

## **Somatic Hypermutation in B Cells: An Optimal Control Treatment**

THOMAS B. KEPLER<sup>†</sup> AND ALAN S. PERELSON<sup>‡</sup>

<sup>†</sup> *The Santa Fe Institute, 1660 Old Pecos Trail, Santa Fe, NM 87501,*  
and <sup>‡</sup> *Theoretical Division, Los Alamos National Laboratory, Los Alamos,*  
*NM 87545, U.S.A.*

*(Received on 10 August 1992, Accepted in revised form on 14 January 1993)*

The vertebrate immune system generates high-affinity antibodies to external antigens through a process of somatic hypermutation that takes place in germinal centers formed in the secondary lymphoid tissues. B cells proliferating in these germinal centers experience random mutations in the genes encoding the variable region of their immunoglobulin molecules and are subsequently selected for high-affinity binding to antigen. These germinal center reactions last for only about 2 weeks, yet in that time typically produce multiple point mutations resulting in affinity increases of factors of ten to a hundred or more. We have attempted to understand this extraordinary effectiveness by casting the problem of affinity maturation as an optimization problem in which a quantity that we call the total affinity is maximized as a functional of  $\mu(t)$ , the mutation rate as a function of time. We have developed a single-compartment model for the process and an optimization algorithm based on the Pontryagin maximum principle. Our results show that the optimum mutation schedule is one with brief bursts of high mutation rates interspersed between periods of mutation-free growth. Though this result at first seems highly non-physiological, we show that, in fact, it provides a framework within which the anatomy and kinetics of the germinal center reaction can be understood.

### **1. Introduction**

The immune system recognizes and neutralizes pathogens and discriminates between self and non-self. The system is highly complex and has many intricately co-ordinated subsystems. In this paper we will be concerned with the humoral response, mediated by immunoglobulin molecules or antibodies, and the B cells that produce them.

Immunoglobulin (Ig) is a protein that binds to specific sites (epitopes) on foreign material (antigen). Once bound, this binding facilitates the phagocytosis of the antigen, it can activate a series of enzymatic events within the complement system leading to the lysis of cellular antigens, or if the epitope is a receptor or adhesion molecule, it can competitively inhibit its function. The immune system cannot know what epitopes it may encounter, so it is faced with the formidable challenge of providing Ig that will recognize and bind anything nature might throw at it. Unraveling the strategy by which this is accomplished has occupied much of modern immunology (Silverstein, 1989), and a significant part of the picture is now clear.

When, at the beginning of this century, Landsteiner showed that immune sera are highly specific for the antigens that produced them, immunologists realized that the antibody repertoire had to be enormous. [At the time a figure of  $10^{11}$  different specificities was often cited (Silverstein, 1989).] Several of the theories devised to account for this great diversity were *instructionist* theories. These held that the antigen itself directed the formation of the appropriate antibody shape. These theories fell out of favor when the "central dogma" of molecular biology, declaring that information flows from genes to proteins and not the other way around, was enunciated by Crick. Instructionist theories were replaced by *selection* theories. The *clonal* selection theory of Burnet (1959), which has been validated extensively by experiment and observation, is the current paradigm in immunology.

Clonal selection holds that diversity pre-exists in vast numbers of small clones of resting B cells. These clones are specific in that the B cells in a clone all express Ig molecules with the same variable region on their plasma membranes, which they use as antigen receptors. If a cell binds antigen with its surface Ig, it becomes activated and (after receiving further signals from specific "helper" T cells) sets into motion a complex sequence of events that leads, among other things, to rapid proliferation of the clone, copious secretion of Ig and generation of immune memory. If the cell does not encounter an appropriate antigen, it dies in a few days. Thus, there is rapid turnover of the circulating B-cell population and vigorous development of new clones in the bone marrow.

The impressive diversity exhibited by this *naïve* repertoire is generated by an extraordinary mechanism during development. The Ig molecules are not individually coded in the germline. As was appreciated by the instructionists, such an encoding would have been prohibitively large. Instead, the coding sequences are assembled by recombination of randomly chosen gene segments stored in germline libraries (see, for example, Klein, 1990). So the body anticipates novel antigens by expressing large numbers of randomly chosen Ig specificities, thereby covering the space of possible epitope shapes (Perelson & Oster, 1979).

But even this enormous diversity is inadequate for the production of sufficiently high-affinity antibodies to a given antigen. Analysis of the naïve B-cell repertoire suggests that the immune system's strategy is to use relatively low affinity, broadly reactive Ig in the naïve repertoire to ensure efficient coverage (Chua *et al.*, 1987). This hypothesis is supported by recent evidence that accessory molecules may function to enhance the stimulation of low-affinity B cells (Carter & Fearon, 1992). The generation of high-affinity Ig then must be handled by a separate mechanism. This mechanism is somatic hypermutation (Berek & Milstein, 1987, 1988; French *et al.*, 1989; Berek *et al.*, 1991; Steele, 1991), first hypothesized by Weigert *et al.* (1970).

A macroscopic description of somatic hypermutation may be somewhat reminiscent of the instructionist theories, although the detailed process is substantially different. Once a B-cell clone is activated, some of these activated cells migrate to newly forming germinal centers (GCs) in the lymph nodes and spleen (Jacob *et al.*, 1991; Leanderson *et al.*, 1992; Liu *et al.*, 1992; Nossal, 1992), where they undergo rapid proliferation and extensive mutation. This process of germinal center formation with extensive B-cell proliferation is called the germinal center reaction (GCR).

Remarkably, hypermutation occurs only in the genes coding for the variable region of the Ig light and heavy chains and their flanking regions. These genes determine the structure and specificity of the antigen-binding site. The molecular mechanisms responsible for hypermutation and its regulation are as yet unknown (Steele *et al.*, 1991). Among the mutants generated in this process, some may express Ig with a higher affinity for antigen. The cells bind more antigen and proliferate more than the naive populations. They also deplete the available antigen, whether free in solution or held on follicular dendritic cells (see below) in the GC, thereby reducing the stimulation and growth of the lower-affinity non-mutated cells.

The germinal center reaction lasts for about 2 weeks, and in that time the characteristic affinity of antibody for the immunizing antigen may be increased by one or two orders of magnitude or more, as a result of multiple amino acid substitutions (Berek & Milstein, 1987, 1988; Jacob *et al.*, 1991; Steele, 1991).

The somatic mutation process is remarkably efficient. Given our general expectations that only a small proportion of mutations should yield a selective advantage, the short time actually required to produce significant increases in affinity through multiple advantageous point mutations is quite surprising. This has led to some speculation about the possible occurrence of novel molecular mechanisms and non-Darwinian mutation and selection processes (Manser, 1991). Before dismissing Darwinian mechanisms as being too inefficient to account for the observed number of mutations, which can be as high as ten or 20, it is important to determine the optimal efficiency of Darwinian mechanisms. Thus, in this paper we investigate the extent to which manipulation of the mutation rate as a function of time can enhance the production of large quantities of high-affinity Ig-bearing cells. It is clear that the mutation process is highly regulated: it is directed at a very specific region of the genome, it turns on some time after the onset of the immune response, and turns off again before the response is over (French *et al.*, 1989; Berek *et al.*, 1991). Furthermore, some estimates of mutation rates (McKean *et al.*, 1984; Berek & Milstein, 1987; Gearhart & Levy, 1991), made assuming that these rates are constant, are so high ( $\hat{\mu} = 10^{-3} \text{ bp}^{-1} \text{ generation}^{-1}$ ) that they would prevent any clone from growing at all. As will be shown below, the maximum proliferation rate for fully stimulated B-cell populations is proportional to  $2(1 - \hat{\mu})^N - 1$ , where  $\hat{\mu}$  is the mutation rate per base pair per generation, and  $N$  is the number of nucleotides in the coding region of the Ig variable region. When this expression vanishes, the clone is just maintaining itself; half its daughters are mutants. We call this the critical mutation rate,  $\hat{\mu}_c$ . When the mutation rate is higher than  $\hat{\mu}_c$ , the clone will disappear owing to mutation, while populating neighboring clones. For  $N = 660$ , the approximate number of nucleotides coding for the Ig variable (V) region (Klein, 1990),  $\hat{\mu}_c = 1.05 \times 10^{-3}$ . Above this rate, the clone depletes itself while populating neighboring clones.

Below, we model B-cell clonal proliferation and mutation under the influence of antigen, and apply Pontryagin's maximum principle to find the mutation schedule,  $\mu(t)$ , that maximizes a quantity that we call the *total* affinity, which takes into consideration both the average affinity for the immunizing antigen and the number of B cells involved in the response. We find that the optimal schedule is one with

bursts of high mutation rates between periods of non-mutated growth. We call such a mutation schedule *phasic*.

Our result is related to work by Perelson *et al.* (1976, 1978) who studied the control of B-cell proliferation and differentiation via Pontryagin's maximum principle and showed that, in order to generate the maximal amount of antibody in a fixed time, it is optimal for a B cell to proliferate without differentiating until a critical clone size is reached and then to turn on differentiation to the plasma cell state. Although our problem is a bit more subtle since we are interested in multi-step mutants, our result is also related to the work of Agur *et al.* (1991), who treated a greatly simplified version of the affinity maturation problem and showed that to generate a particular one-step mutant, it is better to let the progenitor population grow to some critical level and then to switch mutation on at its maximum rate, rather than letting mutation operate at some intermediate rate throughout the population's growth phase.

A phasic mutation schedule may at first seem unphysiological. However, we will show that in the context of the anatomy and physiology of the germinal center reaction that a phasic mutation schedule is easy to implement if cells exit and then re-enter a "mutational compartment", such as the dark zone of the germinal center. This leads us to join others who have discussed the importance of possible re-entrant pathways in the GC (Steele *et al.*, 1991; Liu *et al.*, 1992). If such pathways exist, the phasic mutation schedule arises quite naturally as a result of the multi-compartmental structure of the GC. This insight may add significantly to our understanding of the regulation of affinity maturation *in vivo*.

## 2. Model

We model a single-compartment in which B cells are stimulated by antigen binding and proliferate in response to antigen with some probability of generating point mutations at every division. Cells that are not stimulated die. So we consider two processes:

### (i) PROLIFERATION/MUTATION

Let  $X_i$  denote a B cell expressing surface Ig with a variable region of genotype  $i$ . The proliferation/mutation reaction is



and the rate of this reaction is given by  $k_p h_i m_{ij} m_{ik}$ , where  $k_p$  is the rate of cell division at maximum stimulation,  $h_i$  is the degree of activation, dependent on antigen binding in a manner to be specified, and  $m_{ij}$  is the probability that a cell of type  $j$  arises as a daughter of a cell of type  $i$ . The matrix  $[m_{ij}]$  is determined once we specify the overall mutation rate  $\mu$  and how mutation converts one genotype into another. To conserve probability, we must have for all  $i$ ,  $\sum_j m_{ij} = 1$ . When  $k = j = i$ , eqn (1) models  $X_i \rightarrow X_i + X_i$ , i.e. cell proliferation.

## (ii) CELL DEATH

Any cell that is not stimulated dies (Liu *et al.*, 1989), so the rate of cell death in clonal population  $i$  is  $k_d(1 - h_i)$ , where  $k_d$  is the rate constant characterizing cell death.

## 2.1. ANTIGENIC STIMULATION

The interaction of cell surface immunoglobulin with antigen is reversible and may be written as



where  $r_i$  is a free receptor of type  $i$ ,  $a$  is unbound antigen, and  $r_i^*$  is the bound receptor-antigen complex†. The equilibrium constant for this reaction is  $K_i$ . We assume that on the time scale of cell division, these physicochemical reactions are always at equilibrium. At this equilibrium, the concentration of bound receptors is given by

$$[r_i^*] = \frac{K_i[a][r_i]_0}{1 + K_i[a]} \quad (3)$$

where  $[r_i]_0$  is the total concentration of receptors of type  $i$  and  $[a]$  is the unbound antigen concentration.

If we assume that antigen is conserved, then we have the additional equation

$$[a] + \sum_i [r_i^*] = [a]_0 \quad (4)$$

where  $[a]_0$  is the total concentration of antigen. At first sight, this assumption seems unusual, but it may be a fairly good approximation. There is a population of cells, the follicular dendritic cells (FDC) (Tew *et al.*, 1990), that are critical for the formation and maintenance of GCs (Kosco *et al.*, 1992). These cells bind antigen in the form of antigen-antibody complexes on their surfaces for presentation to B cells. Furthermore, they hold these complexes on their surfaces for very long times, perhaps as long as the life-time of the host organism, and are thought to be one of the mechanisms responsible for maintenance of immune memory (Gray & Skarvall, 1988; Gray & Leanderson, 1990; Gray, 1992). One may view our model as assuming that antigen is "borrowed" from the follicular dendritic cells but is not destroyed. A more detailed model would account for antigen degradation as well as antigen influx from the circulation, but including these phenomena makes most sense in the context of a multi-compartment model in which the generation of memory cells is explicitly considered. In the mean time, we believe our model represents the relevant facts adequately.

We are interested in the receptor occupation density only so far as it affects the activation level  $h_i$ . We make the simplifying assumption that  $h_i$  is given by the

† We are ignoring details of cross-linking and antigen valence because at our present level of understanding of the mechanisms underlying B-cell stimulation in GCs, incorporating such information would overly complicate what is meant to be a simple model.

fraction of bound receptors of type  $i$ ,

$$h_i = \frac{K_i[a]}{1 + K_i[a]}. \quad (5)$$

If there are  $\sigma$  receptors per cell and  $[X_i]$  cells of type  $i$ , then  $[r_i]_0 = \sigma[X_i]$  and

$$[a] + \sigma \sum_i h_i [X_i] = [a]_0. \quad (6)$$

## 2.2. DYNAMICAL EQUATIONS

To formulate the dynamics of this system we introduce the following dimensionless variables:  $x_i = \sigma[X_i]/[a]_0$ ,  $\alpha = [a]/[a]_0$  and  $\kappa_i = [a]_0 K_i$ . In these new variables, the activation level is given by

$$h_i = \frac{\alpha \kappa_i}{1 + \alpha \kappa_i} \quad (7)$$

and the algebraic equation expressing conservation of antigen becomes

$$\alpha + \sum_i h_i x_i = 1. \quad (8)$$

The number of cells of type  $i$  varies owing to proliferation, mutation and death. We model the kinetics of these processes by the deterministic differential equations

$$\frac{dx_i}{dt} = x_i \theta_i \{-k_d(1 - h_i) + k_p h_i(2m_{ii} - 1)\} + 2k_p \sum_{j \neq i} m_{ji} h_j x_j \theta_j. \quad (9)$$

Ignoring  $\theta$  for the moment, the first term on the right is the rate of death of unstimulated cells, the second term is the rate of proliferation of cells of type  $i$ , and the third term is the rate at which cells of type  $j$  have daughter cells that have mutated into type  $i$ . In the proliferation term, the factor  $2m_{ii} - 1$  arises from the fact that according to scheme (1), proliferation occurs when  $i = j = k$ , and hence both daughter cells are of type  $i$ , giving the term  $2m_{ii}$ , and one parent of type  $i$  is lost. The factor 2 in the last term is a result of the fact that either of two daughter cells of a cell of type  $j$  may be of type  $i$ .

The factor  $\theta_i$  is shorthand for the step function  $\theta(x_i - \xi)$ , and  $\xi$  is a threshold equivalent to one cell. The rationale for the inclusion of these  $\theta$ -terms is straightforward. We are preventing a clone's proliferative mechanisms from operating until that clone reaches an occupation level equivalent to one cell. In a conventional continuum description, the system size is taken to infinity, so that the concentration assigned to populations with a single cell is infinitesimal. For systems with mutation, the waiting time for the appearance of a given mutation scales as the inverse system size. Thus, in a conventional continuum description, all possible mutants appear immediately. This can lead to the growth and dominance of high-affinity clones long before they would, in a more careful analysis, be expected to have a single member! Owing to the non-dimensionalization, the threshold parameter  $\xi$  scales like  $1/[a]_0$ .

and since  $[a]_0$  fixes the system size, we recover the naive continuum limit as  $[a]_0 \rightarrow \infty$ .

### 2.3. ASYMPTOTIC BEHAVIOR IN THE ABSENCE OF MUTATION

We can already say a few things of some consequence about this model. First, the model realizes strong competition for antigen as discussed above. In the absence of mutation, the clone with the highest affinity and having supra-threshold concentration (i.e. having at least one cell present) will grow at the expense of all other clones, so that the asymptotic solution is always monoclonal (or null, if there is insufficient antigen to stimulate any cells). To see this, note that in the absence of mutation,  $m_{ij} = 0$  for  $j \neq i$ , and since  $\sum_j m_{ij} = 1$ ,  $m_{ii} = 1$ . Thus, eqn (9) becomes

$$\frac{dx_i}{dt} = x_i \theta_i \{ -k_d(1 - h_i) + k_p h_i \}, \quad (10)$$

from which it is clear that the clone with the highest level of stimulation,  $h_i$ , will grow the fastest. Since  $h_i$  increases monotonically with the affinity  $\kappa_i$  [see eqn (7)], the clone with the highest affinity grows the fastest. At equilibrium, the stimulation of the fastest growing clone,  $h = k_d/(k_d + k_p)$ , and hence its rate of growth will just balance its rate of death. Slower growing clones will necessarily die faster than they are growing and hence be eliminated.

If the (dimensionless) affinity of the highest-affinity clone is  $\kappa$ , then from (8) and (10) the asymptotic solution is given by

$$x(\infty) = \frac{k_d + k_p}{k_d} \{ 1 - \alpha(\infty) \} \quad (11)$$

and by (8) the free antigen concentration is

$$\alpha(\infty) = \frac{k_d}{\kappa k_p}. \quad (12)$$

This solution with a single clone is stable. If we perturb this system by adding the equivalent of one or more cells from another clone of affinity  $\kappa'$ , then one of two things can happen. If  $\kappa' < \kappa$ , then the net growth rate of the second clone must be smaller than that of the first. But since the first was balanced between cell division and cell death, the second must decline until it falls below threshold. If, instead,  $\kappa' > \kappa$ , then the second clone grows. In addition, if  $\kappa' > \kappa$ , the second clone is a more thorough exploiter of antigen, so the concentration of free antigen drops and the equilibrium of the first clone is upset, so it begins to decline until it is eliminated. If mutation is present, however, multi-clonality can be maintained. Note that these results depend only on  $h_i$  being a monotonically increasing function of the affinity  $\kappa_i$  and not on the particular functional form for  $h_i$  that we have used.

### 2.4. THE MUTATION MATRIX $[m_{ij}]$

In order to completely determine the model, we must specify the "vector" of affinities  $\kappa_i$  and the mutation matrix  $[m_{ij}]$ . The mutation matrix is fixed by the

relationship between genetic sequences. If we use an  $n$ -nucleotide genetic alphabet and a coding sequence  $N$  nucleotides long, then the matrix is given by

$$m_{ij} = \left( \frac{\hat{\mu}}{n-1} \right)^{d_{ij}} (1 - \hat{\mu})^{N-d_{ij}} \quad (13)$$

where  $d_{ij}$  is the number of sites at which clones  $i$  and  $j$  differ, and  $\hat{\mu}$  is the probability of mutation per nucleotide per generation. Since there are  $n$  possible nucleotides,  $\hat{\mu}/(n-1)$  is the probability of one nucleotide mutating to some other given nucleotide. Since we are modeling a standard genetic code,  $n = 4$ .

## 2.5. THE AFFINITIES $\kappa_i$

If for each Ig genotype  $i$ , we specify an affinity  $\kappa_i$ , then our model is complete. However, specifying the affinities  $\kappa_i$  is not straightforward. The determination of the binding affinity to a given antigen as a function of the Ig primary sequence is a very complicated unsolved problem. In fact, before one can adequately address this problem, one needs to be able to predict an antibody's structure from its primary sequence, and even this problem is notoriously intractable with today's methods. Furthermore, to keep track of individual clones, even restricting oneself to rather small neighborhoods of any given clone, is computationally infeasible. The coding sequence for the variable regions of the Ig molecule is about 660 bases long (Klein,

1990), so there are  $(n-1)^3 \left( \frac{N}{3} \right) > 10^9$  clones to track among just the three-mutant variants. Clearly, some heuristic is needed.

## 2.6. AFFINITY CLASSES

Rather than deal with individual clones, which requires us to solve the problem of assigning affinities to each Ig genotype, we shall move to a different level of description where we lump all clones that have similar affinities into a single category, which we call an *affinity class*. At this new level of description,  $x_i$  will be the concentration of cells in affinity class  $i$ . Some mutations will move cells from one affinity class to another. By keeping track of the dynamics of cells moving between affinity classes we shall be able to account for the changes in affinity that accompany affinity maturation by somatic mutation.

Let us be precise about the definition of affinity class. For example, we could define affinity classes as corresponding to affinities of  $10^4 \text{ M}^{-1}$ ,  $10^5 \text{ M}^{-1}$ ,  $10^6 \text{ M}^{-1}$ , etc, where a clone would be assigned to affinity class  $10^5 \text{ M}^{-1}$  if it had an affinity  $5 \times 10^4 \leq K < 5 \times 10^5$ . Rather than specifying the affinity classes explicitly, we take a more general perspective and let  $\kappa_i$  denote a (dimensionless) affinity class and arrange the affinity classes in ascending order so that  $\kappa_{i+1} > \kappa_i$ . We assume that the affinity distribution is sufficiently broad that clones will fall into each of our assigned classes.

Once affinity classes are defined, we need to specify how mutations will affect the affinity class of a clone. Here we shall call a mutation *silent* if it does not change the



affinity class of the antibody. A mutation that does lead to an affinity class change will be said to be *expressed*. If a mutation is expressed, the affinity change can be to an antibody that loses all ability to bind antigen. We shall treat such cases specially, and call these mutations *lethal*. Let  $p_s$  be the probability that a mutation is silent. Then, given that a mutation is not silent (probability  $1 - p_s$ ), it can either be lethal (probability  $p_L$ ), i.e. lead to no binding with the antigen, or be non-lethal (probability  $1 - p_L$ ). We assume that any single point mutation that is not silent and is not lethal, will only change the affinity by an amount sufficient for the clone to move up one affinity class or move down one class. Clearly, some mutations can cause a large enough change in affinity so that it would correspond to two or more classes. Here we assume that the affinity classes are defined in such a way that such events are sufficiently rare that they can be ignored.

The probabilities of an expressed non-lethal mutation causing a clone to move up or down one affinity class are determined by the relative number of possible clones (not necessarily occupied) in each affinity class. Here, by possible clone, we mean that if one examines all possible V-region sequences, some number of them,  $\mathcal{N}_i$ , will correspond to affinities in class  $i$ . Given that a single expressed non-lethal mutation has occurred in class  $i$ , the probability that it is advantageous, moving to class  $i + 1$  is simply

$$p(i \rightarrow i + 1) = \frac{\mathcal{N}_{i+1}}{\mathcal{N}_{i-1} + \mathcal{N}_{i+1}}. \quad (14)$$

The fraction of V-region sequences that fall into each class is not known, but it seems clear that high affinities are much rarer than low affinities. To make matters explicit we shall assume that the number of V-region sequences falls by a constant factor  $\Lambda$  as the affinity class is increased by one, i.e. we assume that the number  $\mathcal{N}_i = \mathcal{N}_{i-1}/\Lambda$ , so that

$$\mathcal{N}_i = \mathcal{N}_0 \Lambda^{-i}, \quad (15)$$

for some  $\Lambda > 1$ . Under this assumption

$$p(i \rightarrow i + 1) = \frac{1}{1 + \Lambda^2}. \quad (16)$$

Thus, advantageous mutations are rarer than deleterious ones and the fraction of expressed, non-lethal mutations that are advantageous is given by eqn (16). Notice that the result of assuming that the number of sequences in succeeding classes falls geometrically, eqn (15), is that the fraction of advantageous mutations depends only on  $\Lambda$  and is independent of the affinity class  $i$ . In reality, one might expect that advantageous mutations of a high-affinity antibody are less likely than advantageous mutations of a low-affinity molecule. Generalizations of our model that incorporate this assumption would be interesting to pursue.

## 2.7. THE MUTATION MATRIX $[m_{ij}]$ FOR AFFINITY CLASS CHANGES

In section 2.4 we derived the mutation matrix corresponding to a change in genotype from type  $i$  to type  $j$ . For our lumped version of the model in which

affinities are treated in classes, we need to define a mutation matrix  $[m_{ij}]$  that corresponds to a change in affinity from class  $i$  to class  $j$ .

When a cell divides, multiple point mutations are possible. With  $\hat{\mu}$  the probability of mutation per nucleotide per generation,  $\hat{\mu}(1-p_S)$ , is the probability of an expressed mutation per nucleotide per generation. Now, during a single cell division, mutations can change affinity from class  $i$  to class  $j$ . For such transitions, we have

$$m_{ij} = \binom{N}{|i-j|} (\hat{\mu}(1-p_S)(1-p_L))^{|i-j|} \frac{(1-\hat{\mu}(1-p_S))^{N-|i-j|}}{1+\Lambda^{2(j-i)}} \quad (17)$$

and

$$m_{ii} = \{1-\hat{\mu}(1-p_S)\}^N. \quad (18)$$

In deriving this equation we have assumed that each expressed non-lethal mutation corresponds to a change in one affinity class. Thus, a change from class  $i$  to class  $j$  involves  $|i-j|$  expressed non-lethal mutations in  $|i-j|$  out of  $N$  nucleotides, each of which occurs with probability  $\hat{\mu}(1-p_S)(1-p_L)$ , and no expressed mutations in the remaining  $N-|i-j|$  positions.

If we restrict our attention to a small number of classes, so that  $|i-j| \ll N$ , we may rescale  $\mu$ , letting  $N\hat{\mu}(1-p_S) \rightarrow \mu$ . Then to leading order in  $1/N$ , we have a Poisson process, with

$$m_{ij} = \frac{\{\mu(1-p_L)\}^{|i-j|}}{|i-j|!} \frac{e^{-\mu}}{1+\Lambda^{2(j-i)}} \quad (19)$$

and

$$m_{ii} = e^{-\mu}. \quad (20)$$

### 3. Maximization of Total Affinity and Pontryagin's Maximum Principle

The problem as we have discussed it can be framed in the language of optimal control theory (Pontryagin *et al.*, 1962) and can be solved using Pontryagin's maximum principle. We will not state the principle in its most general form, but will use a version specific to our needs.

As we have mentioned, during the course of an immune response the average affinity of the antibodies in the blood serum for the eliciting antigen increases with time. Here we have chosen not to deal explicitly with the antibody itself, but only with the population dynamics of the B-cell clones that ultimately secrete the antibody. In the context of these clones, we must decide what quantity makes a reasonable figure of merit for the purpose of optimization. We choose to maximize at some fixed time  $T$  after the start of the germinal center reaction (GCR) the product of the total population size of B cells specific for the antigen and the average affinity. We call this quantity the *total affinity*

$$A(t) = \sum_i x_i(t) \kappa_i. \quad (21)$$

The reason for choosing this quantity is that both high average affinity and having many B cells participating in the response are important in antigen elimination.

The maximization of  $A(T)$  is done by varying  $\mu(t)$  acting through  $[m_{ij}]$ , as given by eqns (19) and (20). Fortunately, the choice of the function to be maximized does not seem to be critical, for within a wide range of intuitively reasonable functions, the essential features of the solutions are unchanged. In fact, for the systems we have studied, the  $\mu(t)$  obtained by maximizing  $A(T)$  is indistinguishable by our criteria from that obtained by maximizing a generalized  $A_q(T) \equiv \sum_i x_i(T) \kappa_i^q$  as long as  $q \geq 1$ .

To find the optimal  $\mu(t)$ , we introduce, as required by the Pontryagin procedure, the co-state variables  $z_i$  and the Hamiltonian

$$H(x, z; \mu(t)) = \sum_i z_i f_i(x; \mu(t)), \quad (22)$$

where  $f_i(x; \mu(t))$  is given by the right-hand side of (9). The co-state variables (or, more precisely, the quantities  $z_i - \kappa_i$ ) play a role of Lagrange multipliers for the dynamical equations. In terms of the Hamiltonian, the equations of motion are

$$\frac{dx_i}{dt} = \frac{\partial H}{\partial z_i} = f_i(x, \mu(t)) \quad (23)$$

and

$$\frac{dz_i}{dt} = -\frac{\partial H}{\partial x_i} = -\sum_j z_j \frac{\partial f_j}{\partial x_i}. \quad (24)$$

The initial conditions on  $x$  are fixed by the model. The terminal  $x$  values are free to vary. In this case, according to Pontryagin's maximum principle, the terminal conditions on  $z$  must be fixed. For this problem they are  $z_i(T) = \kappa_i$ . Note that at the terminal time,  $H(T) = \partial A / \partial t(T)$ , but  $H$  and  $\partial A / \partial t$  are generally different at other times. Now a necessary condition that  $\mu(t)$  maximize  $A(T)$  as a functional is that  $\mu(t)$  maximize  $H(t)$  pointwise for all  $t$  between 0 and  $T$ , constrained by  $\mu > 0$ .

The algorithm we have constructed to find the optimal  $\mu(t)$  is quite direct. It starts from an initial guess  $\mu_0(t)$  on a discretized time, and improves this by sequentially solving the equations

$$\frac{dx^{(n)}}{dt} = f(x^{(n)}; \mu_n(t)) \quad (25)$$

with

$$x_i^{(n)}(0) = \begin{cases} 2\xi & \text{if } i = 0 \\ 0 & \text{otherwise} \end{cases} \quad (26)$$

and

$$\frac{dz^{(n)}}{dt} = -\sum_j z_j^{(n)} \frac{\partial f_j}{\partial x} (x^{(n)}, \mu_n(t)) \quad (27)$$

with

$$z_i^{(n)}(T) = \kappa_i, \quad (28)$$

and using these results to find  $\bar{\mu}_n(t)$  such that  $\bar{\mu}_n(t)$  maximizes  $H(x^{(n)}, z^{(n)}, \mu(t))$  over  $\mu(t)$ . Then, the new guess is

$$\mu_{n+1}(t) = \rho \mu_n(t) + (1 - \rho) \bar{\mu}_n(t), \quad (29)$$

where  $\rho$  is a damping parameter,  $0 \leq \rho \leq 1$ . Thus, given  $\mu_n(t)$  we find  $x^{(n)}(t)$  by integrating (25) forward from its fixed initial values [note that (25) does not depend

on  $z$ ]. Then, using  $x^{(n)}$  we find  $z^{(n)}$  integrating (27) backward from its fixed initial values. Using  $x^{(n)}$  and  $z^{(n)}$  in the expression for  $H$ , we find the  $\mu(t)$  that maximizes  $H$  at each of the discrete points of time, and call it  $\bar{\mu}_n(t)$ . Finally, we form the new candidate guess  $\mu_{n+1}(t)$  as a linear combination of  $\bar{\mu}_n(t)$  and  $\mu_n(t)$ . This is repeated until some convergence criterion is met. The convergence criterion we use is that  $\mu(t)$  stops changing very much, or more precisely, that for some small positive number  $\varepsilon$ ,

$$\frac{1}{T} \int_0^T \{\mu_n(t) - \bar{\mu}_n(t)\}^2 dt < \varepsilon. \quad (30)$$

### 3.1. SIMPLER MODELS

In addition to the model developed above, which we will solve numerically, we have developed a family of simpler models that admit analytical solution. These models are useful insofar as they provide a partial bridge between this work and that of Agur *et al.* (1991), increase the power of our intuitions for the more complex systems and highlight the details more closely than is possible in a numerical treatment. They allow us to calculate, for example, the relative advantage conveyed by the use of the globally optimal schedule rather than the simpler constant-mutation schedule. The solutions of the simpler models give us expressions for the parameter dependence of important features of their optimal schedules that we may expect to carry over, in approximation, to the full model as well. These models and their solutions are presented and discussed in the Appendix.

### 3.2. PARAMETER VALUES

- $k_p$  We use  $k_p = 4 \text{ d}^{-1}$  corresponding to a cell cycle of about 6 h (Zhang *et al.*, 1988).
- $k_d$  Cell death rates have not been measured as accurately as have proliferation rates. We use  $k_d = 4 \text{ d}^{-1}$ , consistent with Cohen *et al.* (1992). This high rate is due to apoptosis (see below) and represents a combined rate for the receipt of apoptotic signals and the delay before death occurs.
- $p_L$  Only a small part of the variable region gene actually codes for one of the complementarity determining regions (about 25%). The rest codes for framework regions, mutations to which have a high chance of being lethal. Therefore, we take  $p_L = 0.5 \text{ mutation}^{-1}$ , in close accord with previous model estimates (Litwin, S., personal communication).
- $\kappa_i$  We take the (dimensionless) affinities to increase geometrically, using values taken from the most complete site-directed mutagenesis experiments known to us at this time (Brown *et al.*, 1992):  $\kappa_i = 7.5(7.5)^i$ . The initial seed clone has  $\kappa = 7.5$ .
- $\xi$  The threshold  $\xi$  sets the scale for the populations, since  $\xi$  is equivalent to one cell. In our dimensionless units we set  $\xi = 10^{-4}$  so that using the above values for the death and proliferation rates, the maximum population of a high-affinity clone, from eqn (11), is about 3, which is equivalent to  $3 \times 10^4$  cells, reasonable for the (centroblast) population of a GC (Liu *et al.*, 1992).

- $\Lambda$  We let  $\Lambda$ , which determines the ratio of advantageous to non-lethal but deleterious mutants, vary between 1 and about 70. An estimate given by Manser (1991) is roughly equivalent to  $\Lambda = 30$ .
- $T$  Our approach is to maximize some quantity of merit at a fixed time after the onset of the GCR. For this time we take a typical duration of a primary response and GCR,  $T = 14$  days.

Our initial conditions are  $x_i(0) = 0$  for all clones but the seed clone. For the seed clone,  $x_1(0) = 2\xi$ , corresponding to an initial population of two cells. This is consistent with the known oligoclonality of GCs and their initially small populations (Kroese *et al.*, 1987).

#### 4. Results

The essential results are shown in Figs 1 and 2, where we have used  $\Lambda = 30$  and  $\Lambda = 40$ , respectively. Starting from an initial guess of  $\mu_0(t) = \text{constant}$ , the algorithm produced a *phasic* candidate optimum mutation schedule with four (Fig. 1) or three




FIG. 1. Panel (d) shows  $\mu(t)$ , the optimal rate of expressed mutations per generation per Ig coding region, as a function of time over the course of 14 days. The parameters are as given in the text and with  $\Lambda = 30$ . This  $\mu(t)$  was obtained by iterating the algorithm described in the text from a starting constant rate  $\mu_0(t) = 0.1$ . (a) Shows the log of the total affinity,  $A(t)$ . Panel (b) plots  $\log \alpha$ , the (dimensionless) concentration of free antigen, showing its depletion to successively smaller values as high-affinity clones appear and grow. Panel (c) gives  $\log x$ , the population sizes for all the affinity classes used in the simulation. Note that in these plots, the population variable  $x$  has been normalized by division by  $\xi$  so that  $x$  represents actual number of cells. The lines are labeled by the affinity class index,  $i$ , with  $\kappa_i = 7.5(7.5)^i$ . The seed class corresponds to  $i = 0$ . Of the lower-affinity classes ( $i = -1, -2, -3$ ), only the first is capable of proliferation at the initial free antigen level, so they grow exclusively by mutation from higher classes. The lower boundary of this plot corresponds to  $x = \xi$ , or the one-cell level.

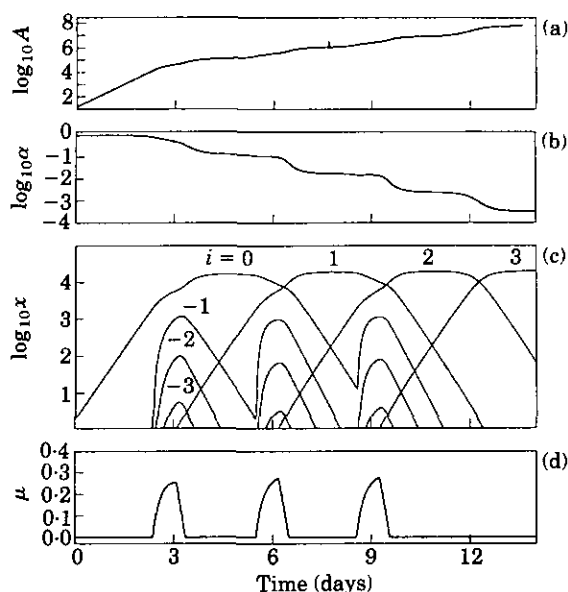


FIG. 2. The same as Fig. 1, but with  $\Lambda = 40$ , i.e. advantageous mutations are rarer.

(Fig. 2) bursts of high mutation rate each lasting about 0.8 days, between longer periods of about 2 days during which there is no mutation at all. Although Pontryagin's maximum principle provides only a necessary condition for optimality, we have confidence that our phasic solutions actually represent global optima. These solutions make sense when compared to the analytic solutions obtained for simplified versions of this model (see Appendix). Furthermore, for these parameter values, the solutions appear to be unique, in that the only other fixed point of our algorithm we have found is the obviously suboptimal  $\mu(t) \equiv 0$  (this solution is unstable). Regardless of our initial guess  $\mu_0(t)$ , as long as it is not identically zero, we find the same fixed point†.

Referring to Figs 1 or 2, interpretation of the optimal solution is straightforward. At the onset of the GCR, the rate of growth of the seed clone [(c),  $i = 0$ ] is maximized by keeping mutation turned off. When it reaches a population size roughly proportional to  $\xi\Lambda^2$  (see Appendix), the clone now has enough cells to virtually assure that if each cell produces one mutant daughter, at least one of these will be advantageous. At this point, it switches on mutation, the growth rate of the seed clone declines and a large number of low-affinity clones ( $i = -1, -2, \dots$ ) are produced. When the higher-affinity variant finally does appear, mutation is no longer desirable, since one now wants to maximize the rate of growth of this new clone. This process is repeated for the duration of the GCR. At no time does the total affinity  $A$

† The stability of fixed points depends on the algorithm parameter  $\rho$ . An unstable fixed point can usually be stabilized by making  $\rho$  closer to 1, but we have no assurance that this is always true; as  $\rho$  gets close to one, the convergence rate falls rapidly, and it becomes impractical to find the optimal solution with our algorithm.

[[a]] decrease. Note the step-wise depletion of antigen [[b]] as higher-affinity clones are produced, illustrating the stringent competition for antigen occurring during this time.

Decreasing the proportion of advantageous mutations (Fig. 2) affects the peak width very little, and increases the duration of the interburst interval from 1.9 days to about 2.2 days. This is roughly in accord with the expectations derived from the analysis of a simplified model (see Appendix), that the interburst interval is proportional to  $\log(\Lambda^2)$  and the peak width is independent of  $\Lambda$ . The peak mutation rate is roughly the same in both cases, about 0.3. Assuming a ratio of silent-to-expressed mutations of 1:2.9 (Jukes & King, 1979) and a coding sequence 660 bases long (Klein, 1990), these rates are roughly equivalent to  $6.1 \times 10^{-4} \text{ bp}^{-1} \text{ generation}^{-1}$ , within the estimated bounds given elsewhere (Berek & Milstein, 1987). The average mutation rates over the total 14-day period are 0.0425 corresponding to  $8.7 \times 10^{-5} \text{ bp}^{-1} \text{ generation}^{-1}$  (Fig. 1) and 0.055 corresponding to  $1.1 \times 10^{-4} \text{ bp}^{-1} \text{ generation}^{-1}$  (Fig. 2).

While certain details of the solutions change as parameters vary, the crucial point is that for *all* reasonable parameter values, *the optimal solutions are phasic*.

#### 4.1. PHASIC $\mu(t)$ IS NOT AN ARTIFACT OF COARSE GRAINING

It may be objected that the phasic solution could be an artifact of our coarse-graining of the affinity distribution. It is important to know that this is not the case. To show this, we have "added back" some fine structure to the affinity distribution of the model. With this modification, each single point-mutation can reach *two* affinity classes forward or backward, rather than just one. Given that a mutation to one of the two higher-affinity mutants has occurred, we assign probability  $p_j$  to the event that it jumps to the higher of the two classes. The point of this is to show that it may be advantageous to *skip* classes. Thus, for example, if in our original coarse-grained model each affinity class corresponded to a change in affinity by a factor of 10, and we now add new classes so that changes of affinity by a factor of 5 are being modeled, we want to know if the optimal solution will always lead to changes in affinity by a factor of 5. If this is the case, then we should add more classes until we have enough classes that it is no longer optimal to always go to the next highest class. If it is always best to advance just to the first available higher-affinity class and no further during a burst of mutation, then coarse graining will lead to spurious results. If, on the other hand, it is sometimes worthwhile to leave mutation on and wait for a still more advantageous variant to appear, then we feel that our discretization of the affinity distribution into classes will not lead to spurious results. This is exactly what we find. The higher of the two accessible higher-affinity clones can be considerably rarer than the lower one, yet the optimal solution is to wait until the higher class becomes occupied before switching mutation off and shifting to "growth mode". This is illustrated in Fig. 3, which shows the optimal schedule for a system with  $p_j = 0.1$ . For this figure, we have used  $\Lambda^2 = 159.1$  so that the population factor for the higher one-step class is exactly the same as that for the system in Fig. 2. Similarly, the affinities are  $\kappa_i = 7.5(7.5)^{i/2}$ . In each of the mutation bursts one can see the trajectory




FIG. 3. For this figure, we have allowed one-step access to *two* affinity classes in either direction from any given class. The relative probability for jumping to the higher-affinity of these two classes (solid line) is 0.1/0.9. Each episode of mutation results in the occupation of both higher-affinity classes, thus illustrating the point that the system is not simply taking the smallest steps possible, but weighing the size of the steps (the relative advantage) against the waiting time to take them. The seed class and all the even-numbered classes are shown in solid lines. The odd-numbered classes appear as dashed lines. The populations for classes lower than the seed class have been omitted for clarity. We have used  $\Lambda^2 = 159.1$  so that the probability of stepping to the higher-affinity one-mutant class is equal to the probability, in Fig. 2, of stepping to the single higher-affinity one-mutant class. Similarly, we use  $\kappa_i = 7.5(7.5)^{i/2}$ .

of the lower of the pair of higher-affinity classes cross threshold before the higher one. It is clear that the additional waiting time required to populate the higher-affinity class is short enough to make this the desired solution. The point is that it is not necessarily the best choice to simply take the easiest step every time. If this were the case, our approximation methods would nullify our results. As it is, though, the advantage gained is weighed against the waiting time, so that it is often best to take a larger step even though it takes longer. It would be an interesting problem to find the optimal  $\mu(t)$  for the case of a *continuous* distribution of affinities, though we believe that the "true" affinity distribution is far from continuous, and that therefore our approximation is genuinely "coarse-graining" and not "discretization".

#### 4.2. COMPARISON WITH SUBOPTIMAL SCHEDULES

The kind of fine control over the instantaneous mutation rate implied by our method seems unrealistic in a biological system; a less perfectly controlled switching mechanism is more likely to be operative. Therefore, it is important to know how sensitive our merit function  $A$  is to perturbations in the optimal mutation schedule. Put another way, if instead of the optimal phasic schedule one were to use a different






FIG. 4. The solid line shows  $A(T)$ , with  $T = 14$  days, our figure of merit, as a function of the amplitude of a square-wave mutation schedule whose switching times are precisely those of the optimal solution shown in Fig. 1. As in Fig. 1, we have  $\Lambda = 30$ . The dashed line shows  $A(T)$  for a constant-rate schedule, where the constant mutation rate,  $\hat{\mu}$  in  $\text{bp}^{-1} \text{gen}^{-1}$ , is given on the x-axis. (Again,  $x$  has been normalized by division by  $\xi$  so as to represent actual population numbers.)

phasic schedule or a constant mutation rate, would one see much difference in the performance of the system? We address both of these cases by generating two suboptimal schedules, and using these in the numerical integration of eqn (9). In addition, we will show that phasic mutation schedules are less sensitive to fluctuations in the maximum mutation rate than are constant-rate schedules. This is clearly a great advantage in itself.

Figure 4 shows the results of using suboptimal mutational schedules. In the first schedule (solid line) we take the globally optimal  $\mu(t)$  and replace each burst with a square pulse of variable amplitude. The second schedule (dashed line) uses a constant mutation rate. The first point to note is that the square wave mutation schedule yields a maximum value of  $A(T)$  which is within 1% of the global maximum, and is four times larger than the maximum achieved by the optimal constant rate. Second, values of  $A(T)$  near maximum are achieved over a wide range of  $\mu$  values when the square wave mutation schedule is used. Finally, the constant rate schedule provides a net improvement over simple mutation-free proliferation for  $\hat{\mu}$  ranging from just positive to a value of about  $9 \times 10^{-4} \text{ bp}^{-1} \text{ generation}^{-1}$ , where mutation becomes severely disadvantageous. In contrast, the phasic schedule provides improvement up to  $\hat{\mu} = 4 \times 10^{-3} \text{ bp}^{-1} \text{ generation}^{-1}$ , and does not become deleterious until well beyond this point.

## 5. Discussion

Although mutation during affinity maturation is certainly regulated (French *et al.*, 1989; Berek *et al.*, 1991), the notion that a phasic mutation schedule is actually

implemented by a set of switches that turn mutation on and off at precise times seems unlikely. So rather than look for esoteric molecular mechanisms that might lead to such unusual switching behavior, we will suggest that movement of cells into and out of a "mutational compartment" in which mutation occurs can, in fact, be an implementation of a phasic mutation schedule.

### 5.1. ANATOMY AND KINETICS OF THE GERMINAL CENTER REACTION

In our model we have assumed that mutation and proliferation occurs in a single compartment. The germinal center (GC) has been implicated as the site of somatic mutation (Jacob *et al.*, 1991). The anatomy of the GC has been well-studied in the past few years (MacLennan, 1991; Leanderson *et al.*, 1992; Liu *et al.*, 1992; Nossal, 1992). Rather than being a single compartment, it is clearly multi-compartmental (see Fig. 5). It is composed of several distinguishable regions within which the B-cell populations display distinct phenotypes. In addition, there are follicular dendritic cells (FDC) (Tew *et al.*, 1990), which populate the lymphoid follicles even when no GCRs are occurring. These cells bind antigen-antibody complexes with very high affinity for long periods of time and have been implicated in the maintenance of immune memory (Gray & Skarvall, 1988; Gray & Leanderson, 1990; Gray, 1992). Recent *in vitro* studies have verified the importance of FDCs to the continued proliferation of GC B cells (Kosco *et al.*, 1992) as well as that of T cells, though their role in that process is less clear.

The GCR starts 1 or 2 days after immunization and is marked by exponential growth of B cell blasts in the spaces of the FDC network. In the next stage, centroblasts appear and begin to form the dark zone of the GC (Fig. 5). These cells divide and mutate rapidly, but surprisingly they express no surface Ig and hence are unlikely to be subject to affinity-based selection. The daughter cells of centroblasts become centrocytes and form the light zone surrounding the dark zone (Fig. 5). In the light zone, the FDC network is particularly dense and the centrocytes here re-express surface Ig. The light zone, presumably, is where selection occurs. Centrocytes that do not bind antigen are lost through apoptosis (Liu *et al.*, 1989), a self-destruction process characterized by fragmentation of nuclear material in the regions between nucleosomes (Cohen *et al.*, 1992). Those centrocytes that do bind antigen may become plasma cells that do not divide but secrete Ig copiously, they may differentiate into memory B cells that remain at rest until restimulated with antigen, or they may follow another, as yet unverified, path that can lead to further mutation as suggested by our phasic mutational schedule.

Liu *et al.* (1992) point out that there is another anatomically distinct region in the CG known as the *outer zone*, which surrounds the light zone (Fig. 5). Some cells leaving the light zone after antigen-based selection might re-enter the dark zone through this outer zone. Re-entry into the dark zone could lead to another round of mutation followed by selection in the light zone. A second possibility, suggested to us by R. W. Anderson, is that memory cells exiting the germinal center may colonize new germinal centers. Sequential colonization also realizes a phasic mutation schedule, and is supported by recent experimental evidence (Vonderheide & Hunt,




FIG. 5. A schematic diagram of the germinal center showing the dark zone of centroblasts, the light zone of centrocytes and the outer zone. The known paths of B cells through the GC are shown as heavy arrows. The dark zone is populated by centroblasts, which proliferate and mutate rapidly, but express no surface Ig. They then enter the light zone as centrocytes where they re-express surface Ig and undergo selection. Those that fail to bind antigen die through apoptosis. Those that do bind antigen pass into the circulation where they become plasma cells or memory cells. Shown in dashed arrows are three proposed alternative pathways for positively selected cells exiting the light zone. First, there is re-entry through the outer zone into the dark zone. Second, there is re-entry through the circulation. Finally, there is sequential colonization, through the circulation, of new germinal centers.

1992) indicating that memory B cells act as germinal center precursors more readily than do other B-cell types. A third possible implementation of a phasic schedule could involve cells entering the circulation after selection in the light zone. Such cells might then re-enter the dark zone of the GC from the circulation. If these re-entrant cells divide during their transit back to the dark zone so that selected populations

expand, then this scenario also implements a phasic mutation schedule through spatial compartmentalization. As pointed out to us by J. Ahouse, this scenario could give rise to multi-clonal germinal centers. More specific assumptions are required to evaluate the expected diversity. Given the known GC oligoclonality (Kroese *et al.*, 1987; Jacob *et al.*, 1991), once these assumptions are established, this scenario will be subject to test against the present data and against more precise measurements of GC clonal diversity possible in future experiments.

The concept of re-entry into a mutational compartment seems unavoidable (Steele *et al.*, 1991) if one considers the large numbers of mutations observed in experimental preparations together with the known anatomy of the germinal center. The number of mutations appearing in the course of the primary response is quite variable and seems to depend on the antigen used and other experimental conditions. Values for the number of mutations in the late primary response run from about three (Berek & Milstein, 1987; Gearhart & Levy, 1991) to as high as ten or more (Sharon *et al.*, 1989) in each of the light and heavy chain variable regions. Manser (1991) quotes an average number of six mutations per V-region gene. Without re-entry, a given cell would undergo one or more divisions with mutation as a centroblast but would not be subject to selection. Its daughter cells, the centrocytes, which re-express surface Ig would then experience selection, but those chosen would not experience further mutation. Mutation would need to be accomplished all in one shot. For example, generation of a high-affinity Ig V-region genes containing six mutations would require generation of a pool of candidate six-mutant V genes large enough to be likely to contain one of high affinity. This seems highly unlikely. Thus our results, suggesting that a phasic mutation schedule is optimal, help explain how a multi-compartmental GC can perform affinity maturation with high efficiency.

## 6. Conclusion

We have examined somatic hypermutation in the immune response of B cells as a problem in optimal control. When modeled as a single-compartment system, the mutation schedule  $\mu(t)$  that maximizes the total affinity (or many other reasonable figures of merit) has periods of high mutation rates interspersed between periods of mutation-free growth, i.e. is phasic. We suggest that a phasic mutation schedule would arise if there were one or more "mutational compartments" in which cells entered periodically. The germinal center is the obvious candidate for a mutational compartment since Jacob *et al.* (1991) have observed somatic mutants in the GC but not elsewhere in the spleen. Further, a phasic mutation schedule arises as a natural consequence of the multi-compartmental structure of the germinal center if there exists a re-entrant pathway for positively selected mutants as suggested by Liu *et al.* (1992), or if the mutants which exit one CG colonize another germinal center.

We thank J. Ahouse, R. W. Anderson, S. Litwin and L. Segel for stimulating discussions and for critical readings of an earlier version of this manuscript. This work was supported by Los Alamos National Laboratory LDRD Program, grants from the National Institutes of Health (AI28433 and RR06555), and the Santa Fe Institute through their Theoretical Immunology Program.

## REFERENCES

- AGUR, Z., MAZOR, G. & MEILUSON, I. (1991). Maturation of the humoral immune response as an optimization problem. *Proc. R. Soc. Lond.* **245**, 147–150.
- BEREK, C., BERGER, A. & APEL, M. (1991). Maturation of the immune response in germinal centers. *Cell* **67**, 1121–1129.
- BEREK, C. & MILSTEIN, C. (1987). Mutation drift and repertoire shift in the maturation of the immune response. *Immunol. Rev.* **96**, 23–41.
- BEREK, C. & MILSTEIN, C. (1988). The dynamic nature of the antibody repertoire. *Immunol. Rev.* **105**, 5–26.
- BROWN, M., STENZEL-POORE, M., STEVENS, S., KONDOLEON, S. K., NG, J., BÄCHINGER, H. & RITTENBERG, M. (1992). Immunologic memory to phosphocholine keyhole limpet hemocyanin: recurrent mutations in the  $\lambda 1$  light chain increase affinity for antigen. *J. Immunol.* **148**, 339–346.
- BURNET, F. (1959). *The Clonal Selection Theory of Acquired Immunity*. Cambridge: Cambridge University Press.
- CARTER, R. & FEARON, D. (1992). CD19: lowering the threshold for antigen receptor stimulation of B lymphocytes. *Science* **256**, 105–107.
- CHUA, M., GOODGALL, S. & KARUSH, F. (1987). Germ-line affinity and germ-line variable-region genes in the B cell response. *J. Immunol.* **138**, 1281–1289.
- COHEN, J., DUKE, R., FADOK, V. & SELLINS, K. (1992). Apoptosis and programmed cell death in immunity. *Ann. Rev. Immunol.* **10**, 267–293.
- FRENCH, D., LASKOV, R. & SCHARFF, M. (1989). The role of somatic hypermutation in generation of antibody diversity. *Science* **244**, 1152–1157.
- GEARHART, P. & LEVY, N. (1991). Kinetics and molecular models for somatic mutation in immunoglobulin genes. In: *Somatic Mutations in V-Regions* (Steele, E. J., ed.) pp. 29–40. Boca Raton, FL: CRC Press.
- GRAY, D. & LEANDERSON, T. (1990). Expansion, selection and maintenance of memory B-cell clones. *Curr. Topics Microbiol. Immunol.* **159**, 1–17.
- GRAY, D. & SKARVALL, H. (1988). B-cell memory is short-lived in the absence of antigen. *Nature, Lond.* **336**, 70–73.
- GRAY, D. (1992). The dynamics of immunological memory. *Sem. Immunol.* **4**, 29–34.
- JACOB, J., KELSEY, G., RAJEWSKY, K. & WEISS, U. (1991). Intracloal generation of antibody mutants in germinal centers. *Nature, Lond.* **354**, 389–392.
- JUKES, J. & KING, J. (1979). Evolutionary nucleotide replacements in DNA. *Nature, Lond.* **281**, 605–607.
- KLEIN, J. (1990). *Immunology*. Boston, MA: Blackwell Scientific.
- KOSCO, M., PFLUGELDER, E. & GRAY, D. (1992). Follicular dendritic cell-dependent adhesion and proliferation of B cells *in vitro*. *J. Immunol.* **92**, 2331–2339.
- KROESE, F., WUBENNA, A., SEIJEN, H. & NIEUWENHUIS, P. (1987). Germinal centers develop oligoclonally. *Eur. J. Immunol.* **17**, 1069–1072.
- LEANDERSON, T., KÄLLBERG, E. & GRAY, D. (1992). Expansion, selection and mutation of antigen-specific B cells in germinal centers. *Immunol. Rev.* **126**, 47–61.
- LIU, Y.-J., JOHNSON, G., GORDON, J. & MACLENNAN, I. (1992). Germinal centres in T-cell dependent antibody responses. *Immunol. Today* **13**, 17–21.
- LIU, Y.-J., JOSHUA, D., WILLIAMS, G., SMITH, C., GORDON, J. & MACLENNAN, I. (1989). Mechanism of antigen-driven selection in germinal centers. *Nature, Lond.* **342**, 929–931.
- MACLENNAN, I. (1991). The centre of hypermutation. *Nature, Lond.* **354**, 352–353.
- MANSER, T. (1991). The efficiency of antibody affinity maturation: can the rate of B-cell division be limiting? *Immunol. Today* **11**, 305–308.
- MCKEAN, D., HUPPI, K., BELL, M., STAUDT, L., GERHARD, W. & WEIGERT, M. (1984). Generation of antibody diversity in the immune response of BALB/c mice to influenza hemagglutinin. *Proc. natn. Acad. Sci. U.S.A.* **81**, 3180–3184.
- NOSSAL, G. (1992). The molecular and cellular basis of affinity maturation in the antibody response. *Cell* **68**, 1–2.
- PERELSON, A. S. & OSTER, G. F. (1979). Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self–nonself discrimination. *J. theor. Biol.* **81**, 645–670.
- PERELSON, A. S., MIRMIRANI, M. & OSTER, G. F. (1976). Optimal strategies in immunology, i. b-cell differentiation and proliferation. *J. math. Biol.* **3**, 325–367.
- PERELSON, A. S., MIRMIRANI, M. & OSTER, G. F. (1978). Optimal strategies in immunology, ii. b memory cell production. *J. math. Biol.* **5**, 213–256.

- PONTRYAGIN, L., BOLTYANSKII, V., GAMKRELIDZE, R. & MISCHENKO, E. (1962). *The Mathematical Theory of Optimal Processes*. New York: Wiley.
- SHARON, J., GEFTER, M., WYCOCKI, L. & MARGOLIES, M. (1989). Recurrent somatic mutations in mouse antibodies to p-azophenylarsonate increases affinity for hapten. *J. Immunol.* **142**, 596–601.
- SILVERSTEIN, A. M. (1989). *A History of Immunology*. San Diego, CA: Academic Press.
- STEELE, E., POLLARD, J., TAYLOR, L. & BOTH, G. (1991). Evaluation of possible mutator mechanisms active on mammalian variable region genes. In: *Somatic Mutations in V-Regions* (Steele, E. J., ed.) pp. 29–40. Boca Raton, FL: CRC Press.
- STEELE, E. J. (1991). *Somatic Mutations in V-Regions*. Boca Raton, FL: CRC Press.
- Tew, J., Kosco, M., Burton, G. & Szakal, A. (1990). Follicular dendritic cells as accessory cells. *Immunol. Rev.* **117**, 185–211.
- VONDERHEIDE, R. & HUNT, S. (1992). Comparison of IgD<sup>+</sup> and IgD<sup>−</sup> thoracic duct B lymphocytes as germinal center precursor cells in the rat. *Int. Immunol.* **3**, 1273–1281.
- WEIGERT, M., CESARI, I., YONKOVITCH, S. & COHN, M. (1970). Variability in the light chain sequences of mouse antibody. *Nature, Lond.* **228**, 1045–1047.
- ZHANG, J., MACLENNAN, I., LIU, Y.-J. & LANE, P. (1988). Is rapid proliferation in B centroblasts linked to somatic mutation in memory B cell clones? *Immunol. Lett.* **18**, 297–300.

## APPENDIX

### A1. Simplified Models

In order to develop an intuition about the results contained in the main body of this paper, and to make contact with the previous work of Agur *et al.* (1991), we present here three models that are sufficiently simple that analytical solutions may be found.

For the simplified models we assume that all disadvantageous mutants are non-binding lethals. Furthermore, we posit that the affinity of the one-step advantageous mutant is much greater than that of the seed clone, so that we are interested only in minimizing the time it takes to generate one such mutation. With  $x_0$  the seed clone and  $x_1$  the advantageous one-step mutant, the equations of motion can be written as

$$\frac{dx_0}{dt} = x_0 \{-k_d + h(x_0)(k_d + k_p(1 - 2\mu))\} \quad (\text{A.1})$$

and

$$\frac{dx_1}{dt} = \frac{2\mu}{\Lambda^2} \phi(\mu) k_p x_0 h(x_0). \quad (\text{A.2})$$

The three models that we consider differ in the choice of functions  $h$  and  $\phi$ .

Note that the terms missing from these equations in comparison with (9) are the terms multiplied by  $\theta_2$ . But since we start with only  $x_0$  populated, these terms are zero until  $x_1$  reaches threshold, and this is, by definition, the terminal event, so they need not be included. Similarly, in the simplified models, the stimulation level  $h_0$  depends only on  $x_0$  since all clones other than the seed clone are negligible, either because they exist in very small concentrations (the higher-affinity clones) or, because by assumption, they do not bind antigen (for the lower-affinity clones,  $h$  is small). Therefore, the sum in (8) gets its only non-negligible contribution from the  $j = 1$  term, and  $h_0 = h_0(x_0)$ .

The function we are maximizing is the negative of  $T$ , the terminal time. The Hamiltonian picks up a constant term to reflect this choice (Pontryagin *et al.*, 1962) and is now given by

$$H = -c + zx_0(-k_d + h(x_0)(k_d + k_p(1 - 2\mu))) + \frac{2\mu}{\Lambda^2} \phi(\mu)k_ph(x_0)x_0. \quad (\text{A.3})$$

The boundary conditions are  $x_0(0) = \xi$ ,  $x_1(0) = 0$ ,  $x_1(T) = \xi$  and  $z(T) = 0$ ;  $c$  is a constant to be determined.

#### A1.1. MODEL 1

The simplest model corresponds most closely to that developed by Agur *et al.* (1991) and is given by  $\phi(\mu) = 1$  and a constant stimulation level,  $h(x_0) \equiv h^0$ . These simplifications are appropriate in the limit of antigen excess and small mutation rates, but we will not enforce this latter condition. Antigen excess is equivalent, in our dimensionless variables, to small  $x_0$ . In the limit where  $x_0 = 0$ , (8) gives us  $h^0 = \kappa_0/(1 + \kappa_0)$ .

If we insist that  $\mu$  be constant, (A.1) and (A.2) can be easily solved, giving

$$x_0(t) = \xi \exp(k_\mu t) \quad (\text{A.4})$$

and

$$x_1(t) = \frac{2\mu k_p h^0}{k_\mu \Lambda^2} (x_0(t) - \xi) \quad (\text{A.5})$$

where

$$k_\mu \equiv k_d(1 - h^0) + k_ph^0(1 - 2\mu).$$

We are interested in the time  $T$  required to reach the threshold condition  $x_1(T) = \xi$ . Inserting this into (A.4) and (A.5) gives

$$T(\mu = \text{constant}) = \frac{1}{k_\mu} \log \left( 1 + \frac{\Lambda^2 k_\mu}{2\mu k_p h^0} \right). \quad (\text{A.6})$$

If we minimize this expression over  $\mu$ , we get, for  $\Lambda$  large

$$\mu_{\text{opt}} = \frac{k_0}{2k_ph^0 \log(\Lambda^2)} \quad (\text{A.7})$$

and

$$T_{\min}(\mu = \text{constant}) = \frac{1}{k_0} (\log \Lambda^2 + \log \log(\Lambda^2)). \quad (\text{A.8})$$

Here,  $k_0 \equiv k_\mu + 2\mu k_p h^0$  is the growth rate when  $\mu = 0$ .

The globally optimal solution can also be found exactly. This is a so-called *bang-bang* solution, in which  $\mu$  switches discontinuously from its minimum value 0, to its maximum value  $\mu_{\max}$  when the co-state variable passes through the value  $z = 1$  from above:  $\mu(t) = \mu_{\max} \theta(1 - z(t))$ . Using the fact that (Pontryagin *et al.*, 1962)  $H(t) = 0$  for all  $0 \leq t \leq T$  and the terminal condition we calculate the switching time,

$$t_1 = \frac{1}{k_0} \log(\Lambda^2) \quad (\text{A.9})$$

and the terminal time

$$T = t_1 + \frac{1}{k_\mu} \log \left( \frac{k_0}{2\mu_{\max} k_p h^0} \right). \quad (\text{A.10})$$

This solution directs the seed clone to grow without mutation until it reaches the size  $\xi\Lambda^2$ , and then to turn on mutation at the maximum rate. The duration of the burst,  $\tau \equiv T - t_1$ , is a decreasing function of  $\mu_{\max}$  and vanishes as  $\mu_{\max} \rightarrow \infty$ , giving

$$T_{\min}(\text{global}) = \frac{1}{k_0} \log(\Lambda^2). \quad (\text{A.11})$$

Of course, this limit is artificial; it is only advantageous because we neglected, in (19), the exponential factor that cuts off  $m_{12}$  for large  $\mu$ . A more reasonable value to take for  $\mu$  is that for which  $k_\mu$  vanishes, in which case,

$$T_{\min} = \frac{1}{k_0} (\log(\Lambda^2) + 1). \quad (\text{A.12})$$

Using  $\Lambda^2 = 900$ ,  $T_{\min}(\mu = \text{constant})$  is about 28% longer than  $T_{\min}(\text{global})$  and 24% longer than the "practical" minimum  $T_{\min}$ .

#### A1.2. MODEL 2

We may introduce some non-linearity by dropping the antigen excess condition. The global optimum solution above required that the seed clone grow to the size  $\xi\Lambda^2$  before switching on mutation. This may not be reasonable when  $\Lambda^2$  becomes large. To take this into account, let  $h(x_0) = h^0(1 - x_0/\chi)$ . This logistic term can be obtained as an approximation of (8) by expanding  $h_0(x_0)$  (recall that  $h_0$  depends only on  $x_0$ ) in a power series and neglecting terms of order higher than the first. Or one can choose  $\chi$  so that the asymptotic value of  $x_0$  agrees with (11).

Since  $\mu$  still enters linearly, the optimal solution is still bang-bang. Again, the differential equations can be solved analytically. If the switching time is  $t_1$ , then

$$x(t_1) = \frac{\chi_0 \xi e^{k_0 t_1}}{\chi_0 + \xi(e^{k_0 t_1} - 1)} \quad (\text{A.13})$$

where  $\chi_0 \equiv \chi k_0 / (k_0 + k_d)$ . The final value attained is

$$x(T) = \frac{\chi_\mu x(t_1) e^{k_\mu \tau}}{\chi_\mu + x(t_1)(e^{k_\mu \tau} - 1)} \quad (\text{A.14})$$

where  $\chi_\mu \equiv \chi k_\mu / (k_\mu + k_d)$ . The optimal solution is given by the pair of equations

$$k_0 x(t_1) \left( 1 - \frac{x(t_1)}{\chi_0} \right) = 2\mu k_p h^0 x(T) \left( 1 - \frac{x(T)}{\chi_\mu} \right) \quad (\text{A.15})$$

and

$$\frac{k_\mu \Lambda^2 \xi}{2\mu k_p h^0 \chi_\mu} = \left( 1 - \frac{\chi_\mu}{\chi} \right) \log \left( \frac{\chi_\mu - x(t_1)}{\chi_\mu - x(T)} \right) + \frac{1}{\chi} (x(T) - x(t_1)) \quad (\text{A.16})$$






FIG. A1. This plot shows the relationship between the terminal time  $T$  and  $\Lambda^2$ , whose reciprocal gives the frequency of advantageous mutants. The solid line is for simple model 2, with  $\mu_{\max} = 0.5$  and the dashed line is for a constant  $\mu$  model, where  $\mu$  has been chosen ( $\mu = 0.065$ ) to minimize  $T$  at  $\Lambda^2 = 900$ . All figures for the Appendix were generated using parameters as given in the text unless noted.

subject to obvious consistency conditions. If  $\mu_{\max}$  is large enough to make  $k_{\mu}$  negative, this solution only exists for  $\Lambda^2$  not too large. For  $\Lambda^2$  large, the solution involves multiple switching events and will not be treated here. For physiologically realistic parameter values, the above solution is the correct one.

Figure A1 provides a comparison of the optimal solution of this model with one for which the mutation rate is kept constant. Shown is  $\Lambda^2$  plotted against terminal time  $T$ , for the optimal solution (solid line) and the  $\mu$ -constant solution. For the latter case, we have chosen  $\mu$  so as to minimize  $T$  at  $\Lambda^2 = 900$  (the value we have been using in the more complicated model). For the optimal solution model we have used  $\mu_{\max} = 0.5$ . This is a reasonable value, though increasing it will give us a marginally better solution. Note that the advantage enjoyed by the optimal solution actually becomes greater as  $\Lambda^2$  increases. The terminal point on the optimal solution curve is where the single switch solution fails and a more complicated multiple switch solution is required.

### A1.3. MODEL 3

A third simple model is defined by returning to the antigen excess approximation, but allowing for non-linearities in the  $\mu$  terms. In particular, let  $\phi(\mu) = 1 - \mu$ . This corresponds to a small  $\mu$  expansion of (19) and (20), in which the first two terms of each  $\mu$  expression are kept. In contrast to the findings of Agur *et al.* (1991), the optimal-control solution is no longer bang-bang. Instead,  $\mu$  changes continuously

between 0 and 1:

$$\mu = \frac{1}{2}(1-z)\theta(1-z)\theta(1+z) + \theta(-1-z) \quad (\text{A.17})$$

although we find that we do not reach the upper switch point. Substituting this into the expression for  $dz/dt$ , we get

$$\frac{dz}{dt} = -\frac{1}{2}k_p h^0 - (k_0 - k_p h^0)z - \frac{1}{2}k_p h^0 z^2 \quad (\text{A.18})$$

for times  $t$  greater than the switching time  $t_1$ , which is to be determined. Using the fact that  $z(t_1) = 1$ , the solution of this equation is

$$z(t_1 + t) = \frac{q(1 - q^*) + q^*(q - 1)e^{-i\omega t}}{1 - q^* + (q - 1)e^{-i\omega t}} \quad (\text{A.19})$$

where  $q$  and  $q^*$  are the complex conjugate roots of the quadratic equation obtained by setting  $dz/dt = 0$  in (48). The frequency  $\omega$  is

$$\omega = \frac{2k_0 \Re q + k_p h^0(1 + |q|^2 - 2\Re q)}{2\Im q}. \quad (\text{A.20})$$

The terminal boundary condition  $z(T) = 0$  gives

$$\tau = \frac{i}{\omega} \log \left( \frac{q(q^* - 1)}{q^*(q - 1)} \right). \quad (\text{A.21})$$




FIG. A2. The correction term  $-\log(2k_0 k_p h^0 I)$  for simple model 3, against  $h^0$ , the initial stimulation level for the seed clone (see text).




FIG. A3. The terminal time  $T$  and the mutation-on time  $\tau$  against  $h^0$  for simple model 3.




FIG. A4. The shape of the trajectory  $\mu(t)$  for simple model 3 with  $h^0 = 0.8$ . This shape is quite insensitive to changes in  $h^0$ .

The switching time  $t_1$  is found by solving the integral condition ensuring that the high-affinity clone has reached threshold. This condition is

$$t_1 = \frac{1}{k_0} (\log(\Lambda^2) - \log(2k_p h^0 k_0 I)). \quad (\text{A.22})$$

The integral  $I$  is

$$I = \int_0^\tau \frac{\mu(1-\mu) dt}{2k_p h^0 \mu(1-\mu) + z(k_0 - 2\mu k_p h^0)} \quad (\text{A.23})$$

where we use  $\mu$  to mean  $\mu(z(t_1+t))$ , etc. This integral involves no undetermined constants and does not depend on  $t_1$ , but only on  $\tau$ , which we know explicitly. So this integral can be performed numerically without difficulty.

For  $k_p$  and  $k_d$  fixed,  $I$  depends only on  $h^0$  and this only weakly. Figure A2 shows this in a plot of the "correction" term  $-\log(2k_p h^0 k_0 I)$  vs.  $h^0$ , when the rate constants are as given in the fully complicated model. Figure A3 plots the total time,  $T$ , and the burst duration,  $\tau$ , against  $h^0$ .  $T$  diverges as  $h^0$  approaches  $1/2$ , where  $k_0 = 0$ . These times may be interpreted as the interburst interval and the burst duration, respectively, showing qualitative agreement with the numerical analysis of the full model.

The mutation rate  $\mu(t)$  has a characteristic shape (Fig. A4) that changes very little as  $h^0$  varies, but approximately scales with  $\tau$ , rising from a value of 0 slightly convex downward, and reaching a value of  $1/2$  at the terminal time. This shape remains characteristic of the rising phase of the mutational bursts in the phasic solutions of the full model.