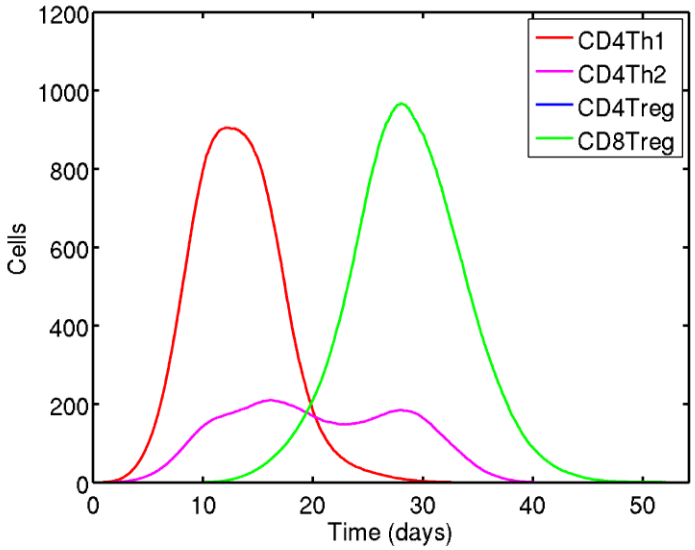

Appendix D: The A-Test

D.1 Effect Sizes for the A-Test

The A-Test is a non-parametric test of effect, used to test the significance of a change effected between two populations where the property of interest is not normally distributed [Vargha and Delaney 2000]. In the context of this thesis we have used the test for assessing the statistical magnitude of the effect of changing one parameter and holding all others constant between two simulator experiments e.g. with Qa-1 expression delayed and with constitutive Qa-1 expression.

The test returns a score which lies in the range 0.00 to 1.00, the score reflecting the statistical size of the observed effect. The effect sizes associated with the range of scores is detailed in Table D.1.

Table D.1: The effect sizes relating to the range of A-Test scores. The descriptions and score boundaries corresponding to them are justified in the paper by [Vargha and Delaney 2000].

A-Test Score	Effect Size
> 0.71	Large
> 0.64	Medium
> 0.56	Small
0.44-0.56	No Difference
> 0.36	Small
> 0.29	Medium
> 0.00	Large

Appendix E: Supporting Data for the CD200-CD200R Negative Signalling Factorial Analyses

E.1 The Graphs Omitted from the Discussion in Section 4.7

i) Effect on Time taken to reach peak CD4 Th1 effector population

Figure E.1: 3-dimensional plot illustrating the time taken to reach the peak CD4 Th1 effector population at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

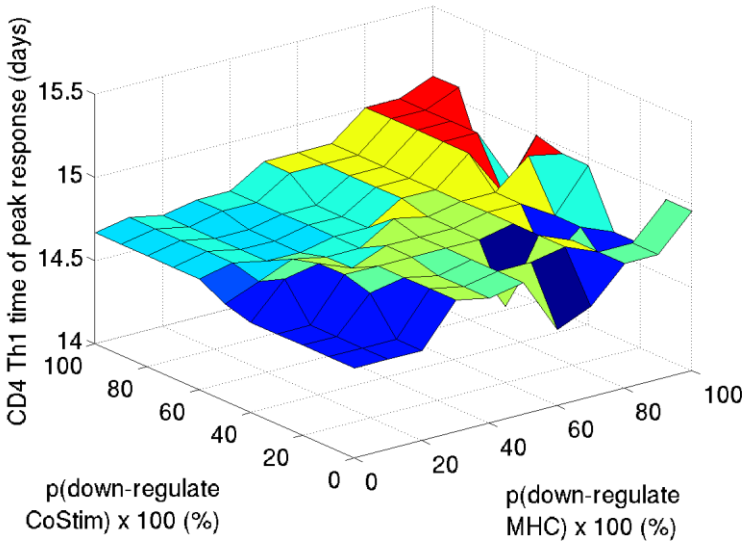

ii) Effect on Peak CD4 Th2 Effector Population

Figure E.2: 3-dimensional plot illustrating the peak CD4 Th2 effector population at all of the possible pairings of probability parameters between 0% and 1 The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

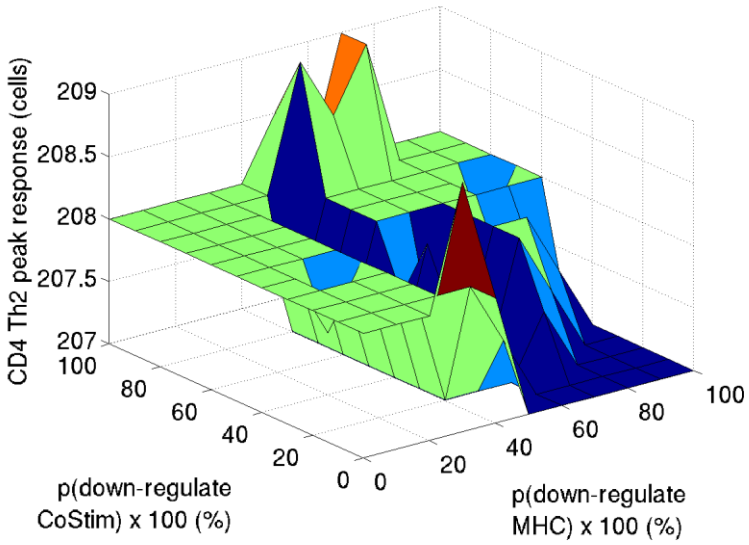

iii) Effect on time taken to reach peak CD4 Th2 Effector population

Figure E.3: 3-dimensional plot illustrating the time taken to reach the peak CD4 Th2 effector population at all of the possible pairings of probability parameters between 0% and 1%. The probability parameters represent the probability that a DC receiving a negative signal from CD200R down regulates MHC expression – p(down-regulate MHC) and the probability that a DC receiving a negative signal from CD200R down-regulates CoStim expression – p(down-regulate CoStim). The point 0, 0 that is nearest to the reader in the lower middle portion of the plot is the current simulator baseline behaviour. The probabilities on the x- and y-axes have been multiplied by 100 for clarity of labelling the axis ticks.

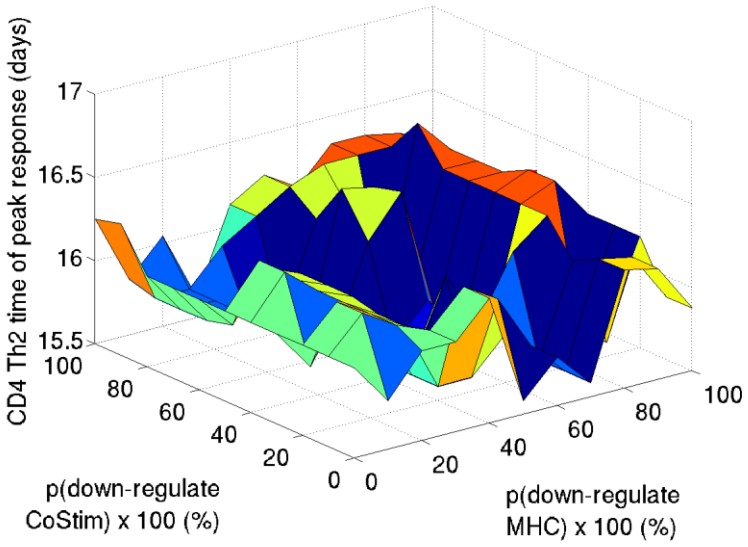

E.2 The Initial Parameter Mapping (0% to 100% in 10% Increments)

E.2.1 A-Test Scores: Comparing Experiments with Equal p(down-regulate CoStim) but Varying p(down-regulate MHC)

Table E.1: A-Test scores for comparison of experiments with p(down-regulate CoStim) held constant and p(down-regulate MHC) varied. p(down-regulate CoStim) is the probability that negative signalling (interaction of CD200 with CD200R) causes the DC bearing CD200R to down-regulate expression of costimulatory molecules. Similarly with p(down-regulate MHC) for MHC expression. The column headings labelled 'Max' indicate the A-Test score for comparing the peak cell population for the named cell type. Columns labelled 'MaxTime' are the scores for comparison of the times taken to reach the peak population. Th1@40 is the score for comparing the CD4 Th1 population at day 40 of the simulation. MaxEAE is the score for comparing the maximum EAE severity scores during the simulation and EAE@40d is the score for comparing EAE severity at day 40 between two experiments. (This Table is continued on the following page)

parameter value (%)	CD4Th1 Max	CD4Th1 MaxTime	CD4Th2 Max	CD4Th2 MaxTime	CD4Treg Max	CD4Treg MaxTime
0	0.500000	0.500000	0.500000	0.500000	0.500000	0.500000
10	0.507706	0.523136	0.419184	0.405426	0.000000	0.039386
20	0.507790	0.519696	0.380486	0.363170	0.000000	0.070310
30	0.518728	0.535204	0.373630	0.339494	0.000000	0.100902
40	0.524850	0.539938	0.362316	0.331794	0.000000	0.109296
50	0.525610	0.555570	0.364720	0.314024	0.000000	0.124714
60	0.531712	0.560296	0.360154	0.303970	0.000000	0.118954
70	0.528210	0.564018	0.366472	0.308928	0.000000	0.144788
80	0.530710	0.553380	0.361682	0.307078	0.000000	0.129618
90	0.534768	0.567352	0.358936	0.310726	0.000000	0.177884
100	0.532738	0.575280	0.359524	0.321834	0.000000	0.132368

Table E.1 (continued from the previous page): A-Test scores for comparison of experiments with p(down-regulate CoStim) held constant and p(down-regulate MHC) varied. p(down-regulate CoStim) is the probability that negative signalling (interaction of CD200 with CD200R) causes the DC bearing CD200R to down-regulate expression of costimulatory molecules. Similarly with p(down-regulate MHC) for MHC expression. The column headings labelled 'Max' indicate the A-Test score for comparing the peak cell population for the named cell type. Columns labelled 'MaxTime' are the scores for comparison of the times taken to reach the peak population. Th1@40 is the score for comparing the CD4 Th1 population at day 40 of the simulation. MaxEAE is the score for comparing the maximum EAE severity scores during the simulation and EAE@40d is the score for comparing EAE severity at day 40 between two experiments.

parameter value (%)	CD8Treg Max	CD8Treg MaxTime	Th1@40d	Max EAE	EAE@40d
0	0.500000	0.500000	0.500000	0.500000	0.500000
10	0.000000	0.033176	0.999836	0.503166	0.510408
20	0.000000	0.045978	0.999984	0.498864	0.521928
30	0.000000	0.052090	0.999998	0.522078	0.522484
40	0.000000	0.091714	1.000000	0.525404	0.519856
50	0.000000	0.060586	1.000000	0.524620	0.528376
60	0.000000	0.083204	1.000000	0.532588	0.534228
70	0.000000	0.097656	0.999960	0.516434	0.521524
80	0.000000	0.083382	1.000000	0.523108	0.542784
90	0.000000	0.079436	1.000000	0.540198	0.540228
100	0.000000	0.111352	1.000000	0.517506	0.532896

E.2.2 A-Test Scores: Comparing Experiments with Equal p(down-regulate MHC) but Varying p(down-regulate CoStim)

Table E.2: A-Test scores for comparison of experiments with p(down-regulate MHC) held constant and p(down-regulate CoStim) varied. p(down-regulate CoStim) is the probability that negative signalling (interaction of CD200 with CD200R) causes the DC bearing CD200R to down-regulate expression of costimulatory molecules. Similarly with p(down-regulate MHC) for MHC expression. The column headings labelled 'Max' indicate the A-Test score for comparing the peak cell population for the named cell type. Columns labelled 'MaxTime' are the scores for comparison of the times taken to reach the peak population. Th1@40 is the score for comparing the CD4 Th1 population at day 40 of the simulation. MaxEAE is the score for comparing the maximum EAE severity scores during the simulation and EAE@40d is the score for comparing EAE severity at day 40 between two experiments. (This Table is continued on the following page).

parameter value (%)	CD4Th1 Max	CD4Th1 MaxTime	CD4Th2 Max	CD4Th2 MaxTime	CD4Treg Max	CD4Treg MaxTime
0	0.500000	0.500000	0.500000	0.500000	0.500000	0.500000
10	0.501582	0.512196	0.501890	0.493508	0.424492	0.537170
20	0.501650	0.515530	0.502052	0.500864	0.391384	0.557110
30	0.509048	0.530504	0.486646	0.492014	0.348824	0.590214
40	0.511748	0.529004	0.502324	0.488778	0.315146	0.629780
50	0.512946	0.535958	0.492400	0.493912	0.295734	0.660070
60	0.513068	0.536054	0.482720	0.481930	0.267788	0.672842
70	0.517040	0.543850	0.485528	0.490268	0.271842	0.680516
80	0.515112	0.533634	0.474712	0.476094	0.251464	0.683260
90	0.517422	0.538278	0.492188	0.510296	0.248208	0.690682
100	0.517522	0.539400	0.493658	0.500876	0.237738	0.701848

Table E.2 (continued from the previous page): A-Test scores for comparison of experiments with p(down-regulate MHC) held constant and p(down-regulate CoStim) varied. p(down-regulate CoStim) is the probability that negative signalling (interaction of CD200 with CD200R) causes the DC bearing CD200R to down-regulate expression of costimulatory molecules. Similarly with p(down-regulate MHC) for MHC expression. The column headings labelled 'Max' indicate the A-Test score for comparing the peak cell population for the named cell type. Columns labelled 'MaxTime' are the scores for comparison of the times taken to reach the peak population. Th1@40 is the score for comparing the CD4 Th1 population at day 40 of the simulation. MaxEAE is the score for comparing the maximum EAE severity scores during the simulation and EAE@40d is the score for comparing EAE severity at day 40 between two experiments.

parameter value (%)	CD8Treg Max	CD8Treg MaxTime	Th1@40d	Max EAE	EAE@40d
0	0.500000	0.500000	0.500000	0.500000	0.500000
10	0.401226	0.550872	0.507656	0.476774	0.478000
20	0.355782	0.593948	0.512812	0.486340	0.493000
30	0.286896	0.627662	0.527040	0.481468	0.495000
40	0.248384	0.652224	0.520228	0.513286	0.506000
50	0.220766	0.680140	0.535756	0.505538	0.494000
60	0.191364	0.692632	0.540336	0.493636	0.494000
70	0.193134	0.731048	0.548118	0.514604	0.507000
80	0.172830	0.699364	0.556400	0.512942	0.490000
90	0.159002	0.714832	0.557060	0.510678	0.495000
100	0.135610	0.723236	0.566768	0.512384	0.499000

E.3 The Second Parameter Mapping (0% to 1% in 0.1% Increments)

E.3.1 A-Test Scores: Comparing Experiments with Equal p(down-regulate CoStim) but Varying p(down-regulate MHC)

Table E.3: A-Test scores for comparison of experiments with p(down-regulate CoStim) held constant and p(down-regulate MHC) varied. p(down-regulate CoStim) is the probability that negative signalling (interaction of CD200 with CD200R) causes the DC bearing CD200R to down-regulate expression of costimulatory molecules. Similarly with p(down-regulate MHC) for MHC expression. The column headings labelled 'Max' indicate the A-Test score for comparing the peak cell population for the named cell type. Columns labelled 'MaxTime' are the scores for comparison of the times taken to reach the peak population. Th1@40 is the score for comparing the CD4 Th1 population at day 40 of the simulation. MaxEAE is the score for comparing the maximum EAE severity scores during the simulation and EAE@40d is the score for comparing EAE severity at day 40 between two experiments. (This Table is continued on the following page).

parameter value x 100 (%)	CD4Th1 Max	CD4Th1 MaxTime	CD4Th2 Max	CD4Th2 MaxTime	CD4Treg Max	CD4Treg MaxTime
0	0.500000	0.500000	0.500000	0.500000	0.500000	0.500000
10	0.501932	0.501572	0.498926	0.498678	0.311550	0.356856
20	0.502022	0.505780	0.498808	0.481066	0.194218	0.275556
30	0.504628	0.509142	0.499680	0.488492	0.130162	0.223092
40	0.503012	0.503852	0.483568	0.476254	0.083034	0.198960
50	0.503568	0.506058	0.485768	0.483976	0.061702	0.161288
60	0.501308	0.503116	0.508066	0.492698	0.039374	0.140530
70	0.502022	0.503366	0.516686	0.510340	0.028486	0.131072
80	0.504006	0.504636	0.502488	0.487826	0.019028	0.118398
90	0.503484	0.505764	0.485170	0.483774	0.013740	0.110824
100	0.503478	0.509196	0.501560	0.507476	0.009842	0.096466

Table E.3 (continued from the previous page): A-Test scores for comparison of experiments with p(down-regulate CoStim) held constant and p(down-regulate MHC) varied. p(down-regulate CoStim) is the probability that negative signalling (interaction of CD200 with CD200R) causes the DC bearing CD200R to down-regulate expression of costimulatory molecules. Similarly with p(down-regulate MHC) for MHC expression. The column headings labelled 'Max' indicate the A-Test score for comparing the peak cell population for the named cell type. Columns labelled 'MaxTime' are the scores for comparison of the times taken to reach the peak population. Th1@40 is the score for comparing the CD4 Th1 population at day 40 of the simulation. MaxEAE is the score for comparing the maximum EAE severity scores during the simulation and EAE@40d is the score for comparing EAE severity at day 40 between two experiments.

parameter value x 100 (%)	CD8Treg Max	CD8Treg MaxTime	Th1@40d	Max EAE	EAE@40d
0	0.500000	0.500000	0.500000	0.500000	0.500000
10	0.311118	0.353784	0.526306	0.499672	0.499000
20	0.196198	0.275970	0.558474	0.494870	0.496000
30	0.114834	0.224868	0.610494	0.486004	0.489000
40	0.073854	0.203982	0.643398	0.484688	0.490000
50	0.048158	0.176650	0.669044	0.474960	0.486000
60	0.030080	0.149998	0.692798	0.478622	0.493000
70	0.019506	0.129692	0.761152	0.492220	0.500000
80	0.010858	0.120848	0.774034	0.496186	0.499000
90	0.008530	0.114460	0.802424	0.488274	0.497000
100	0.005308	0.098662	0.814916	0.490828	0.490000

E.3.2 A-Test Scores: Comparing Experiments with Equal p(down-regulate MHC) but Varying p(down-regulate CoStim)

Table E.4: A-Test scores for comparison of experiments with p(down-regulate MHC) held constant and p(down-regulate CoStim) varied. p(down-regulate CoStim) is the probability that negative signalling (interaction of CD200 with CD200R) causes the DC bearing CD200R to down-regulate expression of costimulatory molecules. Similarly with p(down-regulate MHC) for MHC expression. The column headings labelled 'Max' indicate the A-Test score for comparing the peak cell population for the named cell type. Columns labelled 'MaxTime' are the scores for comparison of the times taken to reach the peak population. Th1@40 is the score for comparing the CD4 Th1 population at day 40 of the simulation. MaxEAE is the score for comparing the maximum EAE severity scores during the simulation and EAE@40d is the score for comparing EAE severity at day 40 between two experiments. (This Table is continued on the following page).

parameter value x 100 (%)	CD4Th1 Max	CD4Th1 MaxTime	CD4Th2 Max	CD4Th2 MaxTime	CD4Treg Max	CD4Treg MaxTime
0	0.500000	0.500000	0.500000	0.500000	0.500000	0.500000
10	0.500216	0.499680	0.503292	0.509718	0.502804	0.483094
20	0.499716	0.499364	0.505904	0.501896	0.499714	0.468330
30	0.500782	0.501306	0.510912	0.498490	0.511886	0.484136
40	0.501036	0.500334	0.525940	0.506372	0.510466	0.494934
50	0.501516	0.503122	0.519098	0.510134	0.507666	0.499596
60	0.501836	0.502730	0.520352	0.515270	0.503068	0.476712
70	0.502634	0.503924	0.513628	0.528814	0.510000	0.506094
80	0.502702	0.504368	0.517840	0.513676	0.493594	0.492716
90	0.503356	0.504402	0.506354	0.506230	0.493154	0.487664
100	0.503376	0.504322	0.529254	0.524930	0.487680	0.487432

Table E.4 (continued from the previous page): A-Test scores for comparison of experiments with p(down-regulate MHC) held constant and p(down-regulate CoStim) varied. p(down-regulate CoStim) is the probability that negative signalling (interaction of CD200 with CD200R) causes the DC bearing CD200R to down-regulate expression of costimulatory molecules. Similarly with p(down-regulate MHC) for MHC expression. The column headings labelled 'Max' indicate the A-Test score for comparing the peak cell population for the named cell type. Columns labelled 'MaxTime' are the scores for comparison of the times taken to reach the peak population. Th1@40 is the score for comparing the CD4 Th1 population at day 40 of the simulation. MaxEAE is the score for comparing the maximum EAE severity scores during the simulation and EAE@40d is the score for comparing EAE severity at day 40 between two experiments.

parameter value x 100 (%)	CD8Treg Max	CD8Treg MaxTime	Th1@40d	Max EAE	EAE@40d
0	0.500000	0.500000	0.500000	0.500000	0.500000
10	0.492524	0.499834	0.504778	0.506592	0.504000
20	0.493992	0.496610	0.481120	0.505646	0.502000
30	0.500508	0.480476	0.498954	0.499270	0.498000
40	0.511922	0.503724	0.492432	0.496632	0.496000
50	0.501148	0.496266	0.486122	0.497326	0.495000
60	0.511442	0.501090	0.487026	0.492840	0.491000
70	0.510000	0.518940	0.486270	0.492914	0.495000
80	0.502756	0.505886	0.498776	0.494954	0.491000
90	0.496330	0.494158	0.489290	0.488580	0.490000
100	0.492420	0.507452	0.486926	0.490320	0.487000

Abbreviations Used Within the Thesis

Immunological and Biological:

MS	Multiple Sclerosis
EAE	Experimental Autoimmune (or Allergic) Encephalomyelitis
CNS	Central Nervous System
BBB	Blood Brain Barrier
SLO	Secondary Lymphoid Organ
CLN	Cervical Lymph Node
DC	Dendritic Cell
APC	Antigen Presenting Cell
Treg	Regulatory T-cell
IFN	Interferon
TNF	Tumour Necrosis Factor

Computational:

ABM(S)	Agent Based Modelling (and Simulation)
ABM	Agent Based Model
UML	Unified Modelling Language
CA	Cellular Automaton
ODE	Ordinary Differential Equation
MASON	Multi-Agent Simulation of Neighbourhoods / Networks
FLAME	Flexible Large-scale Agent-based Modelling Environment
CoSMoS	**Co**mplex **S**ystems **Mo**delling and **S**imulation

Bibliography

[Alberts *et al.* 2008] B. Alberts, D. Bray, J. Lewis, M. Raff, K. Roberts and J.D. Watson, *Molecular Biology of the Cell. 5th Edition*, New York, Abingdon: Garland Science, 2008.

[An 2008] G. An, "Introduction of an agent-based multi-scale modular architecture for dynamic knowledge representation of acute inflammation." *Theoretical Biology and Medical Modelling*, vol. 5, pp. 1-30, 2008.

[Andrews *et al.* 2008] P.S. Andrews, F. Polack, A.T. Sampson, J. Timmis, L. Scott and M. Coles, "Simulating biology: toward understanding what the simulation shows." in *Workshop on Complex System Modelling and Simulation*, Luniver Press, pp. 92-123, 2008.

[Andrews *et al.* 2010] P.S. Andrews, F.A.C. Polack, A.T. Sampson, S. Stepney and J. Timmis, "CoSMoS Process: A Process for the Modelling and Simulation of Complex Systems," University of York, Technical Report YCS-2010-453, 2010.

[Baldazzi *et al.* 2006] V. Baldazzi, F. Castiglione and M. Bernaschi, "An enhanced agent-based model of the immune system response." *Cellular Immunology*, 244, pp. 77-79, 2006.

[Baron *et al.* 1993] J.L. Baron, J.A. Madri, N.H. Ruddle, G. Hashim and C.A. Janeway Jr, "Surface Expression of α4 Integrin by CD4 T Cells is Required for Their Entry into Brain Parenchyma." *Journal of Experimental Medicine*, vol. 177, pp. 57-68, 1993.

[Bauer *et al.* 2009] A.L. Bauer, C.A.A. Beauchemin and A.S. Perelson, "Agent-based modeling of host-pathogen systems: The successes and challenges." *Information Sciences*, vol. 179, pp. 1379-1389, 2009.

[Beeston *et al.* 2010] T. Beeston, T.R.F. Smith, I. Maricic, X. Tang and V. Kumar, "Involvement of IFN-γ and perforin, but not Fas/FasL interactions in regulatory T cell-mediated suppression of experimental autoimmune encephalomyelitis." *Journal of Neuroimmunology*, vol. 229, pp. 91-97, 2010.

[Berg *et al.* 2007] J.M. Berg, J.L. Tymoczko and L. Stryer, *Biochemistry*. 6th Edition, New York: W.H. Freeman and Co., 2007.

[Bernaschi and Castiglione 2001] M. Bernaschi and F. Castiglione, "Design and implementation of an immune system simulator." *Computers in Biology and Medicine*, vol. 31, pp. 303-331, 2001.

[Booch *et al.* 2005] G. Booch, J. Rumbaugh and I. Jacobson, *The Unified Modelling Language User Guide.* Reading, Massachusetts: Addison-Wesley Longman, 2005.

[Bonabeau 2002] E. Bonabeau, "Agent-based modelling: Methods and techniques for modelling human systems." *Proceedings of the National Academy of Science USA*, vol. 99, pp. 7280-7287, 2002.

[Brode and Macary 2004] S. Brode and P.A. Macary, "Cross-presentation: dendritic cells and macrophages bite off more than they can chew!" *Immunology*, vol. 112, pp. 345-351, 2004.

[Cameron *et al.* 2005] C.M. Cameron, J.W. Barrett, L. Liu, A.R. Lucas and G. McFadden, "Myxoma Virus M141R Expresses a Viral CD200 (vOX-2) That Is Responsible for Down-Regulation of Macrophage and T-Cell Activation *In Vivo*." *Journal of Virology*, vol. 79, pp. 6052-6067, 2005.

[Cantor *et al.* 1978] H. Cantor, J. Hugenberger, L. McVay-Boudreau, D.D. Eardley, J. Kemp, F.W. Shen and R.K. Gershon, "Immunoregulatory Circuits Among T cell Sets: Identification of a Subpopulation of T-Helper Cells that Induces Feedback Inhibition." *Journal of Experimental Medicine*, vol. 148, pp. 871-877, 1978.

[Carson 2002] M.J. Carson, "Microglia as Liaisons Between the Immune and Central Nervous Systems: Functional Implications for Multiple Sclerosis." *GLIA*, vol. 40, pp. 218-231, 2002.

[Charcot 1877] J.M. Charcot, *Lecture Notes on the Diseases of the Nervous System.* London: The Sydenham Society, 1877. (English translation)

[Chataway 1989] S.J.S. Chataway, "What's new in the Pathogenesis of Multiple Sclerosis? A Review." *Journal of the Royal Society of Medicine*, vol. 82, pp. 159-162, 1989.

[Chitnis *et al.* 2007] T. Chitnis, J. Imitola, Y. Wang, W. Elyaman, P. Chawla, M. Sharuk, K. Radassi, R.T. Bronson and S.J. Khoury, "Elevated Neuronal Expression of CD200 Protects Wld^S Mice From Inflammation-Mediated Neurodegeneration." *The American Journal of Pathology*, vol. 170, pp. 1695-1712, 2007.

[Claudio *et al.* 1990] L. Claudio, Y. Kress, J. Factor and C.F. Brosnan, "Mechanisms of Edema Formation in Experimental Autoimmune Encephalomyelitis: The Contribution of Inflammatory Cells." *American Journal of Pathology*, vol. 137, pp. 1033-1045, 1990.

[Cohen 2007] I.R. Cohen, "Modelling immune behaviour for experimentalists", *Immunological Reviews*, vol. 216, pp. 232-236, 2007.

[Collier *et al.* 2003] N. Collier, T. Howe and M.J. North, "Onward and Upward: The transition to Repast 2.0." presented at the First Annual North American Association for Computational Social and Organizational Science (Pittsburgh, PA USA). North American Association for Computational Social and Organizational Science, 2003.

[Copland *et al.* 2007] D.A. Copland, C.J. Calder, B.J.E. Raveney, L.B. Nicholson, J. Phillips, H. Cherwinski, N. Jenmalm, J.D. Sedgwick and A.D. Dick, "Monoclonal Antibody-Mediated CD200 Receptor Signaling Suppresses Macrophage Activation and Tissue Damage in Experimental Autoimmune Uveoretinitis." *The American Journal of Pathology*, vol. 171, pp. 580-588, 2007.

[Crawford and Wherry 2009] A. Crawford and E.J. Wherry, "The diversity of costimulatory and inhibitory receptor pathways and the regulation of antiviral T cell responses." *Current Opinion in Immunology*, vol. 21, pp. 179-186, 2009.

[Dancik *et al.* 2010] G.M. Dancik, D.E. Jones and K.S. Dorman, "Parameter estimation and sensitivity analysis in an gent-based model of *Leishmania major* infection." *Journal of Theoretical Biology*, vol. 262, pp. 398-412, 2010.

[Dorland's 2010] *Dorland's Illustrated Medical Dictionary 30th Ed*, Philadelphia: W.B. Saunders & Company, 2010.

[Efroni *et al.* 2005] S. Efroni, D. Harel and I.R. Cohen, "Reactive Animation: Realistic Modeling of Complex Dynamic Systems." *Computer*, vol. 38, pp. 38-47, 2005

[Efroni *et al.* 2007] S. Efroni, D. Harel and I.R. Cohen, "Emergent Dynamics of Thymocyte Development and Lineage Determination" *PloS Computational Biology*, vol. 3(1) e.13, doi:10.1371/journal.pcbi.0030013, 2007.

[Endharti *et al.* 2010] A.T. Endharti, Y. Okuno, Z. Shi, N. Misawa, S. Toyokuni, M. Ito, K. Isobe and H. Suzuki, "$CD8^+CD122^+$ Regulatory T Cells (Treg) and $CD4^+$ Treg Cooperatively Prevent and Cure $CD4^+$ Cell-Induced Colitis." *Journal of Immunology*, vol. 186, pp. 41-52, 2010.

[Feuer 2007] R. Feuer, "Tickling the CD200 Receptor: A Remedy For Those Irritating Macrophages." *The American Journal of Pathology*, vol . 171, pp. 396-398, 2007.

[Forrest and Beauchemin 2007] S. Forrest and C. Beauchemin, "Computer immunology." *Immunological Reviews*, vol. 216, pp. 176-197, 2007.

[Garnett *et al.* 2008] P. Garnett, S. Stepney and O. Leyser, "Towards an Executable Model of Auxin Transport Canalisation." In: S. Stepney, F.A.C. Polack and P. Welch (Eds.), *CoSMoS*, pp. 63-91, Luniver Press, 2008.

[Garrett and Grisham 2004] R.H. Garrett and C.M Grisham, *Biochemistry, 3rd Ed*, Belmont, CA: Thomson Brooks/Cole, 2004.

[Germain 2001] R.N. Germain, "The Art of the Probable: System Control in the Adaptive Immune System." *Science*, vol. 293, pp. 240-245, 2001.

[Germain *et al.* 2011] R.N. Germain, M. Meier-Schellersheim, A. Nita-Lazar and I.D.C. Fraser, "Systems Biology in Immunology: A Computational Modelling Perspective." *Annual Reviews in Immunology*, vol. 29, pp. 527-585, 2011.

[Gorczynski *et al.* 2004a] R.M. Gorczynski, Z. Chen, Y. Kai, L. Lee, S. Wong and P.A. Marsden, "CD200 is a Ligand for All Members of the CD200R Family of Immunoregulatory Molecules." *Journal of Immunology*, vol. 172, pp. 7744-7749, 2004.

[Gorczynski *et al.* 2004b] R.M. Gorczynski, Z. Chen, D.A. Clark, Y. Kai, L. Lee, J. Nachman, S. Wong and P. Marsden, "Structural and Functional Heterogeneity in the CD200R Family of Immunoregulatory Molecules and their Expression at the Feto-maternal Interface." *American Journal of Reproductive Immunology*, vol. 52, pp. 147-163, 2004.

[Hafler 2004] D.A. Hafler, "Multiple Sclerosis." *Journal of Clinical Investigation*, vol. 113, pp. 788-794, 2004.

[Harel 1987] D. Harel, "Statecharts: A Visual Formalism for Complex Systems," *Science of Computer Programming*, vol. 8, pp. 231-274, 1987.

[Hoek *et al.* 2000] R.M. Hoek, S.R. Ruuls, C.A. Murphy, G.J. Wright, R. Goddard, S.M. Zurawski, B. Blom, M.E. Homola, W.J. Streit, M.H. Brown, A.N. Barclay and J.D. Sedgwick, "Down-Regulation of the Macrophage Lineage Through Interaction With OX2 (CD200)." *Science*, vol. 290, pp. 1768-1771, 2000.

[Holcombe 2006] M. Holcombe, "From molecules to insect communities – how formal agent-based computational modelling is uncovering new biological facts." *Scientiae Mathematicae Japonicae Online* e-2006, pp. 765-778, 2006.

[Hone 2009] A. Hone, "On Non-Standard Numerical Integration Methods for Biological Oscillators." In: S. Stepney, P.H. Welch, P.S. Andrews and J. Timmis (Eds.), *CoSMoS*, pp. 45-66, Luniver Press, 2009.

[Huitinga *et al.* 1990] I. Huitinga, N. van Rooijen, C.J.A. De Groot, B.M.J. Uitdehaag and C.D. Dijkstra, "Suppression of Experimental Allergic Encephalomyelitis in Lewis Rats After Elimination of Macrophages." *Journal of Experimental Medicine*, vol. 172, pp. 1025-1033, 1990.

[Janeway *et al.* 2008] C.A. Janeway, P. Travers, M. Walport and M.J. Shlomchick, *Immunobiology: the Immune System in Health and Disease 7th Ed*, Oxford: Garland Science, 2008.

[Kam *et al.* 2001] N. Kam, I.R. Cohen and D. Harel, "The Immune System as a Reactive System:Modeling T Cell Activation With Statecharts." Presented at the Symposium on Human Centric Computing Languages and Environments (HCC), pp. 15-22, IEEE, 2001.

[Keir and Sharpe 2005] M.E. Keir and A.H. Sharpe, "The B7/CD28 costimulatory family in autoimmunity." *Immunological Reviews*, vol. 204, pp. 128-143, 2005.

[Kholodenko 2006] B.N. Kholodenko, "Cell-signalling dynamics in time and space." *Nature Reviews Molecular Cell Biology*, vol. 7, pp. 165-176, 2006.

[Kindt *et al.* 2007] T.J. Kindt, B.A. Osborne and R.A. Goldsby, R.A., *Kuby Immunology 6th Ed*, New York: W.H. Freeman, 2007.

[Kitano 2001] H. Kitano, "Systems Biology: Toward System-level Understanding of Biological Systems." in: H. Kitano(ed), *Foundations of Systems Biology*, pp. 1-36, MIT Press (Cambridge, MA and London, England), 2001.

[Kitano 2002a] H. Kitano, "Computational systems biology." *Nature*, vol. 420, pp. 206-210, 2002.

[Kitano 2002b] H. Kitano, "Systems Biology: A Brief Overview." *Science*, vol. 295, pp. 1662-1664, 2002.

[Kitano 2004] H. Kitano, "Biological Robustness." *Nature Reviews Genetics*, vol. 5, pp. 826-837, 2004.

[Kleinstein and Celada 2000] S.H. Kleinstein and P.E. Seiden, "Simulating the Immune System." *Computing in Science and Engineering*, vol. 2, pp. 69-77, 2000.

[Kohm *et al.* 2002] A.P. Kohm, P.A. Carpentier, H.A. Anger and S.D. Miller, "Cutting Edge $CD4^+CD25^+$ Regulatory T Cells Suppress Antigen-Specific Autoreactive Immune Responses and Central Nervous System Inflammation During Active Experimental Autoimmune Encephalomyelitis." *Journal of Immunology*, vol. 169, pp. 4712-4716, 2002.

[Kumar and Sercarz 2001] V. Kumar and E. Sercarz, "An Integrative Model of Regulation Centred on Recognition of TCR Peptide / MHC Complexes." *Immunological Reviews*, vol. 182, pp. 113-121, 2001.

[Kumar *et al.* 1996] V. Kumar, K. Stellrecht and E. Sercarz, "Inactivation of T Cell Receptor Peptide-specific CD4 Regulatory T Cells Induces Chronic Experimental Autoimmune Encephalomyelitis (EAE)." *Journal of Experimental Medicine*, vol. 184, pp. 1609-1617, 1996.

[Kurts *et al.* 2010] C. Kurts, B. W. S. Robinson and P.A. Knolle, "Cross-priming in health and disease." *Nature Reviews Immunology*, vol. 10, pp. 403-414, 2010.

[Larman 2005] C Larman, *Applying UML and Patterns: An Introduction to Object-Oriented Analysis and Design and Iterative Development.* Westford, Massachusetts: Prentice Hall, 2005.

[Liu *et al.* 2010] Y. Liu, Y. Bando, D. Vargas-Lowy, W. Elyaman, S.J. Khoury, T. Huang, K. Reif and T. Chitnis, "CD200R1 Agonist Attenuates Mechanisms of Chronic Disease in a Murine Model of Multiple Sclerosis." *Journal of Neuroscience*, vol. 30, pp. 2025-2038, 2010.

[Lodish and Darnell 1995] H. Lodish and J.E. Darnell, *Molecular Cell Biology*. 3rd Edition, New York: Scientific American Books, W.H. Freeman and Co., 1995.

[Lublin 1985] F.D. Lublin, "Relapsing Experimental Allergic Encephalomyelitis an Autoimmune Model of Multiple Sclerosis." *Springer Seminars in Immunopathology*, vol. 8, pp. 197-208, 1985.

[Ludowyk *et al.* 1992] P.A. Ludowyk, D.O. Willenborg and C.R. Parish, "Selective localisation of neuro-specific T lymphocytes in the central nervous system." *Journal of Neuroimmunology*, vol. 37, pp. 237-250, 1992.

[Lue *et al.* 2010] L.-F. Lue, Y.-M. Kuo, T. Beach and D.G. Walker, "Microglia Activation and Anti-Inflammatory Regulation in Alzheimer's Disease." *Molecular Neurobiology*, vol. 41, pp. 115-128, 2010.

[Luke *et al.* 2003] S. Luke, G.C. Balan, L. Panait, C. Cioffi-Revilla and S. Paus, "MASON: A Java Multi-Agent Simulation Library." In *Proceedings of the Agent 2003 Conference*, 2003.

[Luke *et al.* 2004] S. Luke, C. Cioffi-Revilla, L. Panait and K. Sullivan, "MASON: A New Multi-Agent Simulation Toolkit." In *Proceedings of the 2004 SwarmFest Workshop*, 2004.

[Luke *et al.* 2005] S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan, "MASON: A Multi-Agent Simulation Environment" in *Simulation: Transactions of the Society for Modeling and Simulation International*, vol. 82(7), pp. 517-527, 2005.

[Macal and North 2005] C.M. Macal and M.J. North, "Tutorial on Agent-based modelling and simulation." Proceedings of the 2005 Winter Simulation Conference, M.E. Kuhl, N.M. Steiger, F.B. Armstrong, J.A. Joines (Eds.)

[Mann and Whitney 1947] H.B. Mann and D.R.Whitney, "On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other." *The Annals of Mathematical Statistics*, vol. 18, pp. 50-60, 1947.

[Mckay *et al.* 2000] M. D. Mckay, R. J. Beckman and W. J. Conover, "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code." *Technometrics*, vol. 42(1) Special 40th Anniversary Issue (Feb., 2000), pp. 55-61, 2000.

[McCombe *et al.* 1994] P.A. McCombe, J. de Jersey and M.P. Pender, "Inflammatory cells, microglia and MHC class-II antigen positive cells in the spinal cord of Lewis rats with acute and chronic relapsing experimental autoimmune encephalomyelitis." *Journal of Neuroimmunology*, vol. 51, pp. 153-167, 1994.

[McDonald 1974] W.I. McDonald, "Pathophysiology in Multiple Sclerosis." *Brain*, vol. 97, pp. 179-196, 1974.

[Melchior *et al.* 2006] B. Melchior, S.S. Puntambekar and M.J. Carson, "Microglia and the control of autoreactive T cell responses." *Neurochemistry International*, vol. 49, pp. 145-153, 2006.

[Mellouli *et al.* 2004] S. Mellouli, B. Moulin and G. Mineau, "Laying Down the Foundations of an Agent Modelling Methodology for Fault-Tolerant Multi-agent Systems." In: A. Omicini, P. Petta and J. Pitt, (Eds), 4th International Workshop of Engineering Societies in the Agents World (ESAW 2003), Lecture Notes in Computer Science, vol. 3071, pages 275-293, Springer, 2004.

[Meuth *et al.* 2008] S.G. Meuth, O.J. Simon, A. Grimm, N. Melzer, A.M. Herrmann, P. Spitzer, P. Landgraf and H. Wiendl, "CNS inflammation and neuronal degeneration is aggravated by impaired CD200-CD200R-mediated macrophage silencing." *Journal of Neuroimmunology*, vol. 194, pp. 62-69, 2008.

[Minar *et al.* 1996] N. Minar, R. Burkhart, C. Langton and M. Askenazi, "The Swarm Simulation System: a toolkit for building multi-agent simulations." Working paper 96-06-042, Sante Fe Institute, Santa Fe, 1996.

[Minas and Liversidge 2006] K. Minas and J. Liversidge, "Is The CD200/CD200 Receptor Interaction More Than Just a Myeloid Cell Inhibitory Signal?" *Critical Reviews in Immunology*, vol. 26, pp. 213-230, 2006.

[Mostarica-Stojkovic *et al.* 1992] M. Mostarica-Stojkovic, S. Vukmanovic, Z. Ramic and M.L. Lukic, "Evidence for target tissue regulation of resistance to the induction of experimental autoimmune encephalomyelitis in AO rats." *Journal of Neuroimmunology*, vol. 41, pp. 97-104, 1992.

[Mustafa *et al.* 1991] M. I. Mustafa, P. Diener, B. Höjeberg, P. van der Meide and T. Olsson, "T cell immunity and interferon-γ secretion during experimental autoimmune encephalomyelitis in Lewis rats." *Journal of Neuroimmunology*, vol. 31, pp. 165-177, 1991.

[OMG] Object Modelling Group (OMG) (undated) available at: http://www.omg.org

[Odell *et al.* 2002] J.J. Odell, H, van Dyke Paranuk, M. Fleischer and S. Brückner, "Modeling Agents and Their Environment." in: *Proceedings of the 3rd International Conference on Agent-oriented software engineering III*, July 15th, 2002, Bologna, Italy.

[Orabona *et al.* 2004] C. Orabona, U. Grohmann, M.L. Belladonna, F. Fallarino, C. Vacca, R. Bianchi, S. Bozza, C. Volpi, B.L. Salomon, M.C. Fioretti, L. Romani and P. Puccetti, "CD28 induces immunostimulatory signals in dendritic cells via CD80 and CD86." *Nature Immunology*, vol. 5, pp. 1134-1142, 2004.

[Pender 1995] M.P. Pender, "Experimental Autoimmune Encephalomyelitis." in *Autoimmune Neurological Disease*, M.P. Pender and P.A. McCombe Eds., Cambridge: Cambridge University Press, 1995, pp 26-88.

[Perelson 2002] A.S. Perelson, "Modelling viral and immune system dynamics." *Nature Reviews Immunology*, vol. 2, pp. 28-36, 2002.

[Perelson and Nelson 1999] A.S. Perelson and P.W. Nelson, "Mathematical Analysis of HIV-2 Dynamics *in vivo*," *Society for Industrial and Applied Mathematics Review*, vol. 41, pp. 3-44, 1999.

[Playfair and Lydyard 1995] J.H.L. Playfair and P.M. Lydyard, *Medical Immunology for Students.* Edinburgh: Churchill Livingstone, 1995.

[Pogson *et al.* 2008] M. Pogson, M. Holcombe, R. Smallwood and E. Qwarnstrom, "Introducing Spatial Information into Predictive NF-κB Modelling – an Agent-based Approach." PLoS ONE 3(6): e2367. doi:10.1371/journal.pone.0002367, 2008.

[Pohjonen and Kelly 2002] R. Pohjonen and S. Kelly, "Domain Specific Modelling" *Dr Dobbs Journal*, vol. 339, pp. 26-35, 2002. http://drdobbs.com Last accessed: 24th August 2011.

[Polack *et al.* 2010] F.A.C. Polack, P.S. Andrews, T. Ghetiu, M. Read, S. Stepney, J. Timmis and A.T. Sampson, "Reflections on the Simulation of Complex Systems for Science." Presented at the 15th IEEE International Conference on Engineering of Complex Computer Systems, 2010.

[Read *et al.* 2009a] M. Read, J. Timmis, P.S. Andrews and V. Kumar, "A Domain Model of Experimental Autoimmune Encephalomyelitis." In: S. Stepney, P.H. Welch, P.S. Andrews and J. Timmis (Eds.), *CoSMoS*, pp. 9-44, Luniver Press, 2009.

[Read *et al.* 2009b] M. Read, J. Timmis, P.S Andrews and V. Kumar, "Using UML to Model EAE and its Regulatory Network," In: P.S. Andrews, J. Timmis, N.D.L. Owens, U. Aickelin, E. Hart, A. Hone and A.M. Tyrrell eds., *ICARIS, vol 5666, Lecture Notes in Computer Science*, pp 4-6, Springer, 2009.

[Read *et al.* 2011] M. Read, P. Andrews, J. Timmis, J. and V. Kumar, "Towards Qualifying the Implications of Epistemic Uncertainty on Simulation-Based Experimentation through Calibration, Uncertainty and Sensitivity Analysis." *Mathematical and Computer Modelling of Dynamical Systems*, (in print).

[Read 2011] M. Read, "Statistical and modelling techniques to investigate immunology through agent-based simulation." PhD Thesis, Department of Computer Science, University of York, 2011.

[Reboldi *et al.* 2009] A. Reboldi, C. Coisne, D. Baumjohann, F. Benvenuto, D. Bottinelli, S. Lira, A. Uccelli, A. Lanzavecchia, B. Engelhardt and F. Salusto, "CC chemokine receptor 6 – regulated entry of TH-17 cells into the CNS through the choroid plexus is required for the initiation of EAE." *Nature Immunology*, vol. 10, pp. 514-523, 2009.

[Redmond and Sherman 2005] W.L. Redmond and L.A. Sherman, "Peripheral Tolerance of T Lymphocytes." *Immunity*, vol. 22, pp. 275-284, 2005.

[Rivers *et al.* 1933] T.M. Rivers, D.H Sprunt and P. Berry, (1933), "Observations on Attempts to Produce Acute Disseminated Encephalomyelitis in Monkeys." *Journal of Experimental Medicine*, vol. 58, pp. 39-53, 1933.

[Roberston 1981] M. Robertson, "Nerves, Myelin and Multiple Sclerosis." *Nature*, vol. 290, pp. 357-358, 1981.

[Santoni *et al.* 2008] D. Santoni, M. Pedicini and F. Castiglione, "Implementation of a regulatory gene network to simulate the TH1/2 differentiation in an agent-based model of hypersensitivity reactions." *Bioinformatics*, vol. 24, pp. 1374-1380, 2008.

[Schwartz 2003] R.H. Schwartz, "T Cell Anergy." *Annual Reviews in Immunology*, vol. 21, pp. 305-334, 2003.

[Segovia-Juarez *et al.* 2004] J.L. Segovia-Juarez, S. Ganguli and D. Kirschner, "Identifying control mechanisms of granuloma formation during *M. tuberculosis* infection using an agent-based model." *Journal of Theoretical Biology*, vol. 231, pp. 357-376, 2004.

[Seiden and Celada 1992] P.E. Seiden and F. Celada, "A Model for Simulating Cognate Recognition and Response in the Immune System." *Journal of Theoretical Biology*, vol. 158, pp. 329-357, 1992.

[Shevach 2006] E.M. Shevach, "From Vanilla to 28 Flavours: Multiple Varieties of T Regulatory Cells." *Immunity*, vol. 25, pp. 195-201, 2006.

[Smith and Kumar 2008a] T.R.F. Smith and V. Kumar, "Immune Suppression by a Novel Population of CD8aa+TCRab+ Regulatory T Cells." In: *Regulatory T Cells and Clinical Application*, D. Jiang (Ed.), New York: Springer, pp. 1-14, 2008.

[Smith and Kumar 2008b] T.R.F. Smith and V. Kumar, "Revival of CD8+ Treg-mediated Suppression." *Trends in Immunology*, vol. 9, pp. 337-342, 2008b.

[Steiner and Wirguin 2000] I. Steiner and I. Wirguin, "Multiple Sclerosis – in need of a critical reappraisal." *Medical Hypotheses,* vol. 54, pp. 99-106, 2000.

[Tang *et al.* 2005] X. Tang, T.R.F. Smith and V. Kumar, "Specific Control of Immunity by Regulatory CD8 T Cells." *Cellular and Molecular Immunology,* vol. 2, pp. 11-19, 2005.

[Tang *et al.* 2006] X. Tang, I. Maricic, N. Purohit, B. Bakamjian, LM Reed-Loisel, T. Beeston, P. Jensen and V. Kumar, "Regulation of immunity by a novel population of Qa-1 restricted CD8aa+TCRab+ T cells." *Journal of Immunology,* vol. 177, pp. 7645-7655, 2006.

[Thorne *et al.* 2007] B.C. Thorne, A.M. Bailey, S.M. Peirce, "Combining experiments with multi-cell agent-based modeling to study biological tissue patterning." *Briefings in Bioinformatics,* vol. 8, pp. 245-257, 2007.

[Tisue and Wilensky 2004] S. Tisue and U. Wilensky, "NetLogo: A Simple Environment For Modeling Complexity." presented at the International Conference on Complex Systems, Boston, 2004.

[van den Bark *et al.* 1985] A.A. van den Bark, T. Gill and H. Offner, "A myelin basic protein-specific T lymphocyte line that mediates experimental autoimmune encephalomyelitis." *Journal of Immunology,* vol. 135, pp. 223-228, 1985.

[Vargha and Delaney 2000] A. Vargha and H.D. Delaney, "A Critique and Improvement of the CL Common Language Effect Size Statistics of McGraw and Wong." *Journal of Educational and Behavioural Statistics,* vol. 25, pp. 101-132, 2000.

[Voet and Voet 2004] D. Voet and J. Voet, *Biochemistry 3rd Edition* New York:John Wiley and Sons Inc., 2004.

[Walker and Southgate 2009] D.C. Walker and J. Southgate, "The virtual cell – a candidate co-ordinator for 'middle-out' modelling of biological systems." *Briefings in Bioinformatics,* vol. 10, pp. 450-461, 2009.

[Walker *et al.* 2004] D.C. Walker, J. Southgate, G. Hill, M. Holcombe, D.R. Hose, S.M. Wood, S. Mac Neil and R.H. Smallwood, "The Epitheliome: Agent-based Modelling of the Social Behaviour of Cells." *BioSystems,* vol. 76, pp. 89-100, 2004.

[Webb and White 2004] K. Webb and T. White, "Cell Modelling Using Agent-Based Formalisms." In: B. Orchard, C. Yang and A. Moonis (Eds.), 17th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (IEA/AIE), *Lecture Notes in Computer Science,* vol. 3029, pp. 128-137, 2004.

[Webb and White 2005] K. Webb and T. White, "UML as a Cell and Biochemistry Modelling Language." *BioSystems*, vol. 80, pp. 283-302, 2005.

[Williams 2010a] R.A. Williams, "Modelling the Glycolytic and Citrate Pathways using the Unified Modelling Language." Department of Biology, University of York, Masters Degree Project Report, 2010.

[Williams 2010b] R.A. Williams, "*In Silico* Experimentation using Simulation of Experimental Autoimmune Encephalomyelitis." Department of Computer Science, University of York, Masters Degree Project Report, 2010.

[Williams *et al.* 2011] R.A. Williams, M. Read, J. Timmis, P.S. Andrews and V. Kumar, "*In Silico* Investigation into CD8 Treg Mediated Recovery in Murine Experimental Autoimmune Encephalomyelitis." Presented at ICARIS 2011, *Lecture Notes in Computer Science*, vol. 6825, pp. 51-54, 2011.

[Wishart *et al.* 2004] D.S. Wishart, R. Yang, D. Arndt, P. Tang and J. Cruz, "Dynamic cellular automata: a simple but powerful approach to cellular simulation." *In Silico Biology*, vol. 5, 0015, 2004.

[Wright *et al.* 2003] G.J. Wright, H. Cherwinski, M. Foster-Cuevas, G. Brooke, M.J. Puklavec, M. Bigler, Y. Song, M. Jenmalm, D. Gorman, T. McClanahan, M.-R. Liu, M.H. Brown, J.D. Sedgwick, J.H. Phillips and A.N. Barclay, "Characterization of the CD200 Receptor Family in Mice and Humans and Their Interactions With CD200." *Journal of Immunology*, vol. 171, pp. 3034-3046, 2003.

[Zozulya *et al.* 2009] A.L. Zozulya, S. Ortler, Z. Fabry, M. Sandor and H. Wiendl, "The level of B7 homologue 1 expression on brain DC is decisive for CD8 Treg cell recruitment into the CNS during EAE." *European Journal of Immunology*, vol. 39, pp. 1536-1543, 2009.