

An Engineering-Informed Modelling Approach to AIS

Emma Hart and Despina Davoudani

Edinburgh Napier University, Edinburgh, Scotland, UK
{e.hart, d.davoudani}@napier.ac.uk

Abstract. A recent shift in thinking in Artificial Immune Systems (AIS) advocates developing a greater understanding of the underlying biological systems that serve as inspiration for engineering such systems by developing abstract computational models of the immune system in order to better understand the natural biology. We propose a refinement to existing frameworks which requires development of such models to be driven by the engineering problem being considered; the constraints of the engineered system must inform not only the model development, but also its validation. Using a case-study, we present a methodology which enables an abstract model of dendritic-cell trafficking to be developed with the purpose of building a self-organising wireless sensor network for temperature monitoring and maintenance. The methodology enables the development of a model which is consistent with the application constraints from the outset and can be validated in terms of the functional requirements of the application. Although the result models are not likely to be biologically faithful, they enable the engineer to better exploit the underlying metaphor, ultimately leading to reduced development time of the engineered system.

1 Introduction

The field of Artificial Immune Systems (AIS) has derived inspiration from many different elements of the natural immune system in order to develop engineered systems that operate in environments with constraints similar to those faced by the immune system [1]. Encouraged perhaps by views such as that of Timmis *et al* [2] that suggest rapid development in AIS is being hindered by both the lack of theoretical foundation and the lack of challenging application areas, a shift in focus has been apparent in the literature over recent years. On the one hand, we see the use of immune-inspired algorithms arising in increasingly more complex application areas, e.g swarm robotics [3] and wireless sensor networks [4]. Such applications embody the principles laid down in [5] which attempted to define the characteristics of applications which could properly benefit from the use of immune-inspiration. Simultaneously, significantly more research in AIS is under-pinned by the use of modelling techniques (both computational and mathematical) which provide theoretical foundations for algorithms, and proper immunological grounding. We propose that this shift towards theoretical

development is necessary to drive innovations and development in future complex applications. In this paper, we examine the role that modelling can play in informing the development of a complex engineered system through the use of a case-study. The chosen study examines the creation of an immune-inspired prototype for controlling a *specknet* [6] — a particular type of ad-hoc wireless network which enables programmable computational networks to be constructed from thousands of miniature devices which sense, process and store information.

We consider the use of one particular aspect of immunology in association with specknets, that of dendritic-cell trafficking, and use an agent-based modelling approach to bridge the gap between the two complex systems. Specifically, we draw on a study which concerns mapping the data-collection and monitoring functionality of dendritic cells into an application which monitors and controls the environment inside a building. Although modelling as a tool has been used extensively by immunologists in efforts to either explain experimental data or to suggest novel lines of experimentation, our approach differs in that we specifically seek to understand how immune-inspired algorithms might be used within a complex engineered environment through the model. We deliberately make no attempt to produce an immunologically meaningful model that could be expected to provide insights to biologists (an approach also advocated by [7]). Instead, we propose that the model fulfills the following roles:

- it captures the high-level functionality of the biological system *and* the constraints of the engineered system
- it enables exploration of a range of state-of-the-art mechanisms from the biological literature, which may often be conflicting
- it enables validation of the functionality of the model according to metrics or measures appropriate to the engineered system

Through the use of a case-study, we show how a model can be developed with the above in mind. The proposed methodology is generically applicable to other domains propose a methodology and successfully bridges the gap between complex domains.; the results suggest that modelling is a useful exercise in tackling complex engineering problems, even if it leads to models which are biologically unfaithful.

2 On immunology, engineering & modelling

Cohen [8] notes that there are three categories of people pursuing research which relates computing and immunology:

1. those of the *literal* school that build systems in silico to try and do what the actual immune system does (e.g. build computer security systems that discriminate between self and non-self),
2. those of the *metaphorical* school that look for inspiration from the immune system and build computational systems with the immune system in mind, and

3. those who aim to understand immunity through the development of computer and mathematical models.

This work is rooted in the second category, in that our goal is to build engineered systems. In line with this, we take a brief look at the literature which advocates the use of modelling to facilitate achieving this goal.

One of the seminal papers responsible for advocating a swing towards a theoretical approach to AIS was published by Stepney *et al* in [9], in which a Conceptual Framework for AIS development was presented. The framework suggests a three-phase process which begins with biology; in Phase 1, the underlying biology system which is being used for inspiration is probed, utilising biological observations and experiments to provide a partial view of the biological system. Phase 2 concerns the development of two types of model: firstly, abstract models of the biology are built from which analytical computational frameworks developed which can be analysed and validated and provide a set of *design principles*. In Phase 3, the principles are used to construct bio-inspired algorithms applicable to non-biological domains.

While following such a process undoubtedly leads to a greater understanding of the immunology on the part of the computer scientist, to date, there remain few algorithms in the literature which credit their derivation to following this framework. A few notable exceptions exist. Owens *et al* [10, 11] refer to biological literature in Phase 1 of the framework, then show how models of T-cell receptor signalling, developed and validated in Phase 2, can be used to derive a Kernel Density Estimation method for anomaly detection. Andrews [12] evaluates the framework itself in his PhD, first examining immunological literature for inspiration, then building models of a receptor degeneracy which inspire the creation of a pattern classification algorithm. [13] begin with Phase 2 and develop a model of self-assertion in the immune system and analyse the model properties in depth. They suggest that the model can be transferred to data-mining applications, although they fall short of actually deriving an algorithm, i.e. Phase 3. Work by [14] is also worth mentioning — they present a lineage of work on the Dendritic Cell Algorithm (DCA) which resulted from very close interactions with lab-based immunologists. In fact, early versions of the algorithm even derived values for numeric parameters from results observed in the laboratory.¹ In this work, a modelling phase did not occur; the jump was made directly from Phase 1 to algorithm development in Phase 3.

However, it remains clear that there are some issue with the framework as originally proposed. Firstly, little guidance is given as to how to usefully follow the framework, as noted by [12] who proposes that *an approach would be to bias the route through the Conceptual Framework by having an application in mind*. Secondly, it seems clear that the computational and abstracted models which are formulated in Phase 2 may be inconsistent with engineering requirements of an application. This point is taken up by Timmis *et al* in an approach dubbed

¹ Note, however, that in subsequent refinements of the algorithm much of the immunological detail has been removed.

immuno-engineering [15], which proposes a mechanism for the abstraction, adaptation and application of immune principles to engineering, taking into account the differences between artificial systems and biological systems. This could for example account for the different numbers, kinds, and rates of signals that need to be monitored and processed; the different kinds of decisions that need to be made; the different effectors available to support and implement those decisions; and the different constraints of embodiment, either physically or virtually engineered. We are unaware however of any cited examples of this approach being followed.

We therefore propose a modification to the Conceptual Framework in which the engineering requirements of the application in question are explicitly identified before Phase 1 and Phase 2 are undertaken. With engineering functionality and constraints clearly identified, the immunology literature or immunology expert can be consulted; an abstract model of the immunology can be developed at an appropriate level; this can be then translated into computational model which is consistent with the engineering constraints. The resulting translation and validation process may result in models which necessitate the introduction of mechanisms inconsistent with the immunological system. Further, the modelling process may further illuminate aspects of the desired functionality that are not appropriate to derived from biological metaphors.

An obvious question arises of how such a model can be validated given that it is no longer appropriate to calibrate it against experimental immunological data. We propose that the validation process shifts to validating *functionality* rather than either mechanistic detail or the ability to replicate experimental results. We propose instead a shift in focus to achieving similar functionality in the two systems. In this respect, modelling can be viewed as explorative, rather than restrictive, with no requirement for models to be biologically plausible.

3 The Application Domain: SpeckNets

This section provides a brief overview of specks in order to illustrate the complexity of the engineering environment under consideration. Details of the engineering aspects of specks, including technical specifications, are discussed in [6]. A detailed motivation of the use of a dendritic cell (DC) analogy is given in [4]. For the purpose of the paper, it is sufficient to focus on the constraints imposed by this complex engineering system, rather than focus on the implementation details.

Specknets are *large scale* — they can consist of thousands of nodes. As they are designed to be sprayed into an environment, they necessarily have an *ad-hoc topology*. Minimal battery power means communication is *limited in both range and in terms of power* as sending and receiving messages is extremely energy draining. The use of on-board energy saving mechanisms means that the nodes *cannot be guaranteed to be functioning* at any moment time. Finally, lack of a central control system forces the network to operate in a *decentralised and asynchronous fashion*.

4 Scoping the model by consideration of engineering

Consider an application in which specks sprayed onto the walls of a large building monitor temperature in their local area. Individual radiators can be remotely switched on and off to control temperature in a local area and the goal is to maintain a constant ambient temperature in the building. From an engineering perspective, we require a system which is capable of *monitoring* the environment and *effecting decisions*. Control must be *de-centralised* and there is a desire to *minimise communication overhead* in order to save power. Turning to immunology for inspiration, the role played by dendritic-cell trafficking mechanisms appears a plausible metaphor for meeting these requirements: these cells serve as *sentinels* of the immune system, acting as scouts which both monitor an environment and prime the adaptive immune system by presenting it with the information regarding how and where to mount a response.

A possible mapping between the two systems is to enable radio signals to function as dendritic cells; signals propagating around the network can sample data from specks and return collected information to specks designated as 'lymph nodes'. The lymph specks evaluate returned data and modulate the temperature as required. At this level of abstraction, the mapping is straightforward. However, when one looks in detail at the immunological literature describing mechanisms which enable trafficking of dendritic cells it becomes clear that the processes are both complex and incompletely understood. In particular, the number of actors is immense, and furthermore, the literature is ambiguous regarding the conditions under which cells are triggered to migrate to lymph nodes (the reader is referred to [16] for an overview).

Given the engineering constraints of the application in terms of minimising power and overhead, it is clear that a model which utilises all the actors present in the biological system will fail. Therefore, in line with the views of Forrest [1], it is necessary to abstract the immune functionality and capture generic properties, rather than intricate detail. In order to do this, it is necessary to clearly articulate the high-level *functionality* of the immune system which map to the engineered requirements:

- Dendritic cells as *scouts*: mobile cells circulate through an environment sampling signals from it.
- Dendritic cells as *reporters*: DCs return to lymph nodes to present a snapshot of the information collected in the tissue on maturation.
- *Routing* via chemical gradients: chemokine gradients are used to route DCs back to lymph node and to direct T cells back to infected sites.
- System *stabilisation*: T cells are released from lymph-nodes if determining a response is required, and are able to participate in reactions which ultimately eliminate infection at local sites.

Note that this is in contrast to the majority of work inspired by DCs in AIS which focuses on the differentiation capabilities of these cells in response to relative amounts of predefined types of signal collected and therefore their

function as *classifiers*. Instead, our emphasis is on the ability of DCs to function as *scouts* and *reporters*. The model is not concerned with the details of how the decision to either tolerate or react to the information returned by DCs is arrived at; this is a property of the adaptive immune system, and is the subject of ongoing research and a further model. It is, however, useful in this model to investigate a crucial property of the response, that it can be directed back to an infected site (stabilisation), assuming a decision to respond had been made as this functionality also relates to DC trafficking.

5 An Engineering Constrained Immunological Model

A high-level description of an *immunological* model of dendritic-cell trafficking which can be analysed and validated in order to determine whether it achieves the required *functionality* can be summarised as follows: a constant size population of immature DCs circulates at random through the tissue. Spontaneous infections randomly arise in the tissue which release danger signals. Naïve DCs differentiate into mature DCs if the amount of danger signal collected from the tissue exceeds a threshold. Immature DCs die naturally on reaching the end of a predefined lifespan. Lymph nodes are found at random locations within the tissue, each of which is surrounded by a cytokine gradient which attracts mature DCs. Mature DCs reaching a lymph node cause the lymph node to release T cells; these circulate at random until they are attracted by a cytokine gradient which surrounds an infection. T cells passing through an infected site reduce the amount of infection at that site until the infection is cleared.

5.1 Validation of the Model

Validation of the model must confirm that the high-level properties defined in section 4, emerge from the model. This contrasts to the majority of work in biological modelling where validation involves comparing the model to experimentally observed results to determine its correctness. According to the engineering application, the validation may be quantitative or qualitative, depending on whether appropriate metrics can be defined.

In this case, the properties listed in section 4 can be measured by the metrics defined in table 1 : these quantities can be extracted from the simulation. Note in some cases the measured quantity is a proxy from another measure. For example, we monitor the *reporting* ability of DCs by directly measuring the number of T cells produced in the system, as T cells are only produced in response to returning DCs.

5.2 Model Implementation

We have chosen to model DC trafficking using an agent-based model (ABM). Bonabeau [17] discusses in depth the advantages of agent-based models which can be summarised as: (i) *they capture emergent phenomena*; (ii) *they provide*

<i>Property</i>	<i>Description</i>	<i>Measure</i>
<i>Scouting</i>	Tissue is adequately sampled	Tissue coverage
<i>Reporting</i>	Mature DCs return to lymph nodes	(1) size of T-cell population (2) no. DCs reaching lymph
<i>Routing</i>	DCs directed back to lymph T cells directed to infected sites	(1) no. DCs at lymph (2) infection level
<i>Stabilisation</i>	t-cells effect a reaction	infection level

Table 1. The table shows quantitative measures which can be used to give an indication of whether the system is exhibiting the desired characteristics.

a natural description of a system; and (iii) they are flexible. More specifically, Forrest and Beauchemin [1] have discussed in depth the effectiveness of applying Agent-Based Modelling techniques (ABM) to immunology. The particular environment adopted is NetLogo [18], due to both its popularity in the community and ease of use.

5.3 Assumptions and Simplifications from an Engineering Perspective

Evaluation of the biological literature to determine what to model must be guided by engineering rather immunological considerations. Further assumptions and simplifications will be necessary in the model implementation — it is vital that the implementation must respect engineering constraints from the outset. In this case, we consider the following aspects:

Production of dendritic cells: Dendritic cells originate in the bone marrow from where they are circulated through the tissue. This is simplified in the immunological model by generating DCs at random locations throughout the tissue, corresponding to generation of radio signals by individual specks in the engineering system.

Trafficking of immature DCs: Mobilisation of dendritic cells into lymphatic vessels is generally assumed to require maturation (either semi or full). However, evidence that immature DCs also gather in the lymph nodes is slowly accumulating, suggesting that, for example, immature Langerhans cells also traffic to lymph nodes [19]. According to Janeway [20], tissue DCs reaching the end of their life-span in body tissues without having been activated by infection can also travel via the lymphatics to local lymphoid tissue where they induce tolerance. The former approach corresponds to returning “pre-classified” information to the lymph, the latter to returning all collected information. Both mechanisms are plausible from the engineering perspective, we implement both trafficking scenarios in the model.

Simplification of cytokine network: Rather than modelling the large numbers of cytokines and chemokines which are known to orchestrate dendritic cell trafficking, we use a single class of cytokine agents as a proxy for those cytokines expressed by mature DCs which attract them back to lymph nodes, and another

single class of cytokine agent to represent cytokines which direct T cells to infected sites within the body. Reducing the numbers, and by implication, volume, of cells trafficking around the system will clearly be of utmost importance in the engineered system.

Placement of lymph nodes: In the body, the secondary lymphoid organs are strategically placed near major organs and important sites. As specific placement of individual specks is unlikely to be feasible in distributed specknets which may consist of thousands of specks, we model a random distribution of lymph nodes in the environment.

Lymph nodes: A lymph node consists of an outermost cortex and inner medulla, the cortex being divided mainly into lymphoid follicles, paracortical (T-cell) areas and germinal centres. We simplify this to a single compartment within which all interactions between cells occur. Again, this simplification is driven by engineering requirements; lymph nodes are represented by individual specks in the network and will undertake processing of all received messages.

Lack of lymphatic vessels: Cells are physically directed to lymph nodes via a system of lymphatic vessels in the natural immune system. Such a physical system cannot be created in a specknet in which messages travel by radio broadcast and cannot be constrained to certain pathways. Therefore, we replace physical lymphatic vessels in the model with cytokine gradients emanating from lymph nodes which serve to direct cells back to the lymph.

Antigens are not explicitly modelled: The focus of the model is to understand how dendritic cell trafficking can be used to collect information from a system, and return that information to a location where collaborative decision making takes place based on all returned information. Therefore, we are not explicitly concerned with modelling responses to specific antigen initiated by T cells finding their cognate antigen in the lymph.² Furthermore, in specknet applications, antigen collected will be specific to the particular application. Therefore, we do not model collection of antigen explicitly.

Lack of specificity of T cells: Activated T cells exhibit particular cytokine profiles depending on the information they received in the lymph node from mature DCs. As the inclusion of T cells in the model only serves to illustrate that a response can be directed back to a specific site, one class of T-cell agent is modelled which is assumed to invoke the correct response at any site.³

A snapshot of the model running is shown in figure 1.

6 Analysis of the Model from an Engineering Perspective

According to the validation criteria listed above, the model was run using various scenarios and appropriate measures gathered. Due to space limitations, we do not present results from the model here. Instead, we present an analysis of how the processes implemented in the model contribute to the functional behaviour

² The learning process taking place in the lymph should be the focus of a separate model.

³ The implications of this are discussed later in the paper.

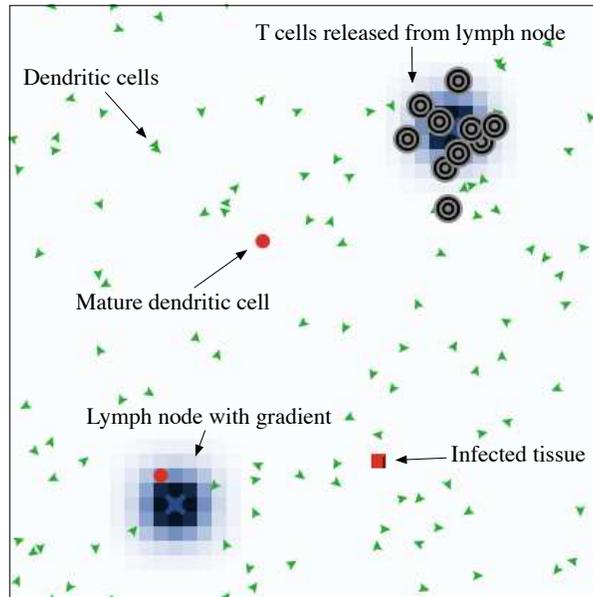


Fig. 1. A snapshot of the model running showing lymph nodes, infections and agents (dendritic cells, mature dendritic cells and T cells).

of the system, while considering the likelihood of transferring these behaviours to the engineered system. The analysis requires the model to be iterated — this must be undertaken in a manner driven by the engineering constraints of the system, rather than a drive to replicate biology, bearing in mind the functional requirements.

6.1 Dendritic Cell Generation

Immunological theory is clear that DCs are generated from bone marrow — this is replicated in the model and results show that it achieves the desired functionality of providing scouts. Note however that this is localised control mechanism; each tissue node releases DCs independently of each of the other nodes. In the context of specknets, this potentially has unwanted consequences — flooding of the network leading to loss of messages can occur through each node behaving independently of others. This could be ameliorated by incorporating a global mechanism to control the release of DCs across the entire system, however, this is antithetical to the vision of specknets which are self-organising and autonomous. Therefore, a balance should be found between local and global control: we propose to delegate control of DC generation to the lymph nodes in the engineered model. Although implausible from a biological perspective, this facilitates a method of controlling the number of DCs (radio messages) by allow-

ing each lymph node to be responsible for when messages are sent out to scout the network.⁴ The balance between local and global control can be altered by changing the number of lymph nodes. Although this has obvious implications in the expected coverage of the environment - areas close to lymph nodes are likely to be more frequently visited by DCs, the metrics can be used to quantify the effect on functionality. The model is therefore updated accordingly results are shown in figure 2. Although functionality is preserved in terms of reporting ability and stabilisation, the time to clear increases when DCs are generated from the lymph nodes.

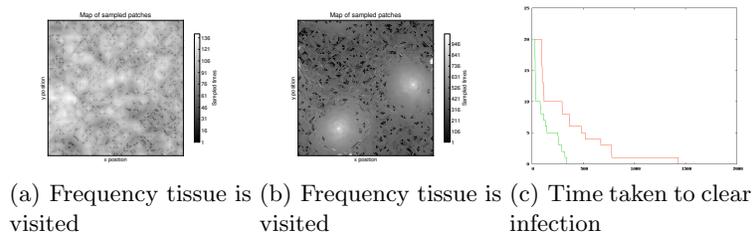


Fig. 2. Figure (a) and (b) contrast the frequency with which tissue patches are visited under tissue generation of DCs (a) and lymph-node DC generation regimes (b) (lighter = more frequent). Figure (c) contrasts the rate of clearance of infection under each scheme (left-hand line is tissue-generation method).

Generating DCs from lymph nodes limits the initial exploratory behaviour of immature DCs to areas surrounding the lymph nodes, as clearly shown in figure 2. As well as affecting clearance time, the resulting pattern of communication can potentially be problematic in specknets: nodes visited more frequently have increased overhead in terms of message sending which can deplete energy resources to the extreme in which nodes die. This potentially leave the lymph nodes isolated from the rest of the network. On the other hand, the pattern illustrated in figure 2(a) reveals a more even distribution of the communication load across the network. The suitability of each method will be ultimately be dictated by the requirements of the application; exploring the impact of approaches both within and beyond the bounds of biological plausibility using the model is however of value.

6.2 Trafficking of Immature Dendritic Cells

Although the immunological literature is ambivalent regarding the trafficking of *immature* DCs back to the lymph nodes, it is useful to examine the options

⁴ Note that this has the potential side-effect in the worst case scenario of introducing *multiple local SPOF (Single Point of Failure)* into the system due to failed lymph specks.

from an engineering perspective. In an engineering context, collection of data, i.e. routing messages around the network is costly in terms of energy. Discarding data by mimicking cell death ‘because it occurs in the immune system’ is therefore particularly wasteful. Routing ‘healthy’ data back to the lymph nodes clearly has an associated cost, but this is likely to be outweighed by the increased volume of data the classification algorithm implemented at the lymph node receives. Arrival of messages denoting healthy network states enables the learning algorithm to build a more accurate picture of the network.

Validation of this hypothesis is beyond the scope of this paper as it involves the *learning and classification* part of the algorithm derived from adaptive immunology, and not programmed as part of this model. However, this process provides an explicit link between the innate and adaptive parts of the system, therefore requires consideration during the innate modelling phase. The engineering constraints are used to evaluate mechanisms from the immunological literature for potential inclusion in the model.

6.3 Routing via Gradients

Dendritic cell trafficking in the natural immune systems occurs as a result of physical and chemical means: a physical network of conduits transport immune cells to lymph nodes through the tissue; cytokines and other chemicals recruit immune cells to the site of infection and promote healing of any damaged tissue following the removal of pathogens. As the former mechanism relies on a physical routing system which cannot be implemented in a WSN, any immune-inspired routing in the speck system must be gained from the chemical gradient based method.

The WSN literature already contains many examples of gradient-based routing techniques which have been optimised in term of minimising energy usage and thereby maximising lifespan of a WSN (e.g. [21, 22]). It is therefore clear from the outset that these can be used directly within the application without further recourse to immunology. While these techniques differ in exact implementation detail, they essentially have common ingredients: a cost associated with each node; a mechanism for setting up the costs (gradient); an optional mechanism which periodically refreshes costs. Although the model is unlikely to offer any further improvements on existing routing mechanisms, it does however offer an opportunity to investigate potential behaviours which may emerge when using gradient based routing.

Routing to Lymph Specknets typically have a static topology, with certain nodes within the network designated as lymph nodes. All lymph nodes perform identical processing. While it is preferable for a scouting message to return to the nearest lymph, it is acceptable for a message to be received by any lymph node. These requirements map exactly to typical situations in any WSN network, and therefore, the model adds no further insight; a routing mechanism can be selected directly from engineering literature.

Routing of T-cells to infected sites Infections (synonymous with an *incident* in a network) occur at random locations. This requires a new gradient to be set up by returning scouting messages each time an infection is encountered. Depending on the application two situations may arise: if all “infections” are identical, any lymph node can send out t-cells to counter the infection. If different types of infections occur however, appropriate t-cells must be routed to a specific infected site, rather than any infected site.

In initial simulations, all infected sites diffuse the *same* chemical to attract t-cells from a lymph node, in recognition of the complexity and energy cost associated with maintaining multiple gradient types within a specknet. Simulations showed that although t-cells are routed to infected sites, they can arrive at *any* infected site as the gradient is non-specific. The problem is exacerbated the further an infection is from the lymph nodes (recall the lymphs are randomly placed). *Stabilisation* of the network will therefore depend whether the application incurs homogeneous or heterogeneous infections. In the latter case, stabilisation would only occur if multiple gradients were maintained. Given the additional overhead this would require, it appears prudent to look to engineering rather than immunology at this point to solve a potential problem. The engineer is free to increase the number of lymph nodes — this simple modification would increase the probability of a message from a lymph node arriving at the closest infected site due to closer proximity.

An alternative approach would be to depart once again from the immunology. The functionality of lymph nodes could be modified, allowing them to take on a dual role as centres where information can be accumulated and classified, but also as *effectors* which can take actions based on that information. In this case, the need for implementing messages which function as T-cells is completely removed and the problem eliminated.

The model has therefore highlighted some possible consequences of using gradient based routing mechanisms in terms of the required functionality of the system, illuminating aspects of the model in which immunology can provide the required inspiration, and those in which engineering techniques are more appropriate.

7 Conclusion

In line with current thinking in the domain of AIS, we have developed a model of a complex immunological mechanism in order to derive a prototype for ultimately monitoring the state of a complex engineered environment, a specknet. From the outset, the model was developed with engineering constraints in mind. Immunological mechanisms were selected for inclusion in the model according to engineering constraints, and the model was validated in terms of functionality, rather than against immunological data. Although this process is likely to yield a biologically unfaithful model, it is of considerable help in the development of the engineered systems. In particular:

- It provides a mechanism for defining the correct level of abstraction of the biological system such that high-level properties of interest are conserved using minimal complexity.
- It is able to illuminate aspects of the model which cannot be transferred to an engineered system for practical reasons.
- It provides an environment in which modifications can be made to the model in light of engineering constraints, and tested to ensure that the model retains similar emergent properties.
- It facilitates the study of the computational aspects of the engineered model in an environment free from complex engineering constraints.

The necessity to produce models which represent computational abstractions of the immune system was identified by Stepney *et al* in the design of the Conceptual Framework. It was further elaborated by Timmis *et al* when proposing Immuno-Engineering, which takes into account differences between engineered and biological system. This paper however makes explicit the requirement to consider engineering constraints from the outset when following a modelling process if a successful result is to be achieved: we propose that the conceptual framework be modified to reflect this. Rather than simply biasing a route through the framework with *an application in mind* as suggested in [12], we show through the case-study that the framework can be of great benefit when the engineering requirements have been explicitly defined before the process begins. Bonabeau [23] recognised *the freedom of the engineer to adapt models in a manner limited only by one's imagination and the available technology*. As advances in technology lead to ever more complex engineered systems, an informed approach to modelling will enable the AIS engineer to develop truly novel algorithms, pushing the boundaries of the current state-of-the-art.

References

1. Forrest, S., Beauchemin, C.: Computer Immunology. *Immunol. Rev.* **216**(1) (2007) 176–197
2. Timmis, J.: Artificial immune systems: Today and tomorrow. *Natural Computing* **6**(1) (Feb. 2007) 1–18
3. Lau, H., Bate, I., Timmis, J.: An immuno-engineering approach for anomaly detection in swarm robotics. In Andrews, P.S., Timmis, J., Owens, N.D.L., Aickelin, U., Hart, E., Hone, A., Tyrrell, A.M., eds.: ICARIS. Volume 5666 of Lecture Notes in Computer Science., Springer (2009) 136–150
4. Davoudani, D., Hart, E., Paechter, B.: Computing the state of specknets: Further analysis of an innate immune-inspired model. In: ICARIS '08: Proceedings of the 7th international conference on Artificial Immune Systems, Berlin, Heidelberg, Springer-Verlag (2008) 95–106
5. Hart, E., Timmis, J.: Application areas of AIS: The past, the present and the future. *Applied Soft Computing* **8**(1) (2008) 191–201
6. Arvind, D., Wong, K.: Speckled computing: Disruptive technology for networked information appliances. In: in Proceedings of the IEEE International Symposium on Consumer Electronics (ISCE'04). (2004) 219–223

7. Ismail, A., Timmis, J.: Aggregation of swarms for fault tolerance in swarm robotics. In: UK Workshop on Computational Intelligence. (2009)
8. Cohen, I.: Real and artificial immune systems: computing the state of the body. *Nature Reviews Immunology* **07** (2007) 569–574
9. Stepney, S., Smith, R., Timmis, J., Tyrrell, A., Neal, M., Hone, A.: Conceptual frameworks for artificial immune systems. *Int. J. Unconventional Computing* **1**(3) (2006) 315–338
10. Owens, N.D., Timmis, J., Greensted, A., Tyrrell, A.: Elucidation of t cell signalling models. *Journal of Theoretical Biology* **262**(3) (2010) 452 – 470
11. Owens, N., Greensted, A., Timmis, J., Tyrrell, A.: T cell receptor signalling inspired kernel density estimation and anomaly detection. In: 8th International Conference on Artificial Immune Systems. Volume 5666 of LNCS., Springer (2009) 122–155
12. Andrews, P.: An Investigation of a Methodology for the Development of Artificial Immune System: A Case Study in Immune Receptor Degeneracy. PhD thesis, University of York (2008)
13. Dilger, W., Strangfeld, S.: Properties of the Bersini experiment on self-assertion. In Cattolico, M., ed.: GECCO, ACM (2006) 95–102
14. Greensmith, J., Aickelin, U., Tedesco, G.: Information fusion for anomaly detection with the dendritic cell algorithm. *Information Fusion* **11**(1) (2010) 21–34
15. Timmis, J., Hart, E., Hone, A., Neal, M., Robins, A., Stepney, S., Tyrrell, A.: Immuno-engineering. In: Biologically-Inspired Collaborative Computing. Volume 268. Springer Boston (2008) 3–17
16. Randolph, G.: Dendritic-cell trafficking to lymph nodes through lymphatic vessels. *Nature Reviews Immunology* **5**(8) (2005) 617–28
17. Bonabeau, E.: Agent-based modeling: Methods and techniques for simulating human systems. *PNAS: Proceedings of the National Academy of Sciences of the United States of America* **99**(suppl. 3) (2002) 7280–7287
18. Willensky, U.: Netlogo. <http://ccl.northwestern.edu/netlogo> (1999) Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston, IL.
19. Randolph, G.: Is maturation required for langerhans cell migration ? *J Exp Med.* **196**(4) (2002) 413–416
20. Janeway., C.A., Paul, T.: Immunobiology: The Immune System in Health and Disease. Third edn. Garland Publishing, New York (1997)
21. Ye, N.F., Chen, F.Y.A.: A scalable solution to minimum cost forwarding in large sensor. In: Computer Communications and Networks. (2001) 304–309
22. Faruque, J., Psounis, K., Helmy, A.: Analysis of gradient-based routing protocols in sensor networks. In: Distributed Computing in Sensor Systems: First IEEE International Conference, Springer (2005) 258–275
23. Bonabeau, E., Dorigo, M., Theraulaz, G.: Swarm Intelligence: From Natural to Artificial Systems. Oxford University Press (1999)