

Clonal expansion without self-replicating entities

Danesh Tarapore^{1,2}, Anders Lyhne Christensen³, Pedro U. Lima¹, and Jorge Carneiro²

¹ Institute for Systems and Robotics (ISR), Instituto Superior Técnico (IST),
Lisbon, Portugal

² Instituto Gulbenkian de Ciência, Oeiras, Portugal

³ Instituto de Telecomunicações & Instituto Universitário de Lisboa (ISCTE-IUL),
Lisbon, Portugal

Abstract. The vertebrate immune system is a complex distributed system capable of learning to tolerate the organisms' tissues, to assimilate a diverse commensal microflora, and to mount specific responses to invading pathogens. These intricate functions are performed almost flawlessly by a self-organised collective of cells. The robust mechanisms of distributed control in the immune system could potentially be deployed to design multiagent systems. However, the essence of the immune system is clonal expansion by cell proliferation, which is difficult to envisage in most artificial multiagent systems. In this paper, we investigate under which conditions proliferation can be approximated by recruitment in fixed-sized agent populations. Our study is the first step towards bringing many of the desirable properties of the adaptive immune system to systems made of agents which are incapable of self-replication. We adopt the crossregulation model of the adaptive immune system. We develop ordinary differential equation models of proliferation-based and recruitment-based systems, and we compare the predictions of these analytical models with results obtained by a stochastic simulation. Our results define the operational parameter regime wherein growth by recruitment retains all the properties a cell proliferation model. We conclude that rich immunological behaviour can be fully recapitulated in sufficiently large multiagent systems based on growth by recruitment.

Keywords: Multiagent system, clonal expansion, agent recruitment, crossregulation model

1 Introduction

Multiagent systems (MAS) comprise a large number of research domains, ranging from software agents to multirobot systems, and play an important role in several applications, such as supply chain management, transportation logistics and network routing. The coordination of agents in a MAS is a major challenge because agent behaviour depends not only on interactions with their immediate environment but also on the behaviour of other agents. A centralised control approach may not always be feasible due to computational and communication

constraints on agents (e.g., [1]). Distributed control, on the other hand, is often complicated to realise because the behavioural rules for the individual units cannot be easily derived from a desired macroscopic behaviour (e.g., [2]). In the design of large scale distributed systems, several researchers have therefore taken inspiration from nature e.g., aggregation of amoeba into slime mold [3], quorum sensing and communication in bacteria [4], division of labour in social insects such as ants [5].

The cell collective that constitutes the adaptive immune system has been extremely successful during the course of evolution as evidenced by its presence in all jawed vertebrate species [6]. Central to this success are the T helper cells that orchestrate the system. These cells are capable of dynamically regulating and differentiating themselves into different sub-lineages (e.g., Th1, Th2, Th17, and regulatory T-cells) to initiate appropriate immune responses (e.g., [7–9]). Importantly, the differentiation and clonal expansion of these cells into different functionally distinct sub-types is decentralised and based solely on the state of the individual cell and local information available to it. Furthermore, cell proliferation and death are fundamental processes underlying clonal expansion, which is the essence of the immune response.

The decentralised and adaptive nature of the immune system, together with the relative simplicity of an individual cell, makes it an appealing model for designers of large scale MAS. Examples of studies that take inspiration from the immune system include distributed intrusion detection systems [10], fault tolerance systems [11], and behaviour arbitration mechanisms in robotics [12, 13]. In these models, a many-to-one analogy between cells and agents is considered, with the behaviour of the agent dictated by the number and type of cells it possesses. While this approach does facilitate some interesting applications, it requires individual agents to have a high degree of complexity and may consequently not be feasible in certain applications, e.g., involving energy constraints on agents. In this study, we explore an alternative one-to-one mapping between a cell and an agent. An obvious obstacle to this approach is the absence of artificial self-replicating agents to implement clonal expansion by cell proliferation. However, as we demonstrate here, this obstacle can be circumvented by a mechanism of “recruitment” that mimics cell proliferation and cell death within a fixed number of agents.

In our investigation, we use an immune system model of *crossregulation* to compare proliferation and recruitment. The crossregulation model (CRM) [14, 15] postulates a dynamics of interactions between T helper cells, that allows the system to discriminate between antigens based solely on their density and persistence in the environment. The system is able to tolerate body antigens (i.e “self”) that are characteristically persistent and abundant, and to mount an immune response to foreign pathogens, that appear as bursts. The model has been used successfully in spam detection (e.g., [16]) and document classification (e.g., [17]) scenarios, making it a good candidate for MAS in other environment classification tasks. In order to determine if and when recruitment can be used to emulate proliferation with reasonably accuracy, we compare two system: one

that relies on proliferation and one that relies on recruitment, respectively. We model both systems analytically and evaluate them in a stochastic simulation environment.

2 The Crossregulation Model

Two principles of multicellular organisation are the foundation of the crossregulation model (CRM). Firstly, the persistence of any cell lineage requires that its cells recurrently interact with other cells in the organism. Cells that fail to interact with other cells eventually die. Secondly, the growth of a cell population involves density-dependent feedback mechanisms controlling individual cell proliferation. These feedback mechanisms involve (i) indirect interactions among cells (such as a competition for limited growth factors), and (ii) direct interactions, such as contact inhibition. Below, we outline the model and highlight its interesting properties that are later replicated with a multiagent system that uses clonal expansion based on recruitment instead of cell proliferation.

The CRM describes the population dynamics of immune cells, based on three mutually interacting cell types: (i) antigen presenting cells (APCs) that display the antigen on their surface; (ii) effector cells T_E that can potentially mount an immune response which, depending on receptor specificity, can be directed against foreign pathogens or self-antigens; and (iii) regulatory cells T_R that suppress the proliferation of T_E cells. Furthermore, individual APCs have a fixed number of binding sites (s) on which T_E and T_R cells can form conjugates.

Dynamics of the T-cell population is regulated by the following density-dependent feedback mechanisms: (i) effector and regulatory cells that are unable to interact with APCs are slowly lost by cell death; (ii) the proliferation of effector and regulatory cells requires conjugation with APCs and direct interactions between T_E and T_R cells co-localised at the APC. The proliferation of the T_E cell population is promoted by the absence of regulatory cells on the APC. In contrast, T_R can only proliferate following co-localisation with effector cells on the same APC. Additionally, T_E and T_R cells interact indirectly by competition for access to conjugation sites on APCs.

2.1 Mathematical formulation of the CRM

The individual T-cells are in one of the following three states; *free* (i.e., not conjugated with an APC), *conjugated* with an APC, and *activated* (about to proliferate). In our mathematical model, effector cells in the free, conjugated and activated states are denoted by Ef , Ec and Ea , respectively. Similarly, regulatory cells in the free, conjugated and activated states are Rf , Rc and Ra , respectively. Consequently, the total number of T_E cells $E = Ef + Ec + Ea$, and the total number of T_R cells $R = Rf + Rc + Ra$.

T-cell–APC conjugation: The dynamics of effector and regulatory T-cells conjugated with APCs is described by the following system of equations. For Ec and Rc we have, respectively:

$$\frac{dEc}{dt} = \gamma_c A E f \left(1 - \left(\frac{Ec + Rc}{As} \right)^s \right) - \gamma_d Ec - \delta Ec \quad (1)$$

$$\frac{dRc}{dt} = \gamma_c A R f \left(1 - \left(\frac{Ec + Rc}{As} \right)^s \right) - \gamma_d Rc - \delta Rc \quad (2)$$

where A is the density of APCs, γ_c and γ_d denote the conjugation and dissociation rate constants between APCs and T-cells, respectively, and δ is the death rate of T-cells (parameters in Table 1).

The equations for Ec (eq 1) and Rc (eq 2) have three terms. The first term represents the conjugation of free T-cells with APCs which at least one unoccupied site. The second term accounts for the dissociation of existing conjugates. Finally, the third term represents the process of cell death, assumed to be a simple exponential decay of T-cells.

T-cell population dynamics with cell *proliferation*: According to the CRM, conjugated effector cells are activated and able to proliferate if and only if there are no regulatory cells conjugated to the same APC. If we let P_e denote the probability that a conjugated effector cell has no neighbouring regulatory cell, the equations describing the dynamics of free Ef and activated Ea effector cells are:

$$\frac{dEf}{dt} = -\gamma_c A E f \left(1 - \left(\frac{Ec + Rc}{As} \right)^s \right) + \gamma_d (1 - P_e) Ec + 2\pi_E Ea - \delta Ef \quad (3)$$

$$\frac{dEa}{dt} = \gamma_d P_e Ec - \pi_E Ea - \delta Ea \quad (4)$$

where π_E is the proliferation rate of activated effector cells (parameters in Table 1).

In the equations for Ef (eq 3) and Ea (eq 4), the conjugated effector cells with no neighbouring regulatory cell on the same APC are selected for activation, and consequently proliferate. By contrast, conjugated effector cells co-localised with one or more neighbouring regulatory cells dissociate without proliferating.

The conjugated regulatory cells can only be activated if at least one effector cell is simultaneously conjugated to the same APC. If we let P_r denote the probability that a conjugated regulatory cell has at least one neighbouring effector cell, the equations describing the dynamics of free Rf and activated Ra regulatory cells are:

$$\frac{dRf}{dt} = -\gamma_c A R f \left(1 - \left(\frac{Ec + Rc}{As} \right)^s \right) + \gamma_d (1 - P_r) Rc + 2\pi_R Ra - \delta Rf \quad (5)$$

$$\frac{dRa}{dt} = \gamma_d P_r Rc - \pi_R Ra - \delta Ra \quad (6)$$

where π_R is the proliferation rate of activated regulatory cells (parameters in Table 1).

Table 1: Parameters for the CRM.

Parameters	Description	Value (a.u.)
A	Density of APCs	—
s	Number of binding sites on an APC	3
E	Total density of effector cells	$Ef + Ec + Ea$
R	Total density of regulatory cells	$Rf + Rc + Ra$
γ_c	Conjugation rate constant of T-cells to APCs	1
γ_d	Dissociation rate constant of T-cells from APCs	10^{-3}
π_E	Proliferation rate of effector cells	10^{-4}
π_R	Proliferation rate of regulatory cells	0.5×10^{-4}
δ	Death rate of effector and regulatory cells	10^{-5}
N	Density of cells, for the recruitment model	30×10^{-3}
γ_r	Recruitment rate constant of cells	10

In the equations for Rf (eq 5) and Ra (eq 6), conjugated regulatory cells with one or more neighbouring effector cells, are selected for activation and consequently proliferate at rate π_R . By contrast, conjugated regulatory cells with no neighbouring effector cell, dissociate without proliferating.

For our system of ODEs, the probability functions P_e and P_r can be expressed with a multinomial approximation [18]. This approximation is reasonable given that the total number of binding sites (summed over all the APCs) is much larger than the number of sites per APC. For 3 binding sites ($s = 3$) on each APC, we have:

$$P_e = \frac{(Rc - 3A)^2}{9A^2} \quad \text{and} \quad P_r = \frac{(6A - Ec)Ec}{9A^2} \quad (7)$$

Behaviour of proliferating T-cell population: The dynamics of the CRM is governed by two key composite parameters representing the effective growth rates of E and R cell populations [14]. These two growth rates are directly proportional to the basic parameters controlling population growth i.e., conjugation and dissociation constants (γ_c and γ_d), proliferation rates of these two types of T-cells (π_E and π_R), and the density of APCs (A). The effective growth rates of the T-cells are also inversely proportional to the death rate (δ) of the corresponding population. The composite E and R growth parameters define four parameter regimes according to the resulting cell population behaviour. Three parameter regimes result in a single stable state that may correspond to either: (i) extinction of all T-cells ($E = 0$, $R = 0$), (ii) immune state ($E > R$), or (iii) tolerant state ($E < R$). The fourth parameter regime corresponds to a bistable system where both immune and tolerant states are stable. A detailed analysis of these parameter regimes is provided in [14].

In the present study, the parameter values have been set so that at low APC densities ($a_E < A < a_R$, Fig. 1a), the system evolves towards a single state composed only of effector cells (immune state). By contrast, at relatively high




Fig. 1: Equilibrium densities of E and R cell populations as a function of APC densities A with proliferation-dependent growth. (a) Bifurcation diagram of the CRM, representing all possible equilibrium states. The lines indicate the total density of T-cells (sum of the variables, $E + R$) as a function of the APC density (parameter A). (b) Equilibria that are actually reached by solving the system with with a fixed seed T-cell population (solid line: $E = 10^{-3}, R = 10^{-3}$; dashed line: $E = 3/4 \times 10^{-3}, R = 1/4 \times 10^{-3}$). Remaining parameter values as in Table 1.

density of APCs ($A > a_R$, Fig. 1a), the system exhibits bistability and can evolve either into the immune or the tolerant equilibrium state. In the bistable regime, the system develops into the regulatory cell dominated state, provided that the seeding population has sufficient T_R cells (Fig. 1b). By contrast, if T_R cells are initially underrepresented than the T_E cells competitively exclude the former from the system (Fig. 1b).

T-cell population dynamics with cell *recruitment*: The T-cell proliferation and death are fundamental processes used by the adaptive immune system to orchestrate an appropriate immune response under varying environmental conditions. Furthermore, these processes are crucial to the density-dependent feedback mechanisms of the CRM. However, in some multiagent systems, new agents cannot be created and existing agents cannot be removed from the system, making it difficult to translate adaptive immune system inspired algorithms to these scenarios. In this section, we address this problem and describe a model of crossregulation in populations with a fixed number of cells that use *recruitment* instead of proliferation. We also characterise the properties of this modified CRM, and highlight the conditions under which it behaves similar to a proliferation model.

In our recruitment model, the activated T-cells can not “proliferate” without the presence of an idle cell for recruitment. Consequently, in addition to T-cells in free, conjugated, and activated states, we define an intermediate state of T-cells (Ei and Ri , for T_E and T_R , respectively), that are scouting for idle cells.

The total number of T_E cells $E = Ef + Ec + Ea + Ei$, and the total number of T_R cells $R = Rf + Rc + Ra + Ri$. We further define a prefixed total density of cells N . Cells that are neither effectors nor regulator are considered as idle I . The density of idle cells I is defined by the conservation equation, $I = N - E - R$.

The dynamics of the interactions between effector and regulator T-cells with APCs is unchanged when recruitment is used instead of proliferation (eqs 1 and eqs 2 for Ec and Rc , respectively). However, the actual process of cell proliferation is now replaced by the recruitment of an idle cell into the proliferating cells' type (effector or regulator), at recruitment rate constant γ_r . In addition, the process of cell death is simulated by the transition of the effector or regulator T-cell into an idle state. Consequently, the equations of the dynamics of T_E cells are now:

$$\frac{dEf}{dt} = -\gamma_c A Ef \left(1 - \left(\frac{Ec + Rc}{As} \right)^s \right) + \gamma_d (1 - P_e) Ec + 2\gamma_r EiI - \delta Ef \quad (8)$$

$$\frac{dEa}{dt} = \gamma_d P_e Ec - \pi_E Ea - \delta Ea \quad (9)$$

$$\frac{dEi}{dt} = \pi_E Ea - \gamma_r EiI - \delta Ei \quad (10)$$

In the equations for recruited Ef (eq 8) and Ea (eq 9), the conjugated effector cells with no neighbouring regulatory cell on the same APC are selected for activation. However, the actual growth of these cells is dependent not only on π_E as with the proliferation model but also the recruitment rate constant γ_r applied to Ei (eq 10), and the density of idle cells I .

Similarly, the equations describing the dynamics of T_R cells are:

$$\frac{dRf}{dt} = -\gamma_c ARf \left(1 - \left(\frac{Ec + Rc}{As} \right)^s \right) + \gamma_d (1 - P_r) Rc + 2\gamma_r RiI - \delta Rf \quad (11)$$

$$\frac{dRa}{dt} = \gamma_d P_r Rc - \pi_R Ra - \delta Ra \quad (12)$$

$$\frac{dRi}{dt} = \pi_R Ra - \gamma_r RiI - \delta Ri \quad (13)$$

where the conjugated regulatory cells with one or more neighbouring effector cells are selected for activation and the growth of these cells is dependent on π_R , γ_r applied to Ri , and I .

Behaviour of *recruiting* T-cell population: In the CRM with recruitment, as with proliferation, the effective growth rates of E and R cell populations is directly proportional to: (i) conjugation and dissociation constants (γ_c and γ_d), (ii) ‘‘proliferation’’ rates of the two types of T-cells (π_E and π_R), and (iii) the density of APCs (A). The effective growth rates of the recruited T-cells are also inversely proportional to the ‘‘death’’ rate (δ). For comparison, these parameters are set to the same values as in the proliferation model (Table 1). In addition, with cell recruitment, the effective growth rates of the T-cells is also




Fig. 2: Equilibrium proportion of recruited effector cells (E/N) as a function of APC densities A for low ($N = 15 \times 10^{-3}$), medium ($N = 30 \times 10^{-3}$) and high ($N = 45 \times 10^{-3}$) total density of T-cells, and recruitment rate $\gamma_r = 10$. Remaining parameter values as in Table 1.

directly proportional to the recruitment rate constant (γ_r), and the total density of cells (N).

In the CRM, the dynamics of the T-cell population is governed by interactions between T_E , T_R cells and APCs. The consequent density dependent feedback mechanisms involved would not operate if the T-cell population can grow no more, i.e. all the T-cells have been recruited. The maximum density of recruited T-cells is investigated by solving the steady state density of recruited effector cells (E) with no inhibitory regulatory cells ($Rf = 0, Rc = 0, Ra = 0$), for different A and N (Fig. 2). In order to operate under conditions wherein density of recruited cells is not saturated, we select parameter value $N = 30 \times 10^{-3}$, $\gamma_r = 10$, and A below 10^{-4} (Fig. 2).

In the chosen parameter regime, at low APC densities ($a_E < A < a_R$, Fig. 3a) the stable node composed only of effector cells (immune state) is globally stable, as was the cases with CRM with proliferation. At relatively high densities of APCs ($a_R < A < 10^{-4}$, Fig. 3a), the system displays bistability and can evolve into either the immune or the tolerant equilibria, consequent to the proportion of effector and regulatory cells in the seeding population (Fig. 3b).

A Jacobian analysis at the biologically relevant equilibrium points (T-cells densities real and positive and $E+R < N$) revealed real and negative eigenvalues for the immune state (stable node). By contrast, at the tolerant state eigenvalues were complex with negative real components, indicating dampened oscillation in the population dynamics. Furthermore, at the unstable equilibrium a mixture of positive and negative eigenvalues were present.

At APC densities above our operational parameter regime ($A > 10^{-4}$), the immune state continues to be stable, with the unstable state and the tolerant state approaching it. At very high APC densities ($A > 7.5 \times 10^{-4}$) the complex eigenvalues of the tolerant state have a mixture of positive and negative real

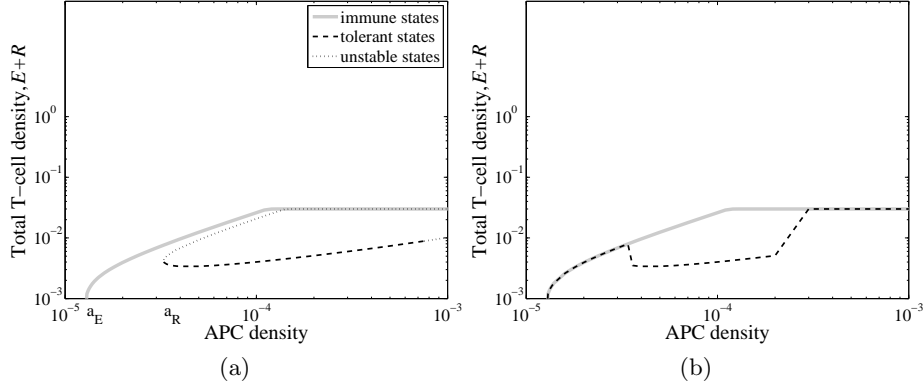


Fig. 3: Equilibrium densities of E and R cell populations as a function of APC densities A , with cell *recruitment*. (a) Bifurcation diagram of the CRM, representing all possible equilibrium states. The lines indicate the total density of recruited T-cells (sum of the variables, $E + R$) as a function of the APC density (parameter A). (b) The lines indicate the equilibria that are actually reached by solving the system with a fixed seed T-cell population (solid line: $E = 10^{-3}, R = 10^{-3}, N = 30 \times 10^{-3}$; dashed line: $E = 3/4 \times 10^{-3}, R = 1/4 \times 10^{-3}, N = 30 \times 10^{-3}$). Remaining parameter values as in Table 1.

components (Fig. 3a). At APC densities above 10^{-3} , the immune state becomes globally stable again (not shown in figure).

3 Stochastic simulation of the model

A stochastic, spatial, discrete-time simulation is used to compare the CRM using proliferation and recruitment, for different number of APCs. The simulated environment is toroidal and has a size of $S = 1000 \times 1000$ units. The MAS is composed of static point-sized entities (APCs) and mobile agents (T-cells), wherein each agent represents an individual cell. Each agent moves at a speed of $v = 0.5$ units/time-step when unattached to the APC, and has probability 0.01 of changing to a new random direction each simulation step. Agents can perceive APCs up to a range $w_c = 1$ unit, while they can perceive other agents up to a range $w_r = 10$ units.

In our experiments, each mobile agent can be an effector (T_E), a regulator (T_R) or idle. Furthermore, effector and regulator agents transition between a *free* mode, *attached* mode and an *activated* mode as follows: (i) At the start of the experiment, the agent is free and explores the environment. When an agent detects an APC with a free binding site, it attaches itself to the APC; (ii) Upon detachment from the APC (probability γ_d at each step), effector agents with no regulator agents on the same APC are activated. Similarly, regulator agents with at least one neighbouring effector agent are also activated. The rest of the

agents return to the free mode and resume random movement; (iii) Activated agents undergo proliferation with probability π_E and π_R , for effectors and regulators, respectively. In proliferation experiments, the parent and daughter agents return to the free mode. By contrast, in recruitment experiments, the effector or regulator agent, scouts for an idle agent, and upon successful recruitment, the recruiter and recruitee agents return to the free mode; and (iv) The effectors and regulators in all modes are subject to death (probability δ).

The process of proliferation is simulated, either by the introduction of a new agent at the proliferating agent’s position (proliferation model), or by the recruitment of an idle agent within range (recruitment model). In both cases, the daughter agent adopts the type of the proliferating/recruiting agent, that is, the daughter immediately transitions either to the effector type or to the regulator type depending on its parent’s type. In addition, death is simulated by the removal of the agent (proliferation model), or by switching the agent to the idle state (recruitment model).

3.1 Calculating agent parameters

In our stochastic simulations, we set the parameters to connect our experiment results to the ODE model. The simulation probabilities of dissociation, proliferation, and death in each iteration step (τ) are set from the corresponding rates of the ODE model (Table 1) as $\gamma_d\tau$, $\pi_E\tau$, $\pi_r\tau$, and $\delta\tau$, respectively. Furthermore, the agent’s speed (v) and its detection range for APCs (w_c) and other agents (w_r), and the environment size S , are set as follows. Let us consider the rates μ_c and μ_r , at which an agent encounters an APC and another agent, respectively. These reaction rates can be computed from the corresponding reaction rate constants γ_c and γ_r of the ODE model ([19]), as $\mu_c = \gamma_c/S$ and $\mu_r = \gamma_r/S$.

As an agent moves through the arena, it sweeps out an area during time interval τ , and will detect an APC or another agent that fall in that area. The detection area can be expressed as is $2w_c v\tau$, for APCs and $2w_r v\tau$ for other agents. Considering that the APCs and agents are uniformly distributed in the arena, the detection rates $\mu_c = 2w_c v\tau/S$ and $\mu_r = 2w_r v\tau/S$.

3.2 Experiments

In experiments with proliferation, the MAS consisted of 1000 T_E agents and 1000 T_R agents. In separate experiments, the number of APCs was set at 20, 40, 60, 80 and 100. The MAS was simulated for 10^7 time-steps (with a constant iteration step $\tau = 1$) at which point the number of effectors and regulators had converged in all experiment conditions. In the case of the recruitment model, an identical set of experiments was conducted, with the total number of cells agents fixed at 30000 ($I = 28000$).

For both proliferation and recruitment models, since the population dynamics appeared similar at high densities of APCs, cell population dynamics has been shown for the extreme cases of 20 and 100 APCs. In these conditions, we compare the density of T_E (E) and T_R (R) cells from the numerical solution of the model, with the mean density of effector and regulator agents across 40 independent replicates.

The results of the stochastic simulation indicate the population of agents reaches one of two stable states depending on the density of APCs (Fig. 4): (i) an immune state, characterised by the presence of only effectors, and (ii) a tolerant state, consisting of a dominant population of regulators and few effectors. In both proliferation and recruitment growth simulations, the system converged to the immune state in the presence of few APCs (20 in Fig. 4a and b). An increase in the number of APCs in the environment (100 cells in Fig. 4c and d) resulted in convergence to the tolerant state. In addition, the numerical solution of our ODE system appears qualitatively similar to the results from the stochastic simulations, for both proliferation and recruitment models. The quantitative discrepancy between the results of numerical solution of the ODEs and of the simulations reflects the asymptotic nature of the former and the discrete finite nature of the latter [20].

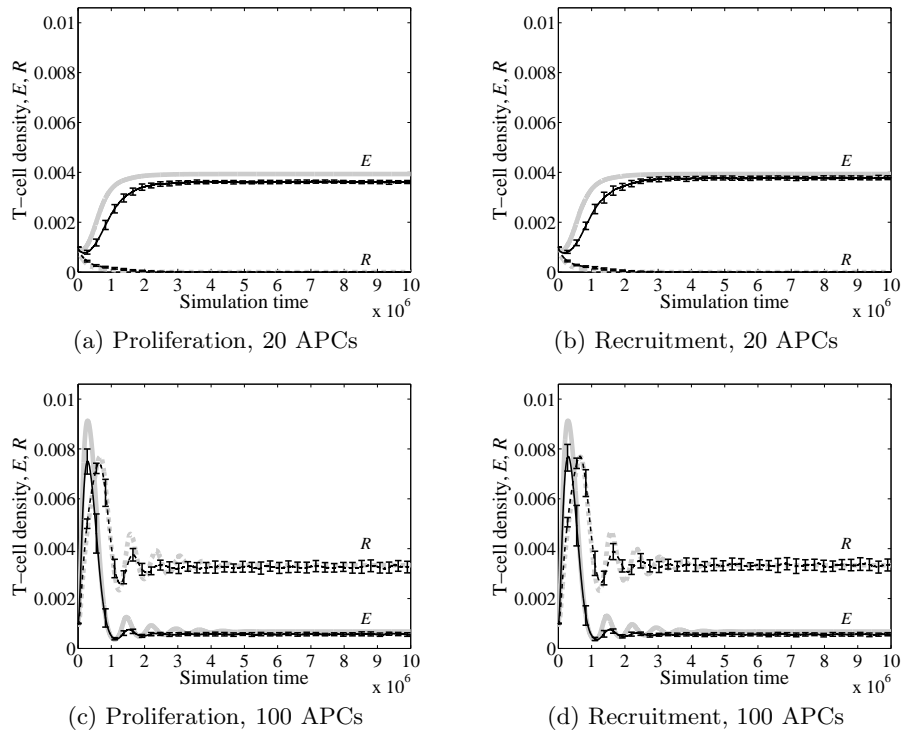


Fig. 4: Density of effector (E) and regulator (R) T-cells from the numerical solutions (light lines) of the ODE system, and the mean density of effector and regulator agents across 40 stochastic simulation replicates (dark lines)

In stochastic simulations operating in the bistable parameter regime ($A > a_R$), the system developed into the tolerant state (e.g., $A = 10^{-4}$, Fig. 4c and




Fig. 5: Trajectory of T-cell population with $A = 10^{-4}$ using proliferation (solid line) and recruitment (dashed line), upon perturbation of tolerant equilibrium state by 50% (seed populations S_1 , cross marker), 75% (seed populations S_2 , square marker) and 95% (seed populations S_3 , star marker).

d). In order to evaluate the stability of this tolerant state, we restarted the proliferation and recruitment growth simulations for the extreme case of (100 APCs) and at time-step 10^7 , with a perturbed density of effector and regulator agents. The new seed populations were computed by incrementing effectors, and decrementing regulators by 50%, 75% and 95%. The perturbed populations were simulated for 4×10^6 time-steps at which point the number of agents had converged in all six (3 perturbed seed populations each for proliferation and recruitment) experiment conditions. In Fig. 5, we have plotted the density of effector and regulator agents of a single replicate, for 4×10^6 simulation steps, and with population perturbations at 50%, 75% and 95%. For seed populations perturbed up to 75%, the system returned to the tolerant state, irrespective of the mechanism of population growth (proliferation and recruitment) (Fig. 5). At relatively high perturbations (e.g., 95%), the system transitioned to an immune state, with the population dominated by effectors, with no regulator agents present (Fig. 5).

4 Discussion and conclusions

In this study, we highlighted some of the challenges in deploying immune system inspired algorithms for MAS. We proposed an alternative to the mechanism of clonal expansion, namely recruitment, that may be used in fixed-sized agent populations. We identified and characterised an operational parameter regime, wherein recruitment can be used while retaining the properties of the model. In this operational regime, the analytical models and stochastic simulations of growth by recruitment were shown to recapitulate the same properties of CRM, where growth is dependent on cell division.

In the natural immune system, the total number of T-cells is maintained fairly constant in healthy individuals by homeostatic mechanisms [21], irrespective of the T-cell specificity or differentiation class (e.g. Th1, Th2, and so on). Antigen-driven clonal expansion of a subpopulation implies erosion of other subpopulations. The net result of clonal expansion in the presence of strong homeostasis is akin to the recruitment-based growth dynamics we proposed here.

An important issue to note when utilising the recruitment model is the large number of agents needed. In our simulations, 30000 agents had to be present in the population. Additionally, an increase in APCs would demand an even higher number of agents. Another potential approach wherein the total number of required agents may be reduced, would be the recruitment of other cell types, besides idle cells, e.g., in the CRM, an activated regulator cell may be able to recruit an effector cell. However, the outcome of this scenario needs to be explored further, particularly in stochastic simulations with fewer agents, where the discrepancies with the mean field model predictions can be not only quantitative (as in the case of the parameter regime studied here) but even qualitative [20]. Also worth exploration is the impact of random perturbations to the number of agents representing the T-cells or the entities representing APCs, as stochastic perturbations could drive anomalous behaviours akin to relapsing autoimmunity [22].

The CRM, for which our recruitment mechanism has been assessed, describes a dynamics of interactions between T-cells of the immune system, that allows the system to discriminate between “self” and “nonself” antigens based solely on their density and persistence in the environment, and largely independent of their intrinsic characteristics. This property of the model opens up very interesting applications involving decentralised feature classification, wherein our recruitment model may be utilised, circumventing the need for self-replicating units.

In summary, we have shown that our approach to applying immune system inspired algorithms for MAS is promising, particularly for applications involving severe constraints on agents or agents that cannot self-replicate, such as speckled computing or collective robotic systems.

Acknowledgment: This study was supported by the FCT grant PTDC/EEACRO/104658/2008.

References

1. Crespi, V., Galstyan, A., Lerman, K.: Top-down vs bottom-up methodologies in multi-agent system design. *Autonomous Robots* **24**(3) (2008) 303–313
2. Hamann, H.: *Space-Time Continuous Models of Swarm Robotic Systems: Supporting Global-to-Local Programming*. Cognitive Systems Monographs. Springer-Verlag, Berlin, Germany (2010)
3. Payton, D., Estkowski, R., Howard, M.: Compound behaviors in pheromone robotics. *Robotics and Autonomous Systems* **44**(34) (2003) 229–240
4. Sahin, E.: *Swarm Robotics: From Sources of Inspiration to Domains of Application*. In: *Swarm Robotics*. Volume 3342. Springer, Berlin, Germany (2005) 10–20
5. O’Grady, R., Groß, R., Christensen, A.L., Dorigo, M.: Self-assembly strategies in a group of autonomous mobile robots. *Autonomous Robots* **28** (2010) 439–455

6. Janeway, C., Travers, P., Walport, M., Shlomchik, M.: *Immunobiology: The Immune System in Health and Disease*. New York: Garland Science (1997)
7. Park, H., Li, Z., Yang, X., Chang, S., Nurieva, R., Wang, Y., Wang, Y., Hood, L., Zhu, Y., Tian, Q., Dong, C.: A distinct lineage of CD4 T cells regulates tissue inflammation by producing interleukin 17. *Nature Immunology* **6** (2005) 1133–1141
8. Mosmann, T., Coffman, R.: Th1 and Th2 cells: different patterns of lymphokine secretion lead to different functional properties. *Annual Review of Immunology* **7** (1989) 145–173
9. Sakaguchi, S.: Naturally arising CD4⁺ regulatory T cells for immunologic self-tolerance and negative control of immune responses. *Annual Review of Immunology* **22** (2004) 531–562
10. Nino, F., Beltran, O.: A change detection software agent based on immune mixed selection. In: *Proceedings of the 2002 Congress on Evolutionary Computation (CEC)*, IEEE Computer Society, Washington, DC (2002) 693–698
11. Bradley, D.W., Tyrrell, A.M.: Immunotronics: Hardware fault tolerance inspired by the immune system. In: *Proceedings of the 3rd International Conference on Evolvable Systems: From Biology to Hardware (ICES)*, Springer-Verlag, Berlin, Germany (2001) 11–20
12. Watanabe, Y., Ishiguro, A., Shirai, Y., Uchikawa, Y.: Emergent construction of behavior arbitration mechanism based on the immune system. In: *Proceedings of the IEEE World Congress on Computational Intelligence*, IEEE Computer Society, Washington, DC (1998) 481–486
13. Singh, S., Thayer, S.: Kilorobot search and rescue using an immunologically inspired approach. In: *Proceedings of the 6th International Symposium on Distributed Autonomous Robotic Systems*, Springer-Verlag, Berlin, Germany (2002) 300–305
14. Leon, K., Perez, P., Lage, A., Farob, J., Carneiro, J.: Modelling t-cell-mediated suppression dependent on interactions in multicellular conjugates. *Journal of Theoretical Biology* **207**(2) (2000) 231 – 254
15. Carneiro, J., Leon, K., Caramalho, I., Van Den Dool, C., Gardner, R., Oliveira, V., Bergman, M., Sepúlveda, N., Paixão, T., Faro, J., Demengeot, J.: When three is not a crowd: a crossregulation model of the dynamics and repertoire selection of regulatory CD4⁺ T cells. *Immunological Reviews* **216**(1) (2007) 48–68
16. Abi-Haidar, A., Rocha, L.: Adaptive spam detection inspired by the immune system. In: *Proceedings of the 11th International Conference on the Simulation and Synthesis of Living Systems, Artificial Life XI*, MIT Press, Cambridge (2008) 1–8
17. Abi-Haidar, A., Rocha, L.: Collective classification of textual documents by guided self-organization in T-cell cross-regulation dynamics. *Evolutionary Intelligence* **4**(2) (2011) 69–80
18. Evans, M., Hastings, N., Peacock, B.: 27. In: *Statistical Distributions*. Third edn. John Wiley & Sons (2000) 134–136
19. Gillespie, D.: A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics* **22** (1976) 403–434
20. Figueroa-Morales, N., Leon, K., Mulet, R.: Stochastic approximation to the t cell mediated specific response of the immune system. *Behavioral Ecology and Sociobiology* **21**(295) (2012) 37–46
21. Sprent, J., Surh, C.: Normal T cell homeostasis: the conversion of naive cells into memory-phenotype cells. *Nature Immunology* **12**(6) (2011) 478–484
22. Velez de Mendizabal, N., Carneiro, J., Sole, R., Goni, J., Bragard, J., Martinez-Forero, I., Martinez-Pasamar, S., Sepulcre, J., Torrealdea, J., Bagnato, F., Garcia-Ojalvo, J., Villoslada, P.: Modeling the effector - regulatory T cell cross-regulation

reveals the intrinsic character of relapses in multiple sclerosis. BMC Systems Biology **5**(1) (2011) 114