A stochastic model of cytotoxic T cell responses

Dennis L. Chao[a], Miles P. Davenport[b], Stephanie Forrest[a,c], Alan S. Perelson[d,*]

[a] Department of Computer Science, University of New Mexico, Albuquerque, NM 87131, USA
[b] Department of Haematology, Prince of Wales Hospital and Centre for Vascular Research, University of New South Wales, Kensington NSW 2052, Australia
[c] The Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA
[d] Theoretical Biology and Biophysics, Los Alamos National Laboratory, T10 MS-K710, Los Alamos, NM 87545, USA

Abstract

We have constructed a stochastic stage-structured model of the cytotoxic T lymphocyte (CTL) response to antigen and the maintenance of immunological memory. The model follows the dynamics of a viral infection and the stimulation, proliferation, and differentiation of naïve $CD8^+$ T cells into effector CTL, which can eliminate virally infected cells. The model is capable of following the dynamics of multiple T cell clones, each with a T cell receptor represented by a digit string. MHC–viral peptide complexes are also represented by strings and a string match rule is used to compute the affinity of a T cell receptor for a viral epitope. The avidities of interactions are also computed by taking into consideration the density of MHC–viral peptides on the surface of an infected cell. Lastly, the model allows the probability of T cell stimulation to depend on avidity but also incorporates the notion of an antigen-independent programmed proliferative response. We compare the model to experimental data on the cytotoxic T cell response to lymphocytic choriomeningitis virus infections.

1. Introduction

Cytotoxic T lymphocytes (CTL) play a crucial role in the immune system's defense against viral infections. After the first exposure to a virus, T cells rapidly replicate and attack infected cells in a *primary response*. The T cell population then decreases, leaving behind a population of long-lived memory cells. These memory cells allow the immune system to respond to subsequent encounters with similar infections more efficiently in a *secondary response*. A lifetime of exposure to pathogens shapes an organism's repertoire of memory cells, making the states of individual immune systems unique.

Based on a variety of experimental results, we constructed a stochastic model of the CTL response to antigens. Our ultimate goal is to use the model to provide insight into the pathology and possible treatments of diseases such as AIDS, influenza, cancer, and autoimmune disorders. In the process of building the model we were able to situate experimental data from multiple experiments in a coherent framework that forms a relatively complete and consistent interpretation of T cell behavior. We take a computationally efficient stage-structured modeling approach which allows us to incorporate biologically realistic features of T cell proliferation and differentiation relatively easily, resulting in a model that makes quantitative predictions. In the sections that follow, we summarize relevant CTL biology, describe our model, then present preliminary results.

2. T cell biology

CTL reside in tissue or circulate through the body via the blood and lymph to detect cells that have been compromised by foreign organisms, such as viruses. We present a summary of the CTL biology relevant to our model. Many essential components of the immune response, such as the innate immune system, dendritic

---

*Corresponding author. Tel.: +1-505-667-5061; fax: +1-505-665-3493.

E-mail address: asp@lanl.gov (A.S. Perelson).

cells, and CD4$^+$ T cells, are intentionally omitted; their roles in facilitating the CTL response are implicit in the model.

2.1. T cell receptors and repertoire

Most of the body's cells process a sample of their internal proteins into short peptide fragments that form complexes with cell surface proteins called major histocompatibility (MHC) class I molecules. There are hundreds of MHC class I alleles in humans (Marsh et al., 2002), and an individual can express as many as six of them. Each MHC type binds a particular set of peptides and is thus capable of presenting a different set of epitopes than other MHCs. The external presentation of cell peptides allows T cells to view a sample of a cell's contents non-invasively. When a cytotoxic T cell binds to peptide–MHC complexes, it can initiate a set of actions that leads to the destruction of the infected cell.

One of the primary factors that determines whether a T cell binds to a cell is the *affinity* of its T cell receptor (TCR) for the peptide–MHC complexes. Each T cell expresses thousands of copies of identical receptors that bind to their cognate peptide–MHC complexes with high affinity. Thus, both the target cell peptides and the particular MHC type that presents the peptide play a role in determining affinity. The set of all TCR specificities in a body, on the order of 10^7 in humans (Arstila et al., 1999) and perhaps 10^6 in mice (Pannetier et al., 1993), comprise the T cell repertoire. *Avidity*, or the sum of the binding interactions between the receptors of a CTL and the surface of a target cell, determines whether a CTL recognizes the target. The number of copies of a particular peptide displayed by a target cell, its expression density, affects the avidity of the interaction. Due to thymic selection, described below, it is unlikely that a T cell will react to an uninfected self cell—infected cells express foreign (e.g., virally encoded) peptides that make them subject to T cell responses. The antigenic peptides that stimulate T cells are known as *epitopes*.

T cell receptors are generated with more or less random specificities, so many potentially harmful self-reactive ones are created. Most are screened out early in their maturation process in the thymus, where they are exposed to a large array of the body's own peptides presented on MHC molecules. During *positive selection*, T cells that have an extremely weak avidity to self-peptides bound to MHC are eliminated (Blackman et al., 1990). This process eliminates T cells that are useless because they probably cannot bind to any peptide–MHC pairs. *Negative selection* ensures that they do not bind too tightly to MHC–self-peptides (Kappler et al., 1987), eliminating T cells that are potentially self-reactive. About 1–3% of pre-selection T cells pass both

these "tests" and leave the thymus to join the peripheral repertoire as naïve T cells (Shortman et al., 1990).

2.2. T cell response

A naïve T cell remains quiescent until it receives antigenic stimulation from its cognate peptide–MHC complex. Larger antigen doses stimulate a greater fraction of naïve cells (perhaps more low affinity T cells) but probably do not affect the degree to which the individual cells are stimulated (Kaech and Ahmed, 2001). After stimulation, naïve cells appear to be committed to a *programmed response* that causes them to divide and acquire effector functions even in the absence of continuing antigenic stimulation (Kaech and Ahmed, 2001; van Stipdonk et al., 2001). For the first 24 h, they do not replicate (Oehen and Brduscha-Riem, 1998; Gett and Hodgkin, 2000; Veiga-Fernandes et al., 2000; van Stipdonk et al., 2001), but after this initial phase, they can rapidly undergo a fixed number of divisions (up to 8 or more) (Kaech and Ahmed, 2001) once every 5–8 h (Murali-Krishna et al., 1998; Gett and Hodgkin, 2000; van Stipdonk et al., 2001). After a few divisions, they acquire effector functions, such as cytotoxicity. Effector CTL kill target cells either by releasing perforins that create holes in the target cell's membrane or by triggering apoptosis (i.e. cell suicide) in the target cell. Even during this period of rapid expansion, the cells have a high death rate, reducing net population growth. After initial expansion, the death rate dominates CTL kinetics and the population declines rapidly (Badovinac et al., 2002). If the infection persists, the remaining T cells can become ineffective and unresponsive to antigen.

2.3. T cell memory

After the activation and proliferation in response to an infection, most of the T cells activated in the response die, but a small subpopulation persists as *memory cells* (Murali-Krishna et al., 1999). Memory cells are able to mount a quicker and more aggressive response in future encounters with the same or closely related pathogens (Dutton et al., 1998). This *secondary response* can clear an infection before significant damage is inflicted upon the body. Immunological memory forms the basis of vaccination, in which an organism is exposed to viral antigens in order to build immune memory to the virus.

The repertoire of memory cells generated by a primary response appears to be similar to the repertoire of the effectors from the response (Sourdive et al., 1998; Busch et al., 1998a; Blattman et al., 2000), indicating that the conversion of the effectors of the primary response to memory cells is independent of their avidity to the antigen. However, their recruitment into a secondary response *is* dependent on their avidity for

antigen, so the secondary response can have a different clonal composition from the primary (Bousso et al., 2000). It can take 2 or 3 weeks for a cell to develop a full memory cell phenotype after the initial infection (Kaech et al., 2002). Therefore, memory cells are not likely to join the immune response that initially generated them.

Upon antigenic stimulation, memory cells begin to proliferate almost immediately and develop cytotoxicity within a few hours (Bachmann et al., 1999; Barber et al., 2003; Byers et al., 2003). They probably have the same antigenic stimulation requirements as naïve cells (Bachmann et al., 1999; Kersh et al., 2003), although some studies found stimulation requirements to be lower (Pihlgren et al., 1996). Their replication rates are approximately the same as recently activated naïve cells. Memory cell-derived effectors die at a slower rate than effectors created in the primary response (Veiga-Fernandes et al., 2000; Grayson et al., 2002), giving them a faster accumulation rate and possibly allowing a larger portion of them to revert to memory. Presumably, the shorter time to acquire effector functions, the larger starting populations, and their faster accumulation rates allow memory cells to clear infected cells much faster than naïve cells.

Homeostatic mechanisms appear to regulate the size of the memory pool, which remains approximately constant in size (Rocha et al., 1989). New memory cells from heterologous infections can displace the memory cells from responses to prior infections (Selin et al., 1996; McNally et al., 2001). In the absence of immune system challenges, memory cells turn over slowly (Dutton et al., 1998; Murali-Krishna et al., 1999).

2.4. Lymphocytic choriomeningitis virus: a model pathogen

Much of what we know about the kinetics of cytotoxic T cell responses in vivo comes from studies of mice infected with lymphocytic choriomeningitis virus (LCMV) (Armstrong and Lillie, 1934; Traub, 1935). LCMV stimulates a well-characterized CTL-mediated immune response, and infection is generally non-lethal and asymptomatic. Although we calibrate the model developed below with mouse data, the model parameters can easily be changed to human-derived parameters as needed. Running our model with mouse parameters allows us to compare our results with the numerous published experimental results based on LCMV.

3. The model

Our model consists of two interacting parts: a stage-structured model of the T cell activation, proliferation and differentiation and a model of viral infection. The models are coupled in that infected cells stimulate
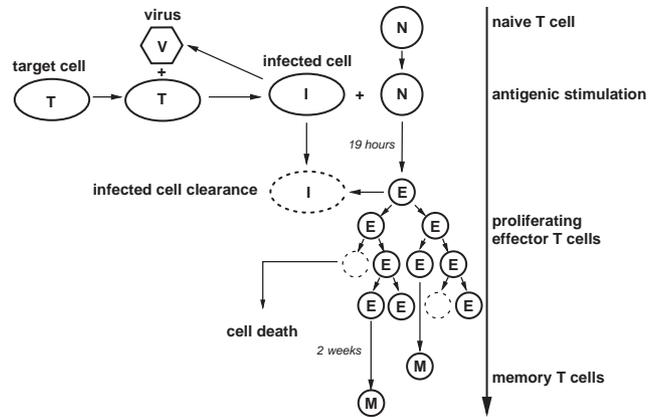

Fig. 1. The process of infection and the life cycle of CTL in our model. Target cells are infected by virus, and these infected cells generate more virus and interact with T cells. Naïve cells, when stimulated by antigen proliferate and become effector cells. The probability of a naïve cell being stimulated by antigen depends on the string distance between the TCR and the antigen–MHC complex. Most effectors die, but about 5% of these proliferating effector cells become memory cells. The memory cells can be stimulated to become effectors in a secondary response (not shown).

naïve T cells and are killed by effector T cells (depicted in Fig. 1). In addition, our model includes a representation of TCR binding and a realistic-sized T cell repertoire.

3.1. Approaches to immunological modeling

Differential equation models have long been used for immune system and viral infection modeling (Bell, 1970; Dibrov et al., 1977; Přikrylová et al., 1992; Perelson and Weisbuch, 1992; Ho et al., 1995; Nowak and Bangham, 1996; Bocharov, 1998; Perelson, 2002). Analytical techniques allow modelers to define regimes of system behavior and their associated parameters and initial conditions. For example, one can determine the model parameters for which an infection is effectively cleared by the immune system (Bocharov, 1998). The solutions capture the average behavior of large populations of perfectly mixed, identical individuals. Many techniques that could make these models more faithful to biological reality, such as adding time delays or age-structured partial differential equations (Antia et al., 2003), complicate solving the models analytically or numerically.

Agent-based simulation is a promising technique made feasible with the advent of greater computer power. These simulations monitor the actions of a large number of simple entities, or agents, in order to observe their aggregate behavior. Each agent consists of state variables and a set of rules that governs its behavior, and agents can interact either directly with each other or indirectly through the environment. Because all individuals in a population are explicitly represented, they can

have unique histories and behaviors. The combined behavior of these agents is observed in a discrete-time or event-driven simulation.

Agent-based modeling has many features suited to modeling the immune response. It is adept at incorporating stochastic events, which appear to be crucial in regulating immune function (Germain, 2001). A single chance event, such as the serendipitous recognition of a cancer antigen by a single cell in the immune system, can determine the fate of an organism (Ochsenbein et al., 2001). The addition of randomness to a model allows one to explore the distribution of possible outcomes, as in Detours and Perelson (2000), as opposed to only the single most likely one addressed by most mathematical models. This is especially valuable when studying immune responses, as even genetically identical individuals can exhibit different responses to the same antigen (Lin and Welsh, 1998). Because small numbers of cells are involved in the beginning of an immune response (Ehl et al., 1998; Bousso et al., 1999), using a discrete model might be more suitable in this context than a continuous one. The existing agent-based models of the immune system, such as IMMSIM (Celada and Seiden, 1992; Seiden and Celada, 1992; Kleinstein and Seiden, 2000), the B cell model of Smith et al. (Smith et al., 1999), and the self-nonself discrimination model of Langman and Cohn (Cohn et al., 2002; Langman et al., 2003), take advantage of these features. Another advantage of agent-based models is that by explicitly representing individual cells, they are in many ways closer to the modeled system. In contrast to population-level models, agent-based model parameters

correspond to actual properties of the cells, and the output of these models can be processed so that they can be observed at any level, from the level of the individual cell to the population level (Table 1).

Unfortunately, agent-based modeling can be computationally expensive. There may be as many as 10^{12} T cells in a human and 10^8 T cells in a mouse. Running a model with this many distinct entities would be prohibitive. To address such problems, agent-based models can be implemented to take advantage of multiple computers, such as the parallel version of IMMSIM known as ParImm (Bernaschi and Castiglione, 2001). Because agent-based models must be run many times to characterize the distribution of outcomes, they should be as simple and efficient as possible without sacrificing essential aspects of the immune response. Using techniques such as lazy evaluation (Smith et al., 1998) allows models to instantiate only cells that participate in an immune response, but a response in a mouse can involve on the order of 10^7 CTL (Butz and Bevan, 1998; Murali-Krishna et al., 1998), which is still a large number to simulate explicitly.

For computational efficiency, we use a stochastic stage-structured approach to modeling the immune response (Chao et al., 2003). Stage-structured models have been used to model populations in ecology (Lefkovitch, 1965; Usher, 1966; Manly, 1990) but have not been applied to immune systems. In stage-structured models, an individual's or cell's life cycle is divided into stages, such as developmental maturity or differentiation states. All individuals in a given stage are assumed to be identical. The transition probabilities

Table 1
A summary of model parameters

Attribute	Value	Reference
Time step (Δt)	10 min	
Naïve cell clone size	10 cells	Casrouge et al. (2000)
Maximum T cell recruitment rate (γ)	1 day^{-1}	
Delay before a stimulated naïve cell becomes an effector (τ_n)	19 h	Oehen and Brduscha-Riem (1998); Gett and Hodgkin (2000); Veiga-Fernandes et al. (2000); van Stipdonk et al. (2001)
Delay before a stimulated memory cell becomes an effector (τ_m)	1 h	Bachmann et al. (1999); Barber et al. (2003)
Naïve-derived active CTL death rate (δ_E)	0.6 day^{-1}	Veiga-Fernandes et al. (2000)
Memory-derived active CTL death rate (δ_{E_m})	0.4 day^{-1}	Veiga-Fernandes et al. (2000)
Time in B phase for CTL	5 h	van Stipdonk et al. (2001)
Average CTL cell cycle time	6 h	van Stipdonk et al. (2001)
Infected cell clearance rate (k^c)	12 day^{-1}	Barchet et al. (2000)
Pre-selection repertoire size	2.5×10^8	
MHC string length	4 digits	Detours et al. (1999)
Peptide string length	6 digits	Detours et al. (1999)
Susceptible cell population (T)	10^6 cells	
Susceptible cell production rate (λ)	10^5 cells/day	
Susceptible cell death rate (δ_T)	0.1 day^{-1}	
Virus infection rate (β)	2×10^{-7}	
Virus production rate (π)	100 day^{-1}	
Virus clearance rate (c)	2.3 day^{-1}	Lehmann-Grube (1988)
Infected cell death rate (δ_I)	0.8 day^{-1}	

between stages are specified, and the model can be used to predict the demographics of a given initial population over time. Stochasticity can be added to the model if needed. Analytical techniques have been developed for studying these models, but when there are interacting populations, in our case T cells and antigens, it is easier to run the model on a computer multiple times. In order to allow T cells to interact with virally infected cells, we run the T cell model and the viral model synchronously and at each time step the populations can interact. For example, at some stages T cells can eliminate infected cells, so at each time step the number of infected cells is reduced by a function of the number of T cells that are in the effector stages.

By using discrete rather than continuous populations and by explicitly specifying the actions and transitions of cells as probabilities per individual cell, our model enforces the realistic behavior of individual cells without the computational cost of representing each cell explicitly. The model attempts to strike a balance between the unrealistically small number of populations used by analytical approaches and the unwieldy one-agent-per-cell implementations of agent-based models. Because we do not intend to solve our system analytically, the model can accommodate multiple cell states. However, to make the model more efficient than an equivalent agent-based model, we must limit the number of possible cell states to a manageable number (described in Section 3.6).

3.2. Receptor binding

Our model uses digit strings and string match rules to represent antigens and affinities, as several immunological models have in the past (Farmer et al., 1986; Celada and Seiden, 1992; Detours et al., 1999; Smith et al., 1999; Bernaschi and Castiglione, 2001). Each antigen in the simulation is associated with one or more digit strings representing its epitopes. The strings loosely represent amino acid sequences. Each epitope also has a scalar value representing its surface expression density, ρ, which affects avidity. We assume that all cells infected with the same pathogen have the same epitope densities, making them all equally antigenic.

Strings also represent the portion of MHC visible to the TCR (Fig. 2). Each MHC allele is represented by a different random string. Self-peptides are represented by 10^3–10^4 random strings of the same length as the antigenic epitope strings. Each antigenic and self-peptide string is assigned a single MHC allele to which it is concatenated, representing the assumption that each peptide is presented by only a single MHC type in the body. We use the length of the strings chosen in Detours et al. (1999): 4 digits for the MHC and 6 digits for the peptides.
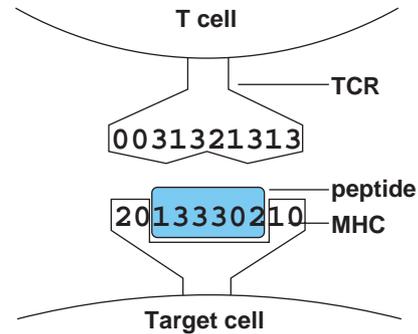

Fig. 2. A digit string representation of receptor binding. Each T cell receptor is represented by a digit string of fixed length. The binding affinity between a T cell receptor is determined by the relationship between this string and the string formed by concatenating the strings associated with a peptide associated with a target cell's peptide and its presenting MHC molecule.

Each CTL clone has a string representing its TCR that is as long as an MHC and a peptide string combined. A string match rule defines the binding affinity between the TCR and the MHC–peptide complexes. The avidity of a TCR for an epitope is a function of its affinity, determined by the string match rule, and the epitope's density on a cell surface. Thus, a high-affinity match can result in low avidity if that particular epitope's density is low. The string distance between newly generated TCR and the MHC–self-peptide complexes determines which T cells survive thymic selection. The string distance between TCR and MHC-foreign peptide complexes determines the behavior of T cells in the periphery. The choice of string match rule determines the affinity distribution of TCR clones from antigens and the effect mutating digits has on affinity.

We have explored a variety of rules that have different properties that may affect the behavior of our model, particularly with regard to the cross-reactivity among responses to different antigens and the mutation of epitopes. For example, the Hamming distance rule, which has been used to represent B cell–ligand binding (Smith et al., 1997), defines the distance between two strings to be the number of positions in which the two strings differ (Hamming, 1950). If there is a single digit "mutation" in one of the strings, then the distance will either increase or decrease by 1 or remain the same. One can interpret this to mean that each position in the TCR string either binds or does not bind to the MHC–peptide complex. The Manhattan distance rule, also known as the first-order Minkowski norm (L_1), defines the distance between two strings to be the sum of distances between their respective digits, with, for example, the distance between 8 and 5 being 3. Therefore, a single digit mutation can change the distance between two strings by as much as α, where α is the size of the alphabet. In other words, each position in the TCR

string can have α degrees of affinity for each corresponding location in the MHC–peptide complex. The xor rule, in which the distance is the sum of the bitwise exclusive ors of the digits of two strings, also allows for a large range of degrees of affinity between two digits but is more computationally efficient than the Manhattan rule (Detours et al., 1999). For all of these distance rules, it is assumed that the binding strengths between corresponding digits are additive and independent of the other digits, which agrees with the observation that amino acid side chains of peptides seem to make independent contributions to the binding energy with the TCR (Parker et al., 1994). It is difficult to determine which rule is most appropriate a priori, and different data sets might require different match rules.

To generate a T cell repertoire, TCR strings are first randomly generated uniformly over the universe of strings. These TCR are subjected to thymic selection (Fig. 3), in which T cells are eliminated if their receptors are too far from all MHC-self-complexes or too close to any MHC-self-complex, representing positive and negative selection respectively. Therefore, the space of TCR sequences is not uniformly populated. The surviving TCR comprise the naïve repertoire.

In order to model the response to a set of pathogens, we do not need to instantiate the whole repertoire of T cell clones, but only the small fraction of cells involved in the immune responses to the pathogens, a technique known as lazy evaluation (Smith et al., 1998). The clones that have no affinity for the antigens (i.e., those that are too antigenically distant and outside the "cross-reactive cutoff") play no role in the response. Therefore, the model considers only the cells that can react to the antigen, which reduces the number of T cells by several orders of magnitude. For each simulation run, random TCR strings are generated within the cross-reactive

cutoff from the antigenic strings to form the pre-selection repertoire, then these strings are subject to the simulated thymic selection process described above. In the past, modelers used artificially small repertoires (Kleinstein and Seiden, 2000), had each agent in the simulation represent more than one cell (Smith et al., 1999), or required substantial computing resources to simulate a realistic-sized repertoire (Detours and Perelson, 2000). In our model, all of the active cells of a realistic-sized repertoire are represented.

3.3. Viral infection

We adopt a standard model of viral infection previously used to describe human immunodeficiency virus (HIV) and hepatitis C virus (HCV) dynamics (Wei et al., 1995; Perelson et al., 1996; Neumann et al., 1998). In the absence of an immune response, the course of a viral infection is described by the following:

$$\dot{T} = \lambda - \delta_T T - \beta TV, \tag{1}$$

$$\dot{I} = \beta TV - \delta_I I, \tag{2}$$

$$\dot{V} = \pi I - cV \tag{3}$$

where T is the number of uninfected (or "target") cells, I is the number of infected cells, V is the number of virus particles, λ is the rate of uninfected cell production, π is the rate of virus production by infected cells, β is the infectivity parameter, δ_T is the death rate for target cells, δ_I is the death rate for infected cells, and c is the clearance rate for free virus. Typically, after infection the viral load and the number of infected cells increases exponentially, peaks, and then declines. Section 3.5 describes how effector T cells clear infected cells I in the model.

In our implementation, we use a difference equation version of the system of ODEs described by Eqs. (1)–(3):

$$\Delta T = (\lambda - \delta_T T - \beta TV)\Delta t, \tag{4}$$

$$\Delta I = (\beta TV - \delta_I I)\Delta t, \tag{5}$$

$$\Delta V = (\pi I - cV)\Delta t, \tag{6}$$

where $\Delta t = 10$ min. In order to include stochasticity, the terms in Eqs. (4)–(6) are randomly drawn from the appropriate distributions at each time step. We assume that the variables are constant over the short interval Δt and are updated at the end of each time step. For the production of uninfected cells and the virus production rate, we assume that they are governed by Poisson processes, and we draw from the Poisson distribution with their expected values as the mean (i.e., $\lambda \Delta t$ and $\pi I \Delta t$, respectively).

To determine stochastically the number of cells out of a population of identical cells that perform a certain action, such as dying, we randomly draw from the
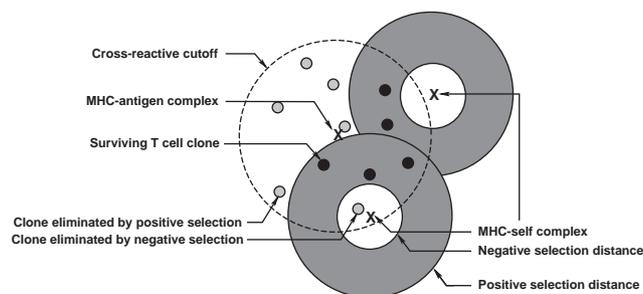

Fig. 3. A representation of thymic selection in a two-dimensional "shape space" (Perelson and Oster, 1979). TCR and antigens are mapped to a plane such that the distance between TCR and antigens is proportional to the affinity between them. TCR are generated around each antigen, and each randomly generated TCR is subjected to positive and negative selection against a set of MHC–self-peptide complexes. All TCR that are too far from all MHC–self-complexes are eliminated in positive selection. Those that are too close to an MHC–self-complex are eliminated in negative selection. The remaining TCR are in a window of affinities from an MHC–self-peptide complex.

binomial distribution. In order to do this, we must convert continuous rates into probabilities that events occur in a time step. If a process occurs at rate r, then the probability that it first occurs at time t is defined by the exponential distribution $\mathscr{E}(r) = re^{-rt}$. The probability that it occurs at or before time t is $1 - e^{-rt}$. Thus, rates r can be converted to probabilities that the processes occur in a time step Δt, $1 - e^{-r\Delta t}$. If there are n cells each with a probability p of performing an action, then drawing from the binomial $\mathscr{B}(n, p)$ is a computationally efficient way to determine the number of cells that perform the action. For example, we compute the number of uninfected cells T that are infected in each time step by converting their infection rate, βV, to the probability that they will become infected in a time step, $1 - e^{-\beta V\Delta t}$, and randomly drawing a value from $\mathscr{B}(T, 1 - e^{-\beta V\Delta t})$. We compared the results of our viral infection model implementation to an alternate version using Gillespie's Direct Method (Gillespie, 1977), which is an exact stochastic simulation technique that explicitly generates all reactions rather than computing how many reactions occur in a given time step. The different implementations produced distributions of outcomes, but our method was orders of magnitude faster.

3.4. Effector recruitment from the naïve and memory cell pools

Antigen stimulates naïve and memory cells, causing them to differentiate into effectors. Because a relatively small number of naïve cells are recruited into an immune response, we assume that they do not compete with each other for antigen, allowing the stimulation of each naïve clone to be computed independently. Stimulation takes the form of a saturating function (De Boer et al., 2001; Davenport et al., 2002):

$$\text{Stimulation} = \frac{\sum e_i I_i / K_i}{1 + \sum e_i I_i / K_i}, \tag{7}$$

where K_i is the amount of antigen i required to generate half-maximal stimulation for the T cell, e_i is epitope density on cells infected by antigen i, and I_i is the number of infected cells expressing antigen i in the system. This expression is in agreement with the observation that CTL recruitment is proportional to epitope density (Wherry et al., 1999). We assume that naïve T cells are recruited into the immune response at a rate of γ multiplied by the stimulation, where $\gamma = 1\,\text{day}^{-1}$ is the maximum recruitment rate of T cells.

Naïve T cells specific to a particular antigen are in the same stage until they are stimulated. Our model accommodates T cells of different antigen specificities by instantiating separate stage-based models for each, but for the purposes of discussion we will assume that there is only one T cell specificity. As naïve cells are

stimulated, they must wait τ_n hours, representing the developmental time before a naïve cell begins its programmed response. To implement this delay, the cells are promoted through a series of $6\tau_n$ stages, with all cells in a stage moving to the next stage at each 10 min time step. The cells in these stages do not interact with infected cells, but when they emerge after τ_n simulation hours, they become effectors and start responding to infected cells and dividing. In our model, we assume T cells take a minimum of 5 h to divide, and that the first T cell divisions take place 24 h after antigenic stimulation (Oehen and Brduscha-Riem, 1998; Gett and Hodgkin, 2000; Veiga-Fernandes et al., 2000; van Stipdonk et al., 2001), so we chose $\tau_n = 19\,\text{h}$.

Memory cells are recruited in the same manner as naïve cells except that we assume it takes only 1 h ($\tau_m = 1\,\text{h}$) for a stimulated memory cell to begin its programmed response, reflecting the rapid response of memory cells to pathogens (Bachmann et al., 1999; Barber et al., 2003).

3.5. Clearance of infected cells

Because the CTL responses to different antigenic epitopes of the same pathogen do not appear to interfere with each other (Vijh et al., 1999), we model the immune response to multiple epitopes as the sum of independent responses to the individual epitopes. Therefore, we need only define the clearance of infected cells expressing a single epitope by many T cell clones. We assume that effector T cells of clone j, E_j, bind to infected cells I in reversible reactions (at rates k^b for binding and k^d for dissociation) to form complexes C_j, and that effectors bound in these complexes clear the infected cells at rate k^c:

$$E_j + I \underset{k_j^d}{\overset{k_j^b}{\rightleftarrows}} C_j \overset{k_j^c}{\rightarrow} E_j. \tag{8}$$

Directly translating the above expression to a differential equation:

$$\dot{C}_j = k_j^b \hat{E}_j \hat{I} - (k_j^d + k_j^c)C_j, \tag{9}$$

where $\hat{E}_j$ and $\hat{I}$ are unbound effectors and infected cells, respectively. Changing variables to total cells and conserving the number of infected cells, as suggested in Borghans et al. (1996), gives

$$\dot{C}_j = k_j^b(E_j - C_j)\left(I - \sum_k C_k\right) - (k_j^d + k_j^c)C_j, \tag{10}$$

where $\sum_k C_k$ is the number of complexes of all effector cells of all specificities with I. Assuming quasi-steady

state:

$$0 = k_j^b \left(E_j I - C_j I - E_j \sum_k C_k + C_j \sum_k C_k \right)$$
$$- (k_j^d + k_j^c) C_j. \qquad (11)$$

Following De Boer and Perelson (1995), we approximate the solution to Eq. (11) by assuming the $C_j C_k$ terms are small enough to be omitted:

$$C_j \approx \frac{E_j I - E_j \sum_k C_k}{I + K_j}, \qquad (12)$$

where $K_j = (k_j^d + k_j^c)/k_j^b$.

Following the derivation from the Appendix of De Boer and Perelson (1995), the solution to Eq. (12) when there are multiple T cell clones is

$$C_j \approx \frac{I E_j}{K_j + I + \sum_k E_k (I + K_j)/(I + K_k)}. \qquad (13)$$

Therefore, the clearance rate of I due to effectors of all specificities is

$$\dot{I} = -\sum_j k_j^c C_j$$
$$\approx \sum_j -k_j^c \frac{I E_j}{K_j + I + \sum_k E_k (I + K_j)/(I + K_k)}, \qquad (14)$$

For a system with only one T cell clone, E:

$$\dot{I} \approx \frac{-k^c I E}{K + I + E}. \qquad (15)$$

Expression (15) yields a dose–response relationship between effector cell numbers and the infected cell clearance rate that saturates at $k^c I$ as the number of effector cells increases, which agrees with experimental findings (Lehmann-Grube, 1988). It also includes a term for inter-clonal competition among the effector cells for infected cells expressing a single epitope. It appears that high- and low-avidity CTL lyse their targets at similar rates (Derby et al., 2001), so we set k^c to be the same for all T cell clones in our model. In LCMV responses, the value of k^c was found to be 12 day^{-1} (Barchet et al., 2000). Smaller populations of T cells might have higher per capita killing rates, but we assume that most of an infection is resolved while the effector cell population is large. In our model, increased avidity, K, affects the ability to detect and bind to infected cells at low concentrations of I. Multiple T cell clones clear infected cells at the rate described by Eq. (14), in which T cells compete for access to infected cells based on their avidities to them. High-avidity clones are more effective at clearing infected cells than low-avidity clones.

We assume that effector cell mediated clearance of infected cells is a Poisson process. From Eq. (15), we can determine the expected number of infected cells to be cleared in a time interval Δt to be $\dot{I} \Delta t$. We compute the number of infected cells that are cleared during Δt by randomly drawing from the Poisson distribution $\mathscr{P}(\dot{I} \Delta t)$

at each time step. This term is subtracted from the right-hand side of Eq. (5) to include the effect of cytotoxic T cell clearance on the infected cell population.

3.6. T cell replication

We implement the programmed divisions of newly activated effector cells by keeping track of the number of times a cell divides. When a naïve cell is first stimulated, it joins the cohort of effector cells that have not yet divided. When it reproduces, it is moved with its daughter to the next division cohort. We adopt the transition probability cell cycle model described by Smith and Martin (1973), which has two phases: an A phase with a variable residence time and a B phase that takes a fixed length of time to traverse. Cells start in phase A, in which the cells do not divide. At each time step, a cell has a constant probability of entering phase B, during which it divides in a fixed amount of time. At the end of the B phase, both the parent cell and the new daughter cell enter the A phase. This two-phase model enforces a minimum time to cell division. Without the fixed length B phase, some cells could divide an arbitrarily large number of times in a time interval, which is a characteristic of continuous models of cell replication.

To implement the Smith–Martin cell cycle model, each division cohort is subdivided into an A phase and a set of B phase sub-cohorts (Fig. 4). To mimic the fixed length of time it takes a cell to traverse the B phase we allocate one B phase sub-cohort per time step that the cells remain in B phase, and move cells from one sub-cohort to the next at each simulation time step (Fig. 4). We use 10-min time steps, so to model cells remaining in the B phase for n hours, we use $6nB$ phase sub-cohorts
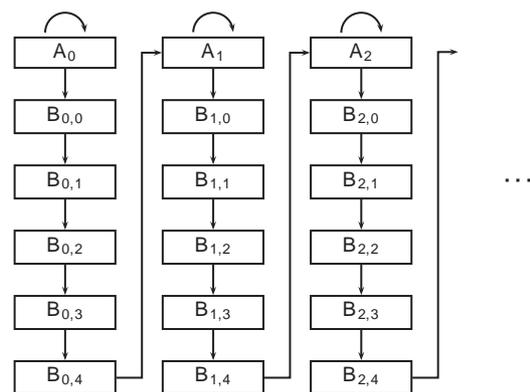

Fig. 4. Implementation of the Smith and Martin two-phase cell cycle model (Smith and Martin, 1973). Each box represents the cells in a given stage, and the arrows represent possible transitions between stages. Note that cells in A phase can either remain in A phase or transition to B phase, while B phase cells progress at a fixed rate until they reach A phase. In this figure, each B sub-stage is 1 h, and in the model implementation each sub-stage is 10 min.

per division cohort. At each time step we draw from a binomial distribution to determine the number of cells in the A phase that transition to the B phase for each division cohort.

We assume that the average cell cycle time of an effector T cell is 6 h and that the minimum time to division is about 5 h (van Stipdonk et al., 2001). Thus, we choose the duration of the B phase to be 5 h and the average duration of the A phase to be 1 h. To simulate a 5 h B phase using 10-min simulation time steps, we use 30 sub-cohorts. To mimic the 1 h average residence in the A phase, we assume the rate at which cells in A phase transition to B phase is $1\,h^{-1}$. We convert this rate to the probability that A phase cells will transition to B phase in a time step in the manner described in Section 3.3 and draw from the binomial distribution to determine how many cells performed the transition. Because T cells with different specificities seem to expand at the same rate in vivo (Busch et al., 1998b), all cells in the model share the same cell cycle parameters. When a death rate of $\delta_E = 0.6\,\mathrm{day}^{-1}$ is included (Veiga-Fernandes et al., 2000), the cell population grows at a rate of $0.092\,h^{-1}$, or about 9-fold per day. T cells divide for about 5 days (Lehmann-Grube, 1988), which implies that a single naïve T cell can generate 60,000 effector cells, which agrees with experiment (Welsh and Selin, 2002). If we assume that a T cell cannot divide more than 100 times, there could to be up to 3100 subpopulations of effector cells per T cell clone, or 100 A phase subpopulations and 3000 B phase subpopulations. These 3100 subpopulations efficiently represent the approximately 600,000 cells (i.e., 10 naïve cells per clone (Casrouge et al., 2000) and 60,000 effectors from each naïve cell) that can originate from a single clone in an immune response.

After their programmed divisions, the cells stop dividing (Badovinac et al., 2002). We assume that during the entire lifetime of the activated T cell, they are subject to the same high death rate δ_E. Thus, cell populations that have stopped dividing are subject to rapid population decline.

3.7. Memory

In our model, effector cells have a 2% per day chance of becoming memory cells after 5 cell divisions (Oehen and Brduscha-Riem, 1998; Opferman et al., 1999), which results in a final memory pool that is about 5% of the peak response (De Boer et al., 2001). Memory cells are dormant until antigenic stimulation. We assume that they have the same sensitivity to antigen as naïve cells, but they enter cell cycle only one hour after antigenic stimulation. Memory-derived effectors have a lower death rate than naïve-derived effectors (Grayson et al., 2002), and we set this rate to be $\delta_{Em} = 0.4\,\mathrm{day}^{-1}$ (Veiga-Fernandes et al., 2000).

4. Results

Our model reproduces population-level phenomena seen experimentally in laboratory mice, and we describe some of these results below. We begin with experiments that illustrate the basic differences between primary and secondary responses and the dynamics of CTL of different affinities, then proceed to describe simulation runs that replicate results found in mouse experiments.

4.1. Primary and secondary immune responses

We simulated the primary and secondary responses to an acute infection (Fig. 5). For this trial, we were not attempting to match our results to a particular laboratory experiment but were instead interested in testing the overall dynamics of the T cell response in the model. We simulated the injection of 500 virus particles into a mouse with a single high-affinity T cell clone of 50 cells. The primary response began after approximately 1 day. The response peaked at day 9 then declined to form a stable memory pool. At day 28, an identical injection was administered, and the secondary response was faster and larger than the primary (Fig. 5). The secondary response began soon after secondary exposure to the virus, and the lower death rate of memory-derived effectors caused the T cell population to increase more rapidly. The secondary response also created a larger pool of stable memory cells. Therefore, the simulated mouse's immune memory could be "boosted" by multiple exposures to the same antigen, making future responses to it even more effective.

4.2. High- and low-avidity responses

To study the clonal composition of the T cell response, we ran the model with a virus with a single
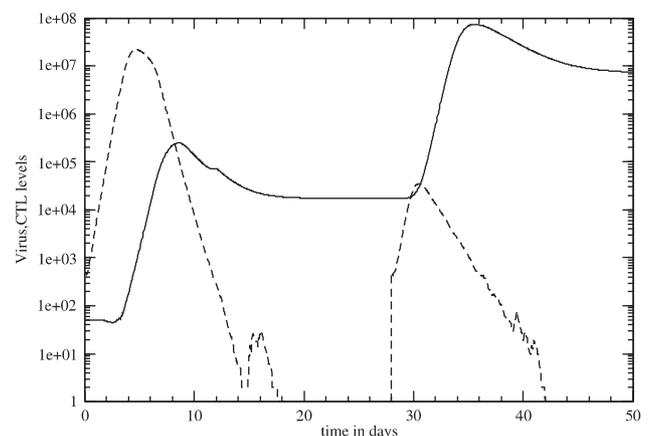

Fig. 5. Primary and secondary CTL responses to viral infection. The dashed lines represent virus levels, with the secondary exposure to the virus at day 28. The solid line represents the number of T cells specific to this virus, including naïve, effector, and memory cells.
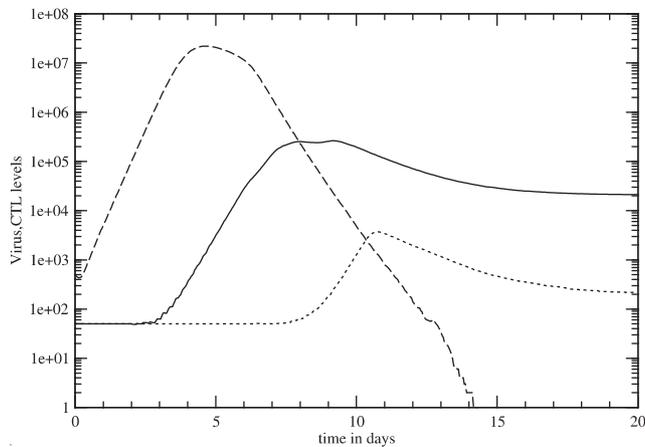
Fig. 6. High- and low-avidity responses. The high-avidity clone (solid line) peaks about 2 days earlier than the low-avidity clone (dotted line). The dashed line represents the simulated viral load.
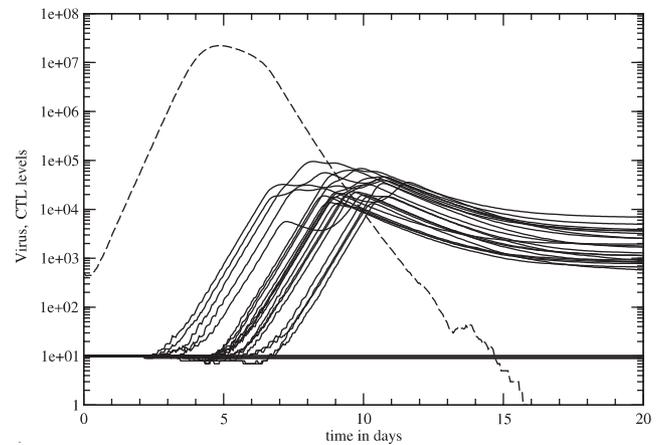

Fig. 7. Simultaneous responses of a large number of clones. The dashed line represents the simulated viral load, and the solid lines represent the responses of T cells from the 29 clones that survived thymic selection.

epitope and two T cell clones with different avidities for this epitope, a high-avidity one ($K = 7.8 \times 10^3$) and a low avidity one ($K = 4.5 \times 10^7$). We assumed both clones initially contained 50 naïve cells each. The peak of the high-avidity clone's response is over one log greater than and over 1 day earlier than the low avidity one (Fig. 6).

One of the strengths of our model is that we can create a large repertoire of CTL with different avidities to various antigens. Perhaps 20 T cell clones respond to a single epitope (Maryanski et al., 1996; Blattman et al., 2002). These clones have affinities not only for the epitope in question, but for all possible epitopes. In a system subjected to heterologous infections, memory cells that cross-react to multiple antigens might be an essential part of our immune responses (Welsh and Selin, 2002). Our digit string implementation, which implicitly defines an affinity between a TCR and any epitope, allows us to model the effect of infections over an organism's lifetime.

We simulated the response of a mouse with a realistic-sized repertoire to a viral infection. We used the xor distance rule with an alphabet size of 128 and set the MHC string length to be 4 digits and the peptide string length to be 6 digits. The simulated mouse had 2.5×10^8 T cell specificities before thymic selection, but only 235 of these specificities were explicitly generated, the remaining clones falling outside the cross-reactive cutoff of the antigen. After thymic selection against 30,000 randomly generated self peptides, 29 of the 235 clones survived to join the naïve repertoire. Fig. 7 shows the dynamics of these clones when exposed to the infection.

4.3. The programmed response

One of the implications of the programmed T cell response is that the immune response is initiated

by antigen but its outcome is antigen-independent. If this is true, then removing antigen after the start of a response should not affect it. This was tested in mice infected by *L. monocytogenes* (Mercado et al., 2000; Badovinac et al., 2002). Antibiotics were administered to eliminate the infection 24 h after inoculation, which quickly removed all antigen. The peak of the T cell response occurred at the same time in the antibiotic-treated mice and in the non-treated mice. The elimination of the infectious agent only caused a small reduction in the magnitude of the response. Therefore, the elimination of antigen did not significantly affect the timing or the magnitude of the T cell response.

In our model, we obtained qualitatively similar results in an LCMV system (Fig. 8). Since antibiotics do not act instantly and do not directly remove bacteria in mice infected with *L. monocytogenes*, we chose to eliminate all LCMV at 36 h post-infection instead of 24. Eliminating the infection caused the peak of the response to occur 1 day earlier and decrease only slightly in magnitude. The reduced response in our model was due to the shortened recruitment time for naïve cells.

Incorporating the programmed response might be essential to modeling the efficacy of vaccinations. Vaccines often use attenuated strains of pathogens that have diminished or no reproductive capacity and are rapidly cleared from the system.

Since the object of vaccination is to induce a large response in order to build a large pool of specific memory cells, then a large dose of an attenuated virus might be effective even if the virus level drops rapidly. If the T cell response were totally antigen-dependent, short periods of antigenic stimulation would not stimulate an adequate response.
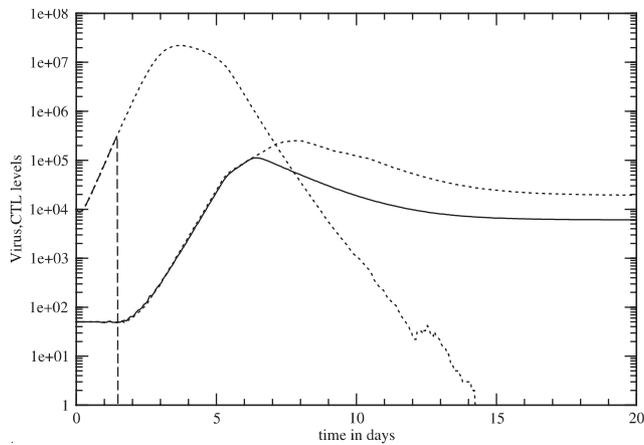
Fig. 8. T cell response to an infection interrupted by treatment. The starting dose of the antigen (dashed line) was 10,000 virus particles. The antigen was removed from the system after 36 h. The T cell response (solid line) is almost unaffected by the removal of antigen. For simplicity, only a single T cell specificity and a single antigenic epitope were used. The antigen and T cell levels of the control case, in which the antigen is not removed, are plotted for comparison (fine dotted lines).
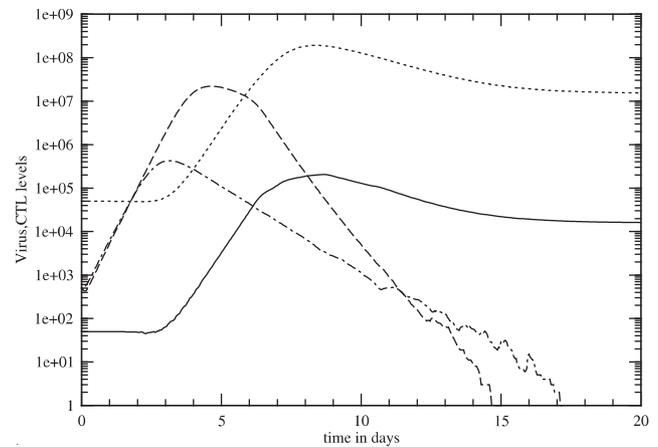

Fig. 9. The effect of increasing the number of naïve cells. One model run was initialized with 50 naïve cells (solid line) and 500 viruses (dashed line). The other model run started with 50,000 naïve cells (dotted line) and the same virus load (dash–dot line).

4.4. Naïve population size effects

The size of the initial naïve cell population can affect the outcome of an infection. Presumably, increasing the number of naïve cells can result in an earlier and larger response to infection. This hypothesis was tested experimentally in mice (Ehl et al., 1998). The number of naïve cells in mice was experimentally increased before infection in order to determine how the number of responding naïve cells affects the T cell response to an acute infection. It was estimated that about 50 naïve cells respond to LCMV in a normal mouse (Ehl et al., 1998), and the number was raised to 50,000 by adoptive transfer. Increasing the number of naïve cells by 1000-fold moved the peak of the infection between 1 and 2 days earlier and reduced the viral load by about 2 logs. In other words, the infection was smaller and eliminated sooner. Our model's results are in agreement with these experiments; after increasing the number of naïve cells from 50 to 50,000, the peak virus load was 1 day earlier and about 2 logs smaller than in the control case (Fig. 9).

4.5. Software

We have released the software for the T cell model used to generate the results presented in this paper (available at http://www.cs.unm.edu/~dlchao). The code is licensed under the GNU General Public License (GPL), which allows anyone to modify and freely distribute the code. The software is written in Java so that it can be run on a wide variety of platforms. On a typical x86 processor running at 1 GHz, the program

takes less than a minute to run for a month of simulated time with one antigen and a realistic number of T cell clones (about 20).

5. Conclusion

We have presented a stochastic stage-structured model of the CTL response to viral infections that features realistic behavior on the level of the individual cells yet is more efficient than standard agent-based approaches to modeling. Our model incorporates antigen- and affinity- dependent stimulation of naïve cells as well as an antigen- independent programmed proliferative response and differentiation into effector cells as suggested by recent experiments. A benefit of our approach is that it allows one to examine the simulated response at many levels, from the total T cell count to the number of responding cells in each clone. Our model takes into consideration cross-reactivity as well as the effects of affinity and peptide-MHC density on the kinetics and clonal composition of the response. Our model can be used to gain qualitative as well as quantitative insights into cellular immune responses. For example, we used the model to analyse the effects of increasing naïve cell frequency and elimination of antigen by drug therapy.

The creation of detailed cellular-level models tests the intuition of immunologists. When the mechanisms of the immune response are studied in isolation, it is easy to lose sight of the system as a whole. Our model integrates the knowledge gained from studying small populations of T cells to form a coherent system that simulates an organism's immune response to a pathogen. This

approach bridges the gap between cellular-level and whole organism studies, and in the future may be useful in designing therapies or in gaining insights into how to modulate the immune response to provide greater protection from disease.

References

Antia, R., Bergstrom, C.T., Pilyugin, S.S., Kaech, S.M., Ahmed, R., 2003. Models of CD8+ responses: 1. What is the antigen-independent proliferation program. J. Theor. Biol. 221, 585–598.

Armstrong, C., Lillie, R., 1934. Experimental lymphocytic choriomeningitis of monkeys and mice produced by a virus encountered in the studies of 1993 St Louis encephalitis epidemic. Publ. Health Rep. 49, 1019–1027, 1CMV.

Arstila, T.P., Casrouge, A., Baron, V., Even, J., Kanellopoulos, J., Kourilsky, P., 1999. A direct estimate of the human $\alpha\beta$ T cell receptor diversity. Science 286, 958–961.

Bachmann, M.F., Barner, M., Viola, A., Kopf, M., 1999. Distinct kinetics of cytokine production and cytolysis in effector and memory T cells after viral infection. Eur. J. Immunol. 29, 291–299.

Badovinac, V.P., Porter, B.B., Harty, J.T., 2002. Programmed contraction of CD8+ T cells after infection. Nat. Immunol. 3, 619–626.

Barber, D.L., Wherry, E.J., Ahmed, R., 2003. Cutting edge: rapid in vivo killing by memory CD8 T cells. J. Immunol. 171, 27–31.

Barchet, W., Oehen, S., Klenerman, P., Wodarz, D., Bocharov, G., Lloyd, A.L., Nowak, M.A., Hengartner, H., Zinkernagel, R.M., Ehl, S., 2000. Direct quantitation of rapid elimination of viral antigen-positive lymphocytes by antiviral CD8+ T cells in vivo. Eur. J. Immunol. 30, 1356–1363.

Bell, G.I., 1970. Mathematical model of clonal selection and antibody production. J. Theor. Biol. 29, 191–232.

Bernaschi, M., Castiglione, F., 2001. Design and implementation of an immune system simulator. Comput. Biol. Med. 31, 303–331.

Blackman, M., Kappler, J., Marrack, P., 1990. The role of the T cell receptor in positive and negative selection of developing T cells. Science 248, 1335–1341.

Blattman, J.N., Sourdive, D.J., Murali-Krishna, K., Ahmed, R., Altman, J.D., 2000. Evolution of the T cell repertoire during primary, memory, and recall responses to viral infection. J. Immunol. 165, 6081–6090.

Blattman, J.N., Antia, R., Sourdive, D.J., Wang, X., Kaech, S.M., Murali-Krishna, K., Altman, J.D., Ahmed, R., 2002. Estimating the precursor frequency of naive antigen-specific CD8 T cells. J. Exp. Med. 195, 657–664.

Bocharov, G.A., 1998. Modelling the dynamics of LCMV infection in mice: conventional and exhaustive CTL responses. J. Theor. Biol. 192, 283–308.

Borghans, J.A., De Boer, R.J., Segel, L.A., 1996. Extending the quasi-steady state approximation by changing variables. Bull. Math. Biol. 58, 43–63.

Bousso, P., Levraud, J.P., Kourilsky, P., Abastado, J.P., 1999. The composition of a primary T cell response is largely determined by the timing of recruitment of individual T cell clones. J. Exp. Med. 189, 1591–1600.

Bousso, P., Lemaitre, F., Bilsborough, J., Kourilsky, P., 2000. Facing two T cell epitopes: a degree of randomness in the primary response is lost upon secondary immunization. J. Immunol. 165, 760–767.

Busch, D.H., Pilip, I., Pamer, E.G., 1998a. Evolution of a complex T cell receptor repertoire during primary and recall bacterial infection. J. Exp. Med. 188, 61–70.

Busch, D.H., Pilip, I.M., Vijh, S., Pamer, E.G., 1998b. Coordinate regulation of complex T cell populations responding to bacterial infection. Immunity 8, 353–362.

Butz, E.A., Bevan, M.J., 1998. Massive expansion of antigen-specific CD8+ T cells during an acute virus infection. Immunity 8, 167–175.

Byers, A.M., Kemball, C.C., Moser, J.M., Lukacher, A.E., 2003. Cutting edge: rapid in vivo CTL activity by polyoma virus-specific effector and memory CD8+ T cells. J. Immunol. 171, 17–21.

Casrouge, A., Beaudoing, E., Dalle, S., Pannetier, C., Kanellopoulos, J., Kourilsky, P., 2000. Size estimate of the $\alpha\beta$ TCR repertoire of naive mouse splenocytes. J. Immunol. 164, 5782–5787.

Celada, F., Seiden, P.E., 1992. A computer model of cellular interactions in the immune system. Immunol. Today 13, 56–62.

Chao, D.L., Davenport, M.P., Forrest, S., Perelson, A.S., 2003. Stochastic stage-structured modeling of the adaptive immune system. In: Proceedings of the IEEE Computer Society Bioinformatics Conference (CSB 2003). IEEE Press, Los Alamitos, CA, pp. 124–131.

Cohn, M., Langman, R.E., Mata, J.J., 2002. A computerized model for the self-non-self discrimination at the level of the T(h) (Th genesis). I. The origin of 'primer' effector T(h) cells. Int. Immunol. 14, 1105–1112.

Davenport, M.P., Fazou, C., McMichael, A.J., Callan, M.F., 2002. Clonal selection, clonal senescence, and clonal succession: the evolution of the T cell response to infection with a persistent virus. J. Immunol. 168, 3309–3317.

De Boer, R.J., Perelson, A.S., 1995. Towards a general function describing T cell proliferation. J. Theor. Biol. 175, 567–576.

De Boer, R.J., Oprea, M., Antia, R., Murali-Krishna, K., Ahmed, R., Perelson, A.S., 2001. Recruitment times, proliferation, and apoptosis rates during the CD8+ T-cell response to lymphocytic choriomeningitis virus. J. Virol. 75, 10663–10669.

Derby, M., Alexander-Miller, M., Tse, R., Berzofsky, J., 2001. High-avidity CTL exploit two complementary mechanisms to provide better protection against viral infection than low-avidity CTL. J. Immunol. 166, 1690–1697.

Detours, V., Perelson, A.S., 2000. The paradox of alloreactivity and self restriction: quantitative analysis and statistics. Proc. Natl Acad. Sci. USA 97, 8479–8483.

Detours, V., Mehr, R., Perelson, A.S., 1999. A quantitative theory of affinity-driven cell repertoire selection. J. Theor. Biol. 200, 389–403.

Dibrov, B.F., Livshits, M.A., Volkenstein, M.V., 1977. Mathematical model of immune processes. J. Theor. Biol. 65, 609–631.

Dutton, R.W., Bradley, L.M., Swain, S.L., 1998. T cell memory. Annu. Rev. Immunol. 16, 201–223.

Ehl, S., Klenerman, P., Zinkernagel, R.M., Bocharov, G., 1998. The impact of variation in the number of CD8$^+$ T-cell precursors on the outcome of virus infection. Cell Immunol. 189, 67–73.

Farmer, J.D., Packard, N.H., Perelson, A.S., 1986. The immune system, adaption and machine learning. Physica D 22, 187–204.

Germain, R.N., 2001. The art of the probable: system control in the adaptive immune system. Science 293, 240–245.

Gett, A.V., Hodgkin, P.D., 2000. A cellular calculus for signal integration by T cells. Nat. Immunol. 1, 239–244.

Gillespie, D.T., 1977. Exact stochastic simulation of coupled chemical reactions. J. Phys. Chem. 81, 2340–2361.

Grayson, J.M., Harrington, L.E., Lanier, J.G., Wherry, E.J., Ahmed, R., 2002. Differential sensitivity of naive and memory CD8$^+$ T cells to apoptosis in vivo. J. Immunol. 169, 3760–3770.

Hamming, R.W., 1950. Error detecting and error correcting codes. Bell Syst. Tech. J. 29, 147–160.

Ho, D.D., Neumann, A.U., Perelson, A.S., Chen, W., Leonard, J.M., Markowitz, M., 1995. Rapid turnover of plasma virions and CD4 lymphocytes in HIV-1 infection. Nature 373, 123–126.

Kaech, S.M., Ahmed, R., 2001. Memory CD8$^+$ T cell differentiation: initial antigen encounter triggers a developmental program in naïve cells. Nat. Immunol. 2, 415–422.

Kaech, S.M., Hemby, S., Kersh, E., Ahmed, R., 2002. Molecular and functional profiling of memory CD8 T cell differentiation. Cell 111, 837–851.

Kappler, J.W., Roehm, N., Marrack, P., 1987. T cell tolerance by clonal elimination in the thymus. Cell 49, 273–280.

Kersh, E.N., Kaech, S.M., Onami, T.M., Moran, M., Wherry, E.J., Miceli, M.C., Ahmed, R., 2003. TCR signal transduction in antigen-specific memory CD8 T cells. J. Immunol. 170, 5455–5463.

Kleinstein, S.H., Seiden, P.E., 2000. Computer simulations: simulating the immune system. Comput. Sci. Eng. 2, 69–77.

Langman, R.E., Mata, J.J., Cohn, M., 2003. A computerized model for the self-non-self discrimination at the level of the T(h) (Th genesis). II. The behavior of the system upon encounter with non-self antigens. Int. Immunol. 15, 593–609.

Lefkovitch, L.P., 1965. The study of population growth in organisms grouped by stages. Biometrics 21, 1–18.

Lehmann-Grube, F., 1988. Mechanism of recovery from acute virus infection. In: Bauer, H., Klenk, H.-D., Scholtissek, C. (Eds.), Modern Trends in Virology. Springer, Berlin, pp. 49–64.

Lin, M.Y., Welsh, R.M., 1998. Stability and diversity of T cell receptor repertoire usage during lymphocytic choriomeningitis virus infection of mice. J. Exp. Med. 188, 1993–2005.

Manly, B.F.J., 1990. Stage-Structured Populations: Sampling, Analysis, and Simulation. Chapman & Hall, London.

Marsh, S.G., Albert, E.D., Bodmer, W.F., Bontrop, R.E., Dupont, B., Erlich, H.A., Geraghty, D.E., Hansen, J.A., Mach, B., Mayr, W.R., Parham, P., Petersdorf, E.W., Sasazuki, T., Schreuder, G.M., Strominger, J.L., Svejgaard, A., Terasaki, P.I., 2002. Nomenclature for factors of the HLA system, 2002. Tissue Antigens 60, 407–464.

Maryanski, J.L., Jongeneel, C.V., Bucher, P., Casanova, J.L., Walker, P.R., 1996. Single-cell PCR analysis of TCR repertoires selected by antigen in vivo: a high magnitude CD8 response is comprised of very few clones. Immunity 4, 47–55.

McNally, J.M., Zarozinski, C.C., Lin, M.Y., Brehm, M.A., Chen, H.D., Welsh, R.M., 2001. Attrition of bystander CD8$^+$ T cells during virus-induced T-cell and interferon responses. J. Virol. 75, 5965–5976.

Mercado, R., Vijh, S., Allen, S.E., Kerksiek, K., Pilip, I.M., Pamer, E.G., 2000. Early programming of T cell populations responding to bacterial infection. J. Immunol. 165, 6833–6839.

Murali-Krishna, K., Altman, J.D., Suresh, M., Sourdive, D.J., Zajac, A.J., Miller, J.D., Slansky, J., Ahmed, R., 1998. Counting antigen-specific CD8$^+$ T cells: a reevaluation of bystander activation during viral infection. Immunity 8, 177–187.

Murali-Krishna, K., Lau, L.L., Sambhara, S., Lemonnier, F., Altman, J., Ahmed, R., 1999. Persistence of memory CD8 T cells in MHC class I-deficient mice. Science 286, 1377–1381.

Neumann, A.U., Lam, N.P., Dahari, H., Gretch, D.R., Wiley, T.E., Layden, T.J., Perelson, A.S., 1998. Hepatitis C viral dynamics in vivo and the antiviral efficacy of interferon-alpha therapy. Science 282, 103–107.

Nowak, M.A., Bangham, C.R., 1996. Population dynamics of immune responses to persistent viruses. Science 272, 74–79.

Ochsenbein, A.F., Sierro, S., Odermatt, B., Pericin, M., Karrer, U., Hermans, J., Hemmi, S., Hengartner, H., Zinkernagel, R.M., 2001. Roles of tumour localization, second signals and cross priming in cytotoxic T-cell induction. Nature 411, 1058–1064.

Oehen, S., Brduscha-Riem, K., 1998. Differentiation of naive CTL to effector and memory CTL: correlation of effector function with phenotype and cell division. J. Immunol. 161, 5338–5346.

Opferman, J.T., Ober, B.T., Ashton-Rickardt, P.G., 1999. Linear differentiation of cytotoxic effectors into memory T lymphocytes. Science 283, 1745–1748.

Pannetier, C., Cochet, M., Darche, S., Casrouge, A., Zoller, M., Kourilsky, P., 1993. The sizes of the CDR3 hypervariable regions of the murine T-cell receptor beta chains vary as a function of the recombined germ-line segments. Proc. Natl Acad. Sci. USA 90, 4319–4323.

Parker, K.C., Bednarek, M.A., Coligan, J.E., 1994. Scheme for ranking potential HLA-A2 binding peptides based on independent binding of individual peptide side-chains. J. Immunol. 152, 163–175.

Perelson, A.S., 2002. Modelling viral and immune system dynamics. Nat. Rev. Immunol. 2, 28–36.

Perelson, A.S., Oster, G.F., 1979. Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. J. Theor. Biol. 81, 645–670.

Perelson, A.S., Weisbuch, G., 1992. Modeling immune reactivity in secondary lymphoid organs. Bull. Math. Biol. 54, 649–672.

Perelson, A.S., Neumann, A.U., Markowitz, M., Leonard, J.M., Ho, D.D., 1996. HIV-1 dynamics in vivo: virion clearance rate, infected cell life-span, and viral generation time. Science 271, 1582–1586.

Pihlgren, M., Dubois, P.M., Tomkowiak, M., Sjogren, T., Marvel, J., 1996. Resting memory CD8$^+$ T cells are hyperreactive to antigenic challenge in vitro. J. Exp. Med. 184, 2141–2151.

Přikrylová, D., Jílek, M., Waniewski, J., 1992. Mathematical Modeling of the Immune Response. CRC Press, Boca Raton, FL.

Rocha, B., Dautigny, N., Pereira, P., 1989. Peripheral T lymphocytes: expansion potential and homeostatic regulation of pool sizes and CD4$^+$/CD8$^+$ ratios in vivo. Eur. J. Immunol. 19, 905–911.

Seiden, P.E., Celada, F., 1992. A model for simulating cognate recognition and response in the immune system. J. Theor. Biol. 158, 329–357.

Selin, L.K., Vergilis, K., Welsh, R.M., Nahill, S.R., 1996. Reduction of otherwise remarkably stable virus-specific cytotoxic T lymphocyte memory by heterologous viral infections. J. Exp. Med. 183, 2489–2499.

Shortman, K., Egerton, M., Spangrude, G.J., Scollay, R., 1990. The generation and fate of thymocytes. Semin. Immunol. 2, 3–12.

Smith, J.A., Martin, L., 1973. Do cells cycle? Proc. Natl Acad. Sci. USA 70, 1263–1267.

Smith, D.J., Forrest, S., Hightower, R.R., Perelson, A.S., 1997. Deriving shape space parameters from immunological data. J. Theor. Biol. 189, 141–150.

Smith, D.J., Forrest, S., Ackley, D.H., Perelson, A.S., 1998. Using lazy evaluation to simulate realistic-size repertoires in models of the immune system. Bull. Math. Biol. 60, 647–658.

Smith, D.J., Forrest, S., Ackley, D.H., Perelson, A.S., 1999. Variable efficacy of repeated annual influenza vaccination. Proc. Natl Acad. Sci. USA 96, 14001–14006.

Sourdive, D.J., Murali-Krishna, K., Altman, J.D., Zajac, A.J., Whitmire, J.K., Pannetier, C., Kourilsky, P., Evavold, B., Sette, A., Ahmed, R., 1998. Conserved T cell receptor repertoire in primary and memory CD8 T cell responses to an acute viral infection. J. Exp. Med. 188, 71–82.

Traub, E., 1935. A filtrable virus recovered from white mice. Science 81, 298–299, lCMV.

Usher, M.B., 1966. A matrix approach to the management of renewable resources, with special reference to selection forests. J. Appl. Ecol. 3, 355–367.

van Stipdonk, M.J., Lemmens, E.E., Schoenberger, S.P., 2001. Naïve CTLs require a single brief period of antigenic stimulation for clonal expansion and differentiation. Nat. Immunol. 2, 423–429.

Veiga-Fernandes, H., Walter, U., Bourgeois, C., McLean, A., Rocha, B., 2000. Response of naive and memory CD8$^+$ T cells to antigen stimulation in vivo. Nat. Immunol. 1, 47–53.

Vijh, S., Pilip, I.M., Pamer, E.G., 1999. Noncompetitive expansion of cytotoxic T lymphocytes specific for different antigens during bacterial infection. Infect. Immun. 67, 1303–1309.

Wei, X., Ghosh, S.K., Taylor, M.E., Johnson, V.A., Emini, E.A., Deutsch, P., Lifson, J.D., Bonhoeffer, S., Nowak, M.A., Hahn, B.H., et al., 1995. Viral dynamics in human immunodeficiency virus type 1 infection. Nature 373, 117–122.

Welsh, R.M., Selin, L.K., 2002. No one is naive: the significance of heterologous T-cell immunity. Nat. Rev. Immunol. 2, 417–426.

Wherry, E.J., Puorro, K.A., Porgador, A., Eisenlohr, L.C., 1999. The induction of virus-specific CTL as a function of increasing epitope expression: responses rise steadily until excessively high levels of epitope are attained. J. Immunol. 163, 3735–3745.