Large-scale cellular automata simulations of the immune system response

M. Bernaschi and S. Succi
Istituto Applicazioni Calcolo, CNR, Viale Policlinico 137, 00161, Roma, Italy

F. Castiglione
Center for Parallel Computing, University of Cologne, Weyertal 80, D-50931 Köln, Germany

(Received 30 April 1999)

The sequential nature of the process allowing the immune system to learn how to withstand pathogen agents is explored by means of large-scale computer simulation of the Celada-Seiden immunological automaton. In accord with our previous results, it is found that the learning process proceeds via a sequential cascade in affinity space.

PACS number(s): 87.16.Ac

The immune system (IS) is a complex organization made of cells and molecules whose collective dynamics exhibits emergent properties that can hardly be inferred from the underlying dynamics of its microconstituents. One major such property is the *learning* ability: how does the IS learn how to mount a specific response against invading entities (antigens, or Ag for short)? This capability is rooted in the specific functions of each of the microconstituents, but it is also true that the way the IS as a whole learns to withstand antigen attacks depends even more on the mutual interactions between these microconstituents. Once the importance of collective behavior is acknowledged, a relevant question becomes whether nonequilibrium statistical mechanics and the theory of (nonlinear) dynamical systems provide a convenient mathematical framework to characterize, at least semiquantitatively, the generic features of the immune system response [2,3].

In a recent paper, we presented preliminary efforts along this line using the Celada-Seiden immunological cellular automaton [1,4]. The Celada-Seiden [5] (CS) cellular automaton is based on a bit-string representation of the biological actors, namely B cells, T cells, antigen presenting cells (macrophages), antigens, antibodies. Details on the rules governing the microdynamics of these actors have been discussed elsewhere [4,5] and will not be repeated here. Within this bit-string model, each cell is characterized by a *bit-matching number m* denoting the number of matching bits with the bit string representing the antigen. Bits match when they are complementary $(0 \leftrightarrow 1)$.

High (low) affinity is therefore to be understood as high (low) values of the matching number m. With a string length l, the *repertoire* of the model is best organized into a hierarchical set of $l+1$ classes of cells characterized by the *matching number* $m=0,1,2 \ldots l$. The generic mth class contains $\binom{l}{m}=l!/(m!(l-m)!)$ elements; the sum over all possible classes involving a total repertoire $\Sigma_{m=0}^{l}\binom{l}{m}=2^l$ possible specificities. This defines the internal phase space of the automaton. The *active region* (defined later) has a structure markedly pyramidal: only one state with perfect match $m=l$, l states with $m=l-1$, and so on down the line.

The population density in this phase space is given by the *actual* occupation number $N_m(x,t)$ of cells in class m at site x at time t. Knowledge of the occupation numbers N_m at each space-time location yields a complete characterization of the dynamical system. The B and T cell binding affinity is expressed via an *''affinity potential''* V_m. The specific form of the affinity potential is not known in detail from biological data, but it is plausible to express it in the form of a sharply increasing function of m above a critical cutoff $m_c<l$, and zero below it. As a result, only a subset of the phase space, characterized by the condition $m_c \leqslant m \leqslant l$, is immunologically active. We shall call this the *active region* of immunological phase-space, whose size $l-m_c+1$ counts the populations competing for the antigen.

We reported [1] numerical evidence that the IS learning process proceeds through a cascade of higher and higher affinity populations (B cell) in which the low-affinity modes indirectly feed the higher-affinity ones in a sequential-like process dubbed *learning cascade*. A quantitative indicator of the aforementioned learning cascade was identified with the Kullback relative entropy, or information gain [6], defined as $G(\tau)=\Sigma_m f_m^2 \ln f_m^2/f_m^1$, where we refer to a transformation taking the system from initial state ''1'' at time $t=t_1$ to final state ''2'' at time $t_2=t_1+\tau$, and $f_m=N_m/\Sigma_{m'}N_{m'}$ is the probability density (normalized to 1) of class m. Ordinary Boltzmann entropy did not show any sign of monotonic behavior, which is not surprising since there is no reason to believe the underlying microdynamics of the CS automaton should obey a Boltzmann H principle.

In accord with the theory of clonal selection, the learning process is basically a shift of the occupation numbers towards the high-affinity region of the spectrum, a bias towards ''smart'' individuals. A prime indicator of this shift is an increase in time of the average matching $\mu(t)=\Sigma_{m,x} m f_m(x,t)$ where x runs over the spatial extension of the system. The gain can be interpreted as the information cost needed to bring the IS from low- to high-affinity states, and under certain assumptions on the shape of f_m it can be expressed as an analytic function of μ [1].

Our previous results pertained to relatively small repertoires, with $l=12$, consisting of $2^{12}=4096$ specificities. This is about four orders of magnitude lower than the *expressed* repertoire of the human immune system. The question, which makes the hard core of this paper, is whether the generic features observed in our previous work do survive once larger repertoires are considered. To this purpose, we have

TABLE I. Affinity potential V_m for convex and concave shapes.

m	Convex V_m	Concave V_m	Concave V_m/ Convex V_m
15	0.05000	0.05000	1.000
16	0.09103	0.50072	5.501
17	0.16572	0.74829	4.515
18	0.30171	0.88428	2.931
19	0.54928	0.95897	1.746
20	1.00000	1.00000	1.000

upgraded our computational tool to take full advantage of *parallel computing* capabilities [4]. Specifically, we have extended the size of the repertoire from 2^{12} up to 2^{20}, namely more than *two* orders of magnitude above our previous work, more than an order of magnitude beyond any previous study we are aware of, and, more importantly, *only an order of magnitude* below the expressed repertoire of the real IS. Being aware that a mere rise in the size of the repertoire does not necessarily imply a corresponding gain of immunological fidelity, we have also included hypermutation, namely, the mechanism by which clones that differentiate from their mother cells may show point mutations in their receptors [7,8].

In the model, hypermutation is represented by a given string s turning into a different string s' as a result of one (or more) bits changing state (zero to one, one to zero). The qualitative effect of hypermutation is to generate cells that would not appear otherwise in the system, thus giving the IS more freedom to explore its phase space. The chief question is whether such freedom is used to help affinity maturation, and, if so, to what extent. This question is genuinely dynamical in nature. On the one hand, affinity-degrading (high-to-low) mutations are more likely than affinity-enhancing (low-to-high) ones simply on account of the pyramidal structure of the active region of the phase space. On the other hand, since high-match cells are more effective in capturing the antigens, once generated they get a chance to reproduce faster than all other competing cells and possibly promote the affinity maturation. Whether such a chance does indeed materialize in actual practice is a nontrivial question that involves a genuinely dynamic nonequilibrium process. Computer simulation is well placed to provide a semiquantitative guidance in this complex territory.

We have performed a series of numerical simulations by varying the string length l, with and without mutation, and the shape of the affinity potential. The simulations are performed on a 16×16 grid, with the following parameters: average lifetime of B cells, $\tau_B=10$, initial population $B(0)=2184$, birth rate $\dot{B}\simeq0.07B(0)$. The Ag are injected at a rate of 1000 individuals per time step. Each time step corresponds to about 8 h in real time. Finally, we assume a time-independent single-bit mutation rate equal to $p=0.02$. We do not address the issue of *optimal* mutation schedule as in Ref. [9]. Our observations are based on a series of simulations with $l=20$, $m_c=15$ with and without mutation, and two typical shapes of the affinity: (a) Convex, (b) Concave.

The values of the affinity functions are reported in Table I.

Each simulation has been performed 40 times with different random seeds to double-check possible dramatic differences in the outcomes. Although no quantitative conclusions can be drawn, we can state that there are no indications of a strong sensitivity to the choice of the random numbers.

Our data show that affinity maturation does take place and in all cases *it proceeds through a sequential cascade from low- to high-affinity populations* (see Fig. 1).

Since no virgin B cells above m_c are allowed, the appearance of active cells in the course of time is necessarily due to hypermutation. The conclusion is that, albeit penalized in the average, hypermutation does have a dramatic effect in promoting affinity maturation. The intuitive picture is that, once a favorable mutation occurs by fluctuation, the higher reactivity of the high-affinity cells allows them to reproduce and survive for a long time.

In all cases, affinity maturation ramps up between $t=50$ and $t=100$ time units, that is, between two and four weeks, respectively, in real time. During this burst, the global affinity $V(t)=\Sigma_m N_m(t)V_m$ grows by almost two orders of magnitude, or more, depending on the shape of the affinity potential, as shown in Fig. 3. This burst is followed by a slower but steady growth associated with an increasing fraction of high-affinity cells.

Runs without hypermutation (see Fig. 2) also show a learning cascade, actually faster than with hypermutation.

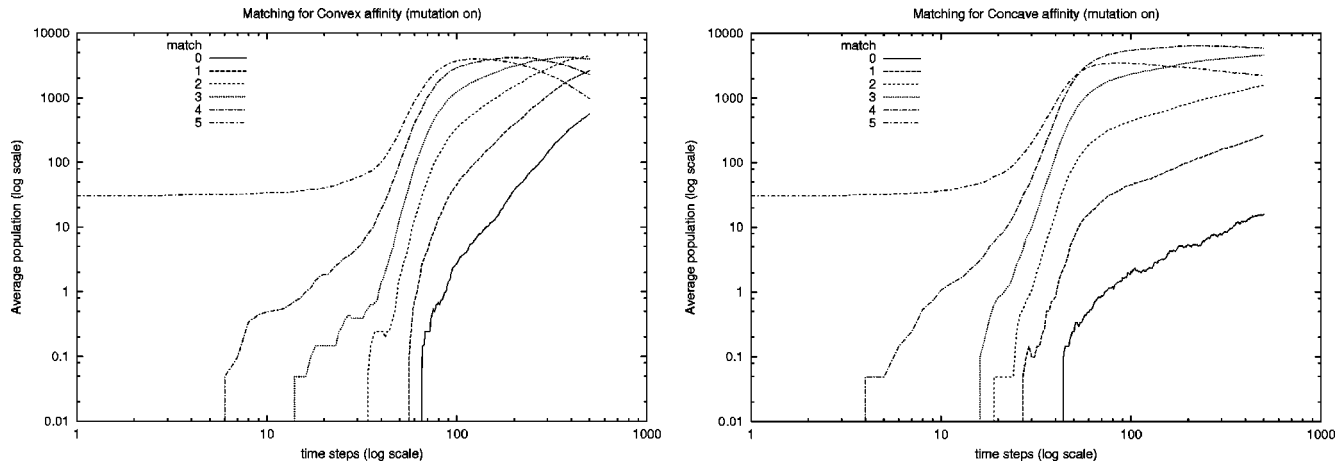

FIG. 1. Population growth for affinity classes of the active region for string length of 20 bits (details in the plots' titles).

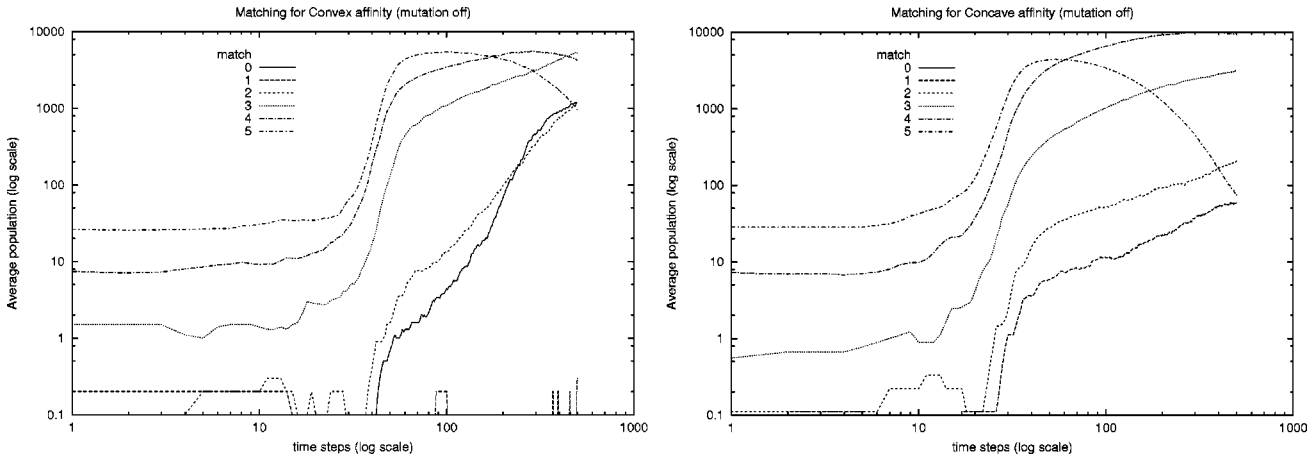

FIG. 2. Same data as in Fig. 1 for runs without mutation.

This means that the chance of generating active cells ($m \geqslant m_c$) by a mutation event is smaller than the corresponding chance of generating it out of the binomial distribution of virgin cells in the hypermutation-free scenario.

Given this intrinsically transient scenario, it is useful to develop a semiquantitative rationale for the role of the shape of the affinity potential.

To this purpose, let us consider all matching bits of a bitstring s (whose length is l) as part of a substring g (''good'' bits) and all other bits of s as part of a substring b (''bad'' bits). String g has length m ($m \geqslant m_c > l/2$). Obviously, b's length $\bar{m}$ is equal to $l-m$.

By definition, mutations on b enhance the affinity whereas mutations on g decrease it.

To compute the total probability of increasing the Hamming distance (i.e., the number of $0 \leftrightarrow 1$ matchings) of n units, we must take into account all possible combinations of mutations in the b and g strings such that $k-j$ is equal to n. Here, k and j are, respectively, the number of mutations in the b and g string and $j<k$ so that $n>0$. The expression of such total probability is

$$P_n^+ = \sum_{j,k:k=n+j} \binom{m}{j} p^j q^{m-j} \binom{\bar{m}}{k} p^k q^{\bar{m}-k}. \qquad (1)$$

In Eq. (1) the block $\binom{m}{j} p^j q^{m-j}$ gives the probability of j mutations in the g string, whereas the block $\binom{\bar{m}}{k} p^k q^{\bar{m}-k}$ gives the probability of k mutations in b. Upon using the relation $k=n+j$, we may recast Eq. (1) in terms of the index j, which runs between 0 and $\bar{m}-n$ (otherwise we would have a degrading mutation). Consequently, we can write

$$P_n^+ = \sum_{j=0}^{\bar{m}-n} \binom{m}{j} \binom{\bar{m}}{n+j} p^{n+2j} q^{l-(n+2j)}. \qquad (2)$$

The probability of affinity-degrading mutations, P_n^-, is obtained by considering $j>k$, swapping m with $\bar{m}$ in the above expression and letting the index run from 0 to $\bar{m}$.

One minute's thought reveals that, since $m_c > l/2$, under the standard condition $m > \bar{m}$, $P = P_n^+ - P_n^-$ is negative, reflecting the intuitive idea that on average, mutation works against high-affinity cells.

It is now instructive to observe that for low mutation rates $p \ll 1$, the summation in expression (2) can be replaced by its first order term. This leads to a very handy expression for the total one-bit affinity-improving mutation rate,

$$P_1^+ \sim \bar{m} p (1-p)^{l-1} = (l-m) p (1-p)^{l-1}, \qquad (3)$$

showing that, within this approximation, the one-bit improving mutation probability decays linearly with m.

Coming back to the interpretation of our results, the picture is as follows. Once an active cell materializes, the chance of capturing an antigen and rapidly duplicating it can be estimated as $A_m = 1 \times V_m$, that is, one cell times its microscopic affinity V_m. How many cells should materialize before the duplication process is actually triggered? The condition is of course $N_m V_m \geqslant 1$, which sets a natural threshold $1/V_m$ for the affinity maturation process to take off. This threshold would of course favor high m's, were it not for the strong penalty set by the mutation probability: a direct jump forward of n matching numbers scales roughly like p^n, which means that the next matching number above m_c is picked up by the condition $N_m V_m p^{m-m_c} > 1$, associated with a critical threshold $N_m^c \sim p^{m_c - m}/V_m$. It is easily seen that N_m^c is a sharply decreasing function of m, unless V_m would grow

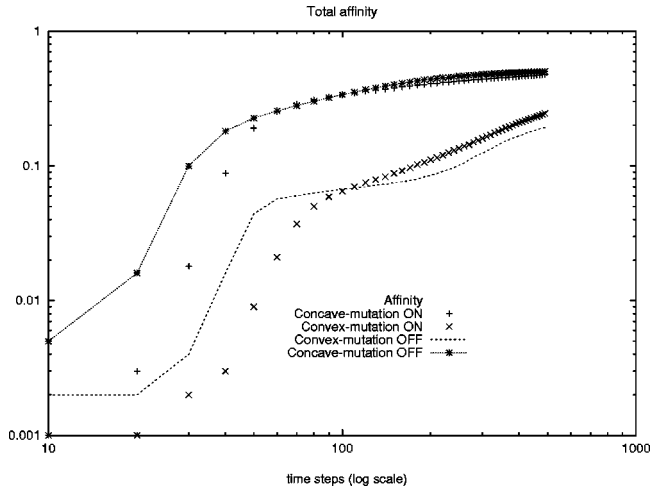

FIG. 3. Total affinity as a function of time (log-log).

faster than the decrease of p^{m-m_c}, not a plausible assumption given the small value of p. This explains the sequential nature of the learning cascade.

The next question relates to the mid-long term dynamics of the response. Here, two competing effects must be balanced.

The probability of cell clonation is proportional to V_m. A stimulated cell duplicates every step during four steps after stimulation, yielding 16 clones in four steps. This exponential growth is contrasted by a mean hypermutation loss proportional to P. There is no way for mutation to compete with such exponential growth in the short term.

It is plausible to assume that the long-term winner is selected by the condition of maximizing $N_m V_m / P$, which leads to a second threshold, $N_m^{c,2} \sim P/V_m$. Now, using the simplified expression P_1^+ given by Eq. (3), it is readily seen that highest-affinity modes are favored unless V_m grows more slowly than P_1^+, i.e., linearly, with m. This supports the intuitive idea that convex potentials favor the development of high-affinity populations as the asymptotic carriers of the IS response.

Of course, the notion of asymptotic, long-term carriers, although interesting from the point of view of statistical mechanics (final attractor of the system) is not necessarily the one most relevant to immunological purposes. To this end, one is probably more interested in the short- and mid-term (days-to-weeks) dynamics of the global affinity $V(t) = \Sigma_m N_m(t) V_m$.

This is shown in Fig. 3 for the two different choices of affinity potential. From these curves we see that, notwithstanding the significant statistical fluctuations in the initial phase, indeed the concave potential yields the quickest and most intense response, especially within the ramp-up period, up to $t \sim 100$. This is obviously due to the much higher affinity of the B cells (see fourth column in Table I). Asymptotically, however, our data suggest that the convex shape might be able to recover due to the emergence of perfect match cells with $m = l$ suffering less competition with other cells as compared to the case of a concave potential.

The above considerations, albeit still semiqualitative, provide a sound background for the interpretation of affinity maturation as a cascade process in affinity space. They also show that the learning cascade is quite robust vis-a-vis the shape of the affinity potential. This latter, however, plays a central role in the short- and mid-term dynamics of the IS response.

In closing, a few considerations on computational performance are in order. The parallel simulator takes about 10 ms/step per grid point, corresponding to about 10,000 s elapsed time for a 500-step-long (about 160 days) simulation on four processors of an UltraSparc Enterprise 4500. Memory requirements peak at about 2 GBytes during the burst of affinity maturation. These figures prove that the numerical investigation of the immune system response via the Celada-Seiden automaton requires substantial amounts of computational resources.

Fruitful discussions with Professors F. Celada and M. Nowak are kindly acknowledged. Special thanks to Dr. P. Seiden for his invaluable comments and suggestions. The Regional Computing Center of the University of Cologne (RRZK) is kindly acknowledged for providing the computational resources.

[1] S. Succi, M. Bernaschi, and F. Castiglione, Phys. Rev. Lett. **79**, 4493 (1997).
[2] A. Perelson and G. Weisbuch, Rev. Mod. Phys. **69**, 1219 (1997).
[3] S. Levin *et al.*, Science **275**, 334 (1997).
[4] F. Castiglione, M. Bernaschi, and S. Succi, Int. J. Mod. Phys. C **8**, 1 (1997); J. Fut. Gen. Comput. Sys. **15**, 333 (1999).
[5] F. Celada and P. Seiden, J. Theor. Biol. **158**, 329 (1992).
[6] R.S. Ellis, *Entropy, Large Deviations and Statistical Mechanics* (Springer Verlag, New York, 1995).
[7] M. Weigert, I. Cesari, S. Yonkovitch, and M. Cohn, Nature (London) **228**, 1045 (1970).
[8] F. Celada and P.E. Seiden, Eur. J. Immunol. **26**, 1350 (1996).
[9] T.B. Kepler and A.S. Perelson, J. Theor. Biol. **164**, 37 (1993).