

Immunological Inspiration for Building a New Generation of Autonomic Systems

Emma Hart
Centre for Emergent
Computing
Napier University
Edinburgh, Scotland
e.hart@napier.ac.uk

Despina Davoudani
Centre for Emergent
Computing
Napier University
Edinburgh, Scotland
d.davoudani@napier.ac.uk

Chris McEwan
Centre for Emergent
Computing
Napier University
Edinburgh, Scotland
c.mcewan@napier.ac.uk

ABSTRACT

Autonomic computing systems of the future will be required to exhibit a number of properties which cannot be engineered using current technologies and algorithms. The most direct inspiration for building such systems is nature, where for example the central nervous system and the immune system function in an autonomic manner. In this paper we show how mechanisms inspired by recent advances in the field of immunology may offer exactly the inspiration required for engineering this new generation of computational systems which are robust, secure, self-organise and self-heal in manner currently unachievable with established software engineering techniques. Immune-inspired mechanisms are often synonymous with providing security in computing applications — however we intend to show that a wider examination of the immune literature offers far greater potential for exploitation of immune-mechanisms and paradigms than simply providing protection to a host. We conclude with a number of case studies, describing work currently in progress, which demonstrate two very different application areas in which the mechanisms described are being applied to illustrate our point.

1. INTRODUCTION

Computational devices are now ubiquitous — in every home, in every office, on every street, numerous small and inexpensive devices are becoming capable of spontaneously networking to each other and to the internet, bringing with this ability the potential of a new age in computing which is data, rather than technology, driven. However, as access to information becomes universal, with vast quantities of information stored transparently across multiple, distributed sources, a paradigm shift in the way that systems are developed and managed is required. Current methodologies in software engineering will not suffice when attempting to develop networks that are self-managing, self-diagnostic, and above all, robust and secure in ever changing dynamic and unpredictable environments. This view is reinforced by

Zambonelli and Panurak who note that *the complexity introduced to software systems by several emerging computing scenarios goes beyond the capabilities of traditional computer science and software engineering abstractions* [45].

At first glance, these hurdles may seem unsurmountable when considered in engineering terms, yet many natural systems deal highly successfully with such problems on a continuous basis. For example, the central nervous system and the immune system function both highly successfully in an autonomic manner, and are both decentralised, scalable, robust, low-maintenance — items on the wish-list of any software-engineer. Therefore, the computational world is increasingly looking to the biological world for solutions.¹ This is evidenced for example by Carreras *et al*, who discuss biological approaches to autonomic communication systems in [15], Chess *et al* who describe immune-based approaches to security in [16], and the work performed as part of the BISON project [2] which aims to explore how a number of biological processes such as reaction-diffusion, chemotaxis, and replication might be modelled to enable the construction of robust and self-organizing information systems for deployment in highly dynamic network environments. Other biological systems have been studied in detail, in search of inspiration; evolution, ant-colonies and swarm systems. In this paper, we turn our attention to one of the most remarkable feats of nature — the vertebrate immune system.

The use of immune-inspired mechanisms in computing environments perhaps pre-dates the advent of autonomic computing itself. In fact, immunology first provided inspiration for building computer systems in the early 1990s, with attempts to build virus detection systems for machines [24, 23]. These were seminal works at the time - showing that ideas inspired by the natural immune system could be transferred with some success to machines in engineering (rather biological) scenarios. Artificial Immune Systems (AIS) went on to evolve in to a field in its own right, with application areas diversifying into classification, optimisation and robotics. However, the holy grail of providing computer systems with their own digital immune system still remains, inspiring a plethora of research papers every year.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AUTONOMICS 2007, 28-30 October 2007, Rome, Italy
Copyright © 2007 ICST 978-963-9799-09-7
DOI 10.4108/ICST.AUTONOMICS2007.2229

¹though it should perhaps be noted that biological systems are often stochastic and difficult to analyse, which are not such desirable properties from an engineering perspective

A quick survey of the AIS literature reveals that the majority of computational mechanisms and application domains are chosen based on the common assumption that immune system's major function is to separate "self" from "non-self", where *self* defines the "normal" body of the immune system host and *non-self* everything else. Indeed, to quote Wikipedia "*The immune system defends the body by recognizing agents that represent self and those that represent non-self, and launching attacks against harmful members of the latter group*". Although the self/non-self distinction can be semantically relabelled as dangerous/acceptable, interesting/useless, normal/abnormal, depending on the particular application in question, all versions of this theory essentially take the same stance: the immune system acquires a definition of itself during its early ontogenesis, and from then learns to ignore anything classified as "self". By default, anything that is then recognised by the immune system must be non-self and is attacked, thus protecting the host. This theory is commonly referred to as "Classical immunology" and forms the basis of the majority of immunology textbooks today, despite the fact that many obvious objections can be raised against it (for example, the relatively high occurrence of auto-immune diseases, the lack of rejection of fetuses by mothers in pregnancy etc.). It is unsurprising then that the majority of computer scientists in the early days of AIS attempted to exploit this metaphor in computational algorithms, and that today, in the context of autonomous computing, the immunological metaphor is generally directly linked with the topics of providing protection and security to future systems.

However, despite the ubiquity of the self/non-self theory of immunology which really became pre-dominant in the 1970s, alternative "camps" exist in the immunological world, each with sizeable followings, and proposing a number of different theories which question not only the mechanisms by which the immune system is held to operate but more fundamentally, the actual role of the immune system itself. While host defence is clearly a critical function, [17, 37] have proposed that it can not be the only function of interest. Indeed, the immune system might be regarded as primarily fulfilling an altogether different role if its phylogeny is carefully examined [38].

In a radical departure from accepted thinking, recent work by leading immunologist Irun Cohen has suggested that the immune system plays a much more fundamental role in the body than simply protecting it from harm, and instead, its function is that of *body maintenance*. Although this may seem a subtle semantic point, in this view (expounded in detail by Cohen in [18]), detection of harmful situations is merely a special case of overall body maintenance. The immune system instead is seen as a *cognitive system* which continuously provides body maintenance. The term maintenance covers a wide and diverse spectrum of functions, ranging from healing of cuts and bruises, to inflammation, to mending broken bones. To achieve these capabilities, the immune system must have detailed knowledge of its current internal state and the external environment, be re-configurable, adaptive, secure, and to provide self-healing functionality, in order to deal with an unpredictable and dynamic environment.

Although this alternative view has yet to filter through to the majority of computational research, we propose that it offers a much richer metaphor with which to tackle the problems inherent in building autonomic systems. As we explain in the following section, the requirements of autonomic computing systems fit well with the cognitive view of the immune system proposed by Cohen, and the metaphor may provide a means of tackling a wider range of problems than just security. Security is of course a crucial issue in autonomic computing, which cannot be ignored, and certainly we do not claim that immune-metaphors will not be useful in this respect. However, the goal of this paper is to elucidate some alternative immune mechanisms which may inspire research which can overcome some of the other hurdles associated with the developing field. We continue with a short review of the requirements of an autonomic system, and then attempt to outline some of the potential mechanisms which may be exploited in the future. This is followed by a number of exemplars of research currently being performed which is attempting to utilise the proposed mechanisms.

2. AUTONOMIC SYSTEMS

Although autonomic computing is still a relatively young field, which is undergoing continuous transformation and mutation as the contributing technologies mature, a number of characteristics can be defined which are likely to be common to autonomic systems of the future, whatever form they may take [28, 1]. In this section, we outline these characteristics, and attempt to map these characteristics to a more holistic view of the immune system.

An autonomic computing system needs to "know itself"

. Discarding the view that the immune system can discriminate neatly between two categories of substance labelled as self and non-self, a number of immunologists have proposed exactly that the immune system "knows only itself", i.e it knows only its own components [40, 27]. This view arises naturally from the proposals of both Jerne and Cohen, described in more detail in the next section. By knowing itself, and via a continuous "conversation" or exchange of signals with the body, the immune system is able to respond to perturbations of itself above a certain threshold. Thus in this view, nothing is foreign *per se* and it is only as external observers of a system that we can apply labels such as "self" and "non-self". This view is much more appealing to engineering autonomic systems, in which it is clearly impossible to label *a priori* system states in one category or another, particularly given the dynamic nature of system components and connection topologies.

An autonomic computing system must configure and re-configure itself under varying and even unpredictable conditions.

The immune system is under constant challenge to respond to an infinite number of unpredictable situations, and must respond with an appropriate behaviour. It is clear that the immune system is capable of adapting and changing accordingly, and that it exhibits a range of behaviours, varying from a full immune response to something akin to a mild irritation.

An interesting insight into what might drive a system to configure and re-configure itself is given by the *theory of autopoiesis*, originally introduced by Chilean biologists Humberto Maturana and Francisco Varela [4]. The term autopoiesis, which literally means "auto (self)-creation" expresses a fundamental dialect between structure and function. Maturana and Varela go on to define an autopoietic machine as "... a machine organized (defined as a unity) as a network of processes of production (transformation and destruction) of components which: (i) through their interactions and transformations continuously regenerate and realize the network of processes (relations) that produced them; and (ii) constitute it (the machine) as a concrete unity in space in which they (the components) exist by specifying the topological domain of its realization as such a network." . The implications of this theory for autonomic computing in general are discussed in section 3.4.

An autonomic computing system must detect, identify and protect itself against various types of attacks to maintain overall system security and integrity. Defending its host against attack is clearly *one* of the functions of the natural immune system, which allows species to survive. The function is equally critical in autonomic systems — the immune metaphor can offer mechanisms for engineering such functionality, but offers much more besides.

An autonomic computing system must perform something akin to healing - it must be able to recover from routine and extraordinary events that might cause some of its parts to malfunction. The immune system is remarkably robust to major perturbations. For example, amputation of a limb (and therefore a major percentage of the immune system's components) in the majority of cases simply results in re-organisation of the system itself. According to Cohen [18], wound healing, tissue repair, and cell regeneration are just some of the maintenance processes in which the immune system is involved. These fundamental processes cannot simply be modelled by a simplistic view of the immune system that considers in isolation only those mechanisms which it allow it to respond to attacks.

An autonomic computing system must know its environment and the context surrounding its activity, and act accordingly. The immune system, like most biological subsystems, does not exist in isolation, but instead is part of a larger whole. Tight coupling and integration with the external environment and other systems means that the immune system must constantly adjust itself in relation to these other systems. In effect, the reactivity of the immune system is determined by context; an event in one context may provoke a particular type of response, the same event in a different context by result in no response at all.

We can again turn to the autopoietic theory of Maturana and Varela for another perspective on the tight relationship between the immune system, the environment and subsequent actions of the immune system. The philosophical slant on the argument given is that the external environment is not an independent, objective entity from the system, but

that rather that the two mutually define each other. Clearly, two different systems with different sensory and motor abilities perceive a very different reality from the same environment, and act accordingly. The choice of action from one system need not make sense in the domain of the other: experience is entirely phenomenological and thus cognition is a process of constructing the world from experienced regularities.

Matzinger [33, 32] offers an additional insight into the relevance of context and the environment, in another controversial (though not mutually exclusive) immune theory which is known as the Danger theory. Matzinger refutes the self/non-self dichotomy, citing the numerous situations which cannot be satisfactorily explained by this theory, and instead proposes that the immune system responds not to foreignness *per se*, but to *danger*. Thus, it is the *context* in which a molecule is detected which is relevant; if the molecule causes damage, it should be responded to, if not, it can be safely ignored. This theory is discussed in more detail in section 3.

An autonomic computing system cannot exist in a hermetic environment. The immune system is itself heterogeneous, and functions in a heterogeneous environment. The components of the immune system are common to all individuals of a particular species, yet the immune system of each individual is unique, having adapted to its own particular local host. This is an appealing feature to engineers attempting to build flexible, rather than proprietary solutions, offering a bottom-up, rather than top-down approach to design.

We note that in addition to the above properties, [1] define two further properties of autonomic system that have no immediate parallel with the immune system. For completeness, these properties are mentioned. The first of these properties relates to the fact that an autonomic computing system always looks for ways to optimize its workings. It is not clear that the immune system is optimized in any sense — in many ways it appears remarkably inefficient, and exhibits a large amount of redundancy. This may simply be an artefact of evolution — unfortunately autonomic systems cannot afford the luxury of time afforded by evolution to the immune system and hence may need to offer solutions that are more optimal in some sense than that provided by the immune system. However, it should be noted perhaps that it is unclear to what extent optimality is really desired in a system — an optimal solution is often a fragile solution, prone to collapse if conditions change even slightly. Often, what is more desirable is a good, robust solution. In the context of a dynamic environment, the issue of optimality also becomes less relevant, as the definition of what is optimal at all changes as the system and environment change.

The remaining property states that an autonomic computing system will anticipate the optimized resources needed while keeping its complexity hidden. Clearly, the complexity of the immune system is transparent to the host that it ultimately protects — whether it anticipates resource usage requires a deeper understanding of the immune system than is appropriate in this article.

3. ALTERNATIVE IMMUNE THEORIES

As already alluded to in the introduction, the computer-scientist attempting to gain inspiration from the world of immunology is immediately confronted with a dilemma: immunologists fail to agree on the method and indeed the purpose of the immune system. It is not the aim of this paper to elucidate all of the potential theories — there are a number of accessible texts which can be referred to for details of the most common theories (e.g. [5, 38, 10]). An interesting overview from the immunological perspective is given in [30] in a volume of the journal *Seminars in Immunology*. This volume brings together competing immunological theories for the first time. Alternatively, the mechanisms are described in great detail from a computational perspective in [8, 39] and are generally discussed in the introductory sections of the majority of the literature in AIS. Instead, the less familiar mechanisms, not yet prolific in the computational literature but which offer a new perspective for exploiting biological mechanisms in the development of autonomic and pervasive computational systems, are summarised. It is intended that by outlining these mechanisms, the paper will stimulate and promote further investigation into the applicability of immune-inspired computing in autonomic system research.

The major immunological models which have provided inspiration for computational research in the past are clonal selection, immune networks and negative selection. Clonal selection and negative selection are the corner-stones of the self/non-self discrimination theory, and form the bulk of computational AIS models. The reader is referred to texts such as [39] for further details than have already been given in this paper. It is notable that clonal selection is concerned mainly with the *adaptive* immune response, that is the specific response of the system to unknown foreign invaders. In fact, there are two basic types of immunity; the *innate* immune system provides a non-specific, general response and can be considered as a first line of attack (in the “classical” immune model). Any invaders not removed by the innate immune system are then tackled by the *adaptive* immune response, which is the response most familiar to computer scientists. This response generates clones which evolve on a short time-scale to specifically target features on rogue pathogens. Although interest tends to focus on adaptive immunity in the digital world, in fact, endowing autonomic systems with a more general *innate* immune response would provide a robust foundation on which to provide systems with maintenance mechanisms and ensure basic security.

Immune networks represent a competing branch of immunological thinking which tends to be dismissed by many immunologists, as the empirical evidence required to justify the theory is somewhat lacking. However, they provide a compelling model for many engineering solutions and in fact may be undergoing something of a renaissance as techniques from complex network theory become more widespread and available for analysing network data (e.g [25]). This model is described in the next section. This is then followed by a description of some of the more recent immunological theories; the *Cognitive Immune Model* of Irun Cohen, and the *Danger Theory* model of Polly Matzinger. The relevant features of these models are described within the context of engineering systems in the future, rather than in detailed immunological

terms.

3.1 The Idiotypic network

In 1974, Niels Jerne proposed that the immune system proposed a radical alternative view to the theory of self/non-self discrimination that had by then become adopted as fact in immunological circles. Jerne proposes that the immune system is in fact a self-organising network in which immune cells recognise each other as well as pathogenic material, and in which all cells are treated equally. In this model, there is no distinction between a “self” cell and a “non-self” cell — the dichotomy collapses and the system is “complete unto itself” [14]. In the classical view, immune cells cannot persist in the absence of pathogenic stimulus — they are constantly randomly created by the bone marrow, and die naturally if not stimulated. In the Jernian view, immune-cells recognising other immune-cells create a self-sustaining and adaptive network which persists in the absence of any stimulation. The network effectively represents the memory of the immune system. The network continuously responds to perturbations in the system — perturbations above a certain threshold are considered as foreign and provoke a response. Foreignness thus depends directly on the current state of the network. Bersini [14, 6] argues strongly that this vision is far more empowering for computer scientists attempting to engineer systems and artefacts. He notes that the theory lead to the possibility of “*creating strongly adaptive systems, both parametrically and structurally, but whose adaptability aimed at satisfying endogenous constraints instead of just responding to exogenous impacts.*”. This leads to certain adaptive advantages: not only might have a system have the capacity to respond to larger diversity of external stimuli, but it offers an economical method of memorizing a repertoire of adapted responses when facing a non-stationary environment.

3.2 Danger theory

As previously noted in section 2, the Danger Theory of immunology was introduced by Matzinger in [33] in yet another challenge to the accepted wisdom of self/non-self discrimination. Although this theory does not deny that the immune system may discriminate between two classes of molecule as part of an immune response, it states that there may be additional factors that help initiate the response. The theory proposes that the immune system reacts to certain chemicals which are given off by distressed cells when they become damaged by invading proteins or die via non-programmed cell-death. It is these chemicals or signals that initiate the response, rather than the invader itself. Matzinger identifies a number of candidate signals in the immune system which may act as the danger markers (for example heat-shock proteins), and proposes that the immune system then associates the danger with nearby cells, thereby classifying certain cells as dangerous. The response then actuates a number of other immune cells to attack the cells classified as dangerous; the cells activated by the response are not necessarily in close proximity to the danger itself. This theory is somewhat more appealing than the self/non-self idea to the engineer; rather than having to label an unquantifiable quantity of data as “normal” or “abnormal” it shifts the emphasis to quantifying the *effect* of an action on the system, rather than the action itself. This model is gaining momentum in the AIS — for example [29, 12] apply the model to

intrusion detection in networks.

3.3 The Cognitive Immune Model

Irwin Cohen's immune model, popularised in [18] identifies the immune system as a complex system. Taking a holistic view of the system, he puts forward a model of a complex, adaptive and reactive system which is capable of *cognition*. This is a fundamentally different viewpoint to the classical immunology of the previous decades. In line with Jerne and Varela, Cohen proposes that the primary function of the immune system is *maintenance*. This may take the form of healing and repair (such as mending broken bones) as discussed in section 2 or even defence. Neuman sums this up neatly in [35] "*Rather than promoting the metaphor of the immune system as a warrior that defends his castle against invaders.... the immune system is the maintenance man of the compartment building we call the organism*". Such functionality demands far more complexity from a system than defence, requiring examining the cells and interactions that comprise the immune system in a completely different light. Cohen reframes the entire problem of attempting to define the proposed goals of immunity by considering the entire immune system as a *computational* system. The computational question is then not "how which cells and interactions comprise the immune system" but rather "what is the immune system computing", which can only be answered by considering the *state* of the immune system [19].

Considering the immune system as cognitive, i.e that it computes states and effect actions given that state, may at first glance seem surprising. Cognition is usually thought of as a conscious process performed by the brain, but Cohen defines three elements that, when integrated, can result in a cognitive system without consciousness. The first element is the ability of the system to exercise *options*, that is, to make *decisions* based on a number of choices. In order to accomplish this, the system must contain *internal images* of its environment, which are updated based on its experience, gained via interactions or *self-organisation*. Update of the internal-images implies an increase of information which is driven through inputs of energy and information by the world, and similar outputs generated by the system. Eventually, these interactions result in choices, and therefore the emergence of cognition.

Cohen describes three important mechanisms which contribute to the emergence of cognition; *co-responsence*, *pleiotropia* and *redundancy*. The *co-responsence* concept suggests that in order to fulfill its role (maintenance, protection), the immune system maintains different classes of immune cells. These cells individually see different aspects of any object that may be of immune interest, from within the body (tissues) or external (antigens). Each class of immune cell informs other immune cells about what it has seen, by expressing co-response signals (cytokines, processed peptides, interaction molecules, antibodies). The effect of these signals, essentially, is that each cell modifies its own response based on the feedback it receives from the other cells [18] — essentially, although it is impossible for a cell to perceive what another cell perceives of its environment, it can perceive how another cells responds, and therefore respond to this response. This is a key point for autonomic systems in which individual components of the system may have little




Figure 1: Immune Computation (adapted from Cohen (2007))

storage capacity and limited communication abilities, rendering it impossible for them to communicate their perceptions in their entirety to other components of the system.

Pleiotropia denotes the capacity of a single immune component to produce several diverse effects. Depending on existing conditions, immune agents elicit different responses and do different, sometimes contradictory things. For example, in the natural immune system, a T cell can kill a target cell and stimulate the growth of another. This is another essential property for components of an autonomic system, in which components must be capable of effecting a number of different responses depending on current environmental conditions.

The final property, *redundancy* is distinguished as simple and degenerate. Simple redundancy designates the existence of multiple copies of the same element (e.g. numerous antibodies are produced during an immune response), while degenerate redundancy describes the situation in which many different immune components perform the same action. Both types are relevant to large, distributed networks containing multiple components. Assuming unreliability of individual network elements, multiple copies of the same *network component* will inevitably exist in any application, and in addition, different components may respond in a similar manner to certain environmental conditions.

This is a short summary of the Cognitive Model — the immunological basis is described in detail in [18] and summarised by Andrews in [11] for the computing community. Perhaps however the theory (and its potential use) is summed up most succinctly by Cohen himself in purely computational terminology in [19] and depicted in 1. "*The immune system functions to compute the state of the body and effects actions accordingly. The input to the immune system is the state of the body which comprises of a collection of molecular signals (indicating for example inflammation or trauma). The output is a particular response state which triggers the appropriate processes. The output is fed back to the tissues (inducing healing for example), but also to the immune system itself which modifies both its structure and behaviour. The response is thus formulated as a result of the cumulative experience of the immune system dealing with both the body and the world.*". In short, the cognitive view is indeed compelling as a framework for building autonomic systems which can continuously compute their state and act accordingly. In section 4, we describe some current research which is attempting to exploit this metaphor.

3.4 Autopoiesis and Enaction

As a final thought, we return to the theory of *autopoiesis* which was mentioned in section 2, in an attempt to define the close-coupling between state, structure and environment. The theory is worth reflecting on, as it offers a useful perspective which may have wide application across many scientific fields, but particularly in the domain of autonomic computing. The crucial distinguishing feature of the theory is that it provides a framework that can be applied from the level of organisation of organic matter, through cellular systems, tissues and organs, to species and societies.

At each level, order emerges through a systems self-organisation, rather than being *a priori* given which the system must then learn. Learning and development at each layer is the unconscious activity of successfully adapting around the order and constraints of the surrounding layers. The term unconscious is used in the sense that the system infers the external reality of itself and its environment through the internal effects on its structure, caused by the cyclic internal processes of action and perception. The system is its own cognitive architecture, defined by its innate features and continued evolution [40].

Most biological systems are autopoietic. In contrast, *mechanical* systems primarily adapt at a behavioural, or macro, level through interaction with the environment and are as such termed *allopoietic*. Most mechanical systems exhibit limited or zero ability to adapt in any physical sense, being constrained by fixed components which cannot grow and evolve unlike in natural systems. This adaptability at a 'hardware' as well as a 'software' level endows biological systems with the properties desirable in an autonomic system — for example, robustness, fault tolerance, and sustainability. This is often overlooked by engineers who tend to attempt to achieve such functionality through software alone. Furthermore, according to Sharkey *et al* [7], the lack of evolvable capability is directly related to a lack of autonomy in mechanical systems. Biological systems on the other hand are fully autopoietic and fully autonomous. Autopoiesis is therefore a desirable property in any autonomic system; as a specific example, in pervasive computing environments, the ability to evolve will be crucial in order to be able to both function in heterogeneous and dynamic environments and to integrate future emerging technologies. An autopoietic perspective may thus bring useful insights to the bio-inspired designer of future systems.

In the following sections, we provide examples of research being performed in two diverse application areas which serve to illustrate how some of the mechanisms just mentioned might be utilised in autonomic applications.

4. CASE STUDY 1: SPECKLED COMPUTING.

Advancements in micro-electro-mechanical systems technology, wireless communications, and digital electronics have enabled the development of exceptionally small mechanical devices that are inexpensive, low-power and sense phenomena in the physical world [13]. Such devices can be connected together in large numbers to form wireless sensor networks (WSNs). The Speckled Computing Consortium [3] is dedi-

cated to the realisation of a new WSN platform of minute dimensions, which promises to create minute semiconductor grains, called *specks*, around the size of one cubic millimetre [42]. Each minute *Speck* contains its own processor, memory, and communication hardware, and can be equipped with one or more sensors such as pressure, temperature or accelerometers. This promises a new generation of "spray-on computers" [42], in which dense networks, *SpeckNets*, consisting of thousands of nodes can be created.

Specks may be scattered or sprayed on the person or surfaces, and act as a "computational aura" [3], opening up a plethora of potential applications, perhaps using tens to thousands of specks. For example, a few tens of specks may be attached to rigid object allowing tracking of the position and orientations of articulated rigid bodies [44]. Normally passive artefacts such as furniture and appliances might have Specks incorporated into their structure, thus enabling them to interact with users — Wong *et al* [43] describe a scenario in which specks integrated into a reading table might detect the removal of a book from a table and automatically turn on a reading light. On detection of a fire, a Speck system might automatically drop thousands of Specks from the ceiling onto the floor, where they would self-organise in order to light up pathways to the nearest exit. The technology clearly promises much, yet despite great advancements in hardware and nano-technology enabling such devices to be built, the software required to enable these devices to function in a useful manner lags far behind. A *SpeckNet* is an autonomous computing machine which must achieve some specified task and simultaneously organise and maintain itself. Although the ultimate goal of the SpeckNet Consortium is to create minuscule semiconductor grains around the size of one cubic millimetre [42], the first generation of these units, currently in existence, have a dimension of approximately five millimetres, and are known as ProSpeckz (Programmable Specks over Zigbee Radio). The main components that constitute a speck are:

- a computer-on-a-chip, that combines a micro-processor and memory (FLASH and RAM),
- a radio chipset, which is compatible with the IEEE 802.15.4 standard for low rate wireless personal area networks [9],
- an antenna that allows communication ranges from a few centimetres to over a few meters,
- a power supply, such as compact rechargeable batteries, and
- a number of sensors, which vary from heat sensors to accelerometers, depending on the type of application.

Sensors carried by each speck, perform data capture, while in-built processing capabilities permit specks to filter data and extract information from the environment and LED components provide specks with a feedback mechanism. Distributed functioning in SpeckNets is enabled by incorporating communication capabilities, which allow the constituent devices to interact with each other. The lack of powerful

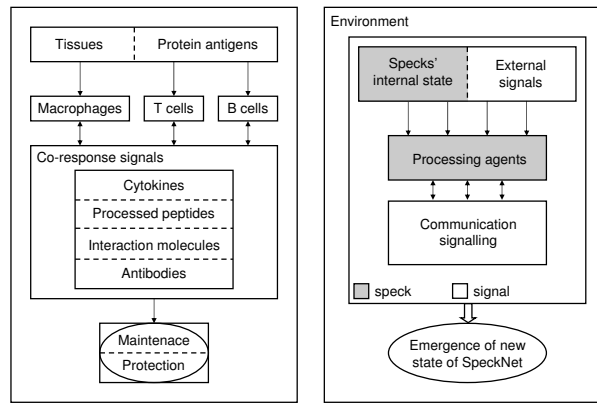


Figure 2: The proposed mapping between Cohen’s cognitive framework and a SpeckNet

base station units forces the network to operate in a decentralised and asynchronous fashion, by sharing tasks and processes between its autonomous units.

The physical characteristics of speck devices inevitably impose restrictions on the functionality of SpeckNets and of any software which might control them, rendering traditional software engineering methodologies impossible. Firstly and foremost, the minute dimensions of a Speck affects all constituent parts of a speck node, and consequently, the overall performance and efficiency of SpeckNets. The small physical dimensions imply limited space for loading on board memory and power supply, and as a result, storage capacity of specks is extremely limited (about less than 10Kb of FLASH memory and 1Kb at most of RAM). Processing power is also restricted because of the above constraints, and especially due to the memory constraints. Furthermore, wireless communication is not only unreliable but is of limited range. To minimise the size of the antenna (13×6mm), operation is set at high frequencies (2.4GHz), which however incurs high path losses [41]. Yet the most crucial constraint in SpeckNets is to minimise power consumption on every design level [20], which is also a common problem many in embedded sensor networks [26]. Technically, the most expensive activity in small sensor nodes is radio usage. Data transmission and reception usually draw significant amounts of power, in comparison with processor and sensors [22, 42]. Sending a single bit can consume the same energy as executing 1000 instructions [31]. Energy drain leads to speck “death”, which is also exacerbated by the placement in open environments, where it is not always possible to fully control operating conditions. Thus, unpredictable node failures can occur for a number of reasons, which has an unpredictable effect on the functioning of the speck network.

Clearly this is a challenging problem. However, in recent work in [21] we describe in detail how the functional challenges faced by a Specknet are common to those faced by the immune system. Furthermore, both the Specknet and immune system have to address these challenges faced with a system which comprises of vast numbers of individually weak, and unreliable components. We show how the cognitive model of Cohen described above potentially provides the foundation for an architecture which will allow a specknet to both *regulate its operation* so that it remains active as

long as it can (i.e. maintain itself) and *fulfill its application goals*.

In the proposed architecture, (depicted in figure 2) as in the immune system, specks perform one or more of three basic functions; they can *sense* information from their environment; they *process* information; and they *communicate* information, i.e. carry signals. In both the immune and specknet environments, cells can be *active* or *resting* (corresponding to an idle state in a speck). In addition, specks can be in a further state which has no direct immunological analogy, of *sleeping* in which a particular component is turned off and therefore does not interact with the system. As we have seen in the discussion above, processes occurring in the tissues of a host provide a context to signals received by the immune system, for instance indicating infection or damage. This can be seen as analogous to the *internal* state of a speck, corresponding to the current state of (for example) its battery, energy levels, malfunctioning or failure of any of its component and its relative location to other specks (which can be calculated by existing algorithms, for example [34]). However, specks exist in dynamic, changing environments which provide them with additional input data, corresponding to antigen data in an immune system. For example, sensors on a speck may provide information regarding humidity, light, pressure, or temperature. This can be viewed as *external* information. External signals may ultimately affect the internal state of a speck, e.g. the environment may renew power supplies of specks (e.g. provide energy captured by solar cells) or may cause permanent damage to the network, by breaking some nodes.

Internal and external information is continuously gathered by each type of speck in a network. This information is filtered, integrated and processed by individual specks according to their designated *type* — similarly, internal and external signals are gathered and processed by a number of different types of cells in the immune system, e.g. macrophages, t-cells and b-cells. This requires the receiving and generation of *signals* which are communicated to local specks and eventually may spread globally. As in the immune system, data is not generated, sent or received in a sequential or orderly fashion. Communication is asynchronous, thus the system is independent of signalling time scale. Collectively, specks undertake the responsibility of collating data (collect broad-

casts), filtering signals (accept some, reject others), and processing them (transform them appropriately, e.g. calculate average value over a specific time window).

The framework allows the system to be continuously active and results in the emergence of a new state of SpeckNet. Depending on the current conditions, this new state can be either of maintenance (e.g. specks work in a power saving mode), in which the aim of the system is to stay alive for the longest possible or of response, in which a noteworthy change in the environment activates the system and if necessary causes a reaction. In this first instance, the framework is being developed in simulation, and experiments are being conducted based on utilising Specks to detect fire in an enclosed space and light up appropriate pathway to an exit. We hope to report further results of this research (currently described in more detail in [21]) in the near future. However, this example serves to illustrate just one potential use of immune-inspired methods in computational systems.

5. CASE STUDY 2: SELF-ORGANISING OVERLAY NETWORKS

Further work being performed within the Centre for Emergent Computing at Napier University (although in early stages) is exploring the applicability of immunological mechanisms to decentralised machine-learning and network management, in order to construct decentralised, self-organising overlay networks to aid information retrieval. Such unstructured logical network substrates need to maintain a functional steady-state while in constant flux of peer membership, exploitable resources and connection topology. This problem domain is synonymous with many of the existing problems in autonomous wired decentralised systems and ad-hoc wireless mobile systems. Such extreme conditions have proven antagonistic to traditional systems engineering, but can be considered the norm in many naturally occurring systems. The work is outlined here again as an exemplar of the potential application of immune mechanisms in a variety of diverse domains.

Large heterogeneous information networks present significant engineering, financial and administrative issues for centralised information retrieval systems: distributed content; volume of information; rate of change; scale of growth; lack of schema; interoperability; maintenance; and not least, the inherent complexity of human semantics. In order to decentralise the process, a dual perspective can be taken. An engineering perspective gives a baseline formalism, but is undesirable as a solution, due in the main to lack of adaptability. A bio-inspired perspective however may separately provide inspiration for addressing two the major issues concerned with this type of application; how to deal with *behaviour* and *learning* in the network. From a behavioural perspective, the network structure must exhibit robust, adaptive self-organisation. However, this must be driven by *learning* in the network — some kind of understanding of the changing information consumption and production patterns at its edges. This is essentially a cyclic process: learning affects behaviour which affects learning — the two are strongly coupled and this work is heavily influenced by the concepts of autopoiesis discussed in the previous section.

Although the biological influence is clear, it remains to be

determined exactly how such learning and behaviour might take place on a large, distributed network. Existing “bio-inspired” and “multi-agent” research often only appeals to anthropomorphic descriptions of its agents motivation, rather than capture any underlying, unconscious mechanism. This type of approach is in reality only a small shift from traditional engineering — instead this research takes insight directly from theoretical immunology. The problem domain is framed in a metaphor of the vertebrate immune-system — itself a robust, adaptive system capable of seemingly cognitive tasks. Ultimately we hope to show the applicability of these mechanisms in decentralised machine-learning, information retrieval and network management. To achieve this, a number of central research questions need to be addressed. These questions, which are listed below and currently directed at information networks, also have more general relevance in the wider domain of autonomic computing.

- Can immune-inspired mechanisms be exploited in order to capture the time and space dynamics of an on-line information system ?
- Can cell binding networks provide an abstract, non-geometric representational space suitable for capturing coarse-grained semantics in an information network ?
- Can the communication and feedback model utilised by the immune system be abstracted into an event-based protocol appropriate for local and distributed communication ?

The ultimate goal of such research is to create a network which can self-organise around the patterns of information which occur on it, by delegating responsibility of usage and maintenance to those who contribute to it, and whose needs the network exists to serve. Such a feature is surely central to pervasive computing systems of the future.

6. CONCLUSION

This paper has outlined the key characteristics that autonomic systems must exhibit, and attempted to map those characteristics to those apparent in the natural immune system. By outlining a number of alternative theories of immunology, we have tried to raise the consciousness of computer scientists to the plethora of immune functions and methods that may be of use when designing bio-inspired systems. In the past, immune-inspired has tended to adopt a somewhat naive approach to extracting metaphors to the immune system — guilty of “reasoning by metaphor” according to Stepney *et al* [36]. This has led perhaps to the tendency of those outside the immunological world to attempt to capitalize on popularist theories of immunology, reflected in the emphasis in applying immunological mechanisms to security and machine protection. The immune system appears to play a far greater role in our bodies than simply that of defence — we hope that by considering this greater role, and in addition, taking a holistic approach which examines the interactions of the immune system with other biological sub-systems such as the neural or endocrine systems, a new era in computing will be kick-started. Exploiting the full potential of the immunological metaphor will require close interaction of computer scientists, engineers, mathematicians

and immunologists. Nevertheless, the field remains ripe for exploration.

7. ACKNOWLEDGEMENTS

The Specknet research is funded by the Specknet consortium, SHEFC HR03007

8. REFERENCES

- [1] <http://research.ibm.com/autonomic/index.html>.
- [2] Biology-inspired techniques for self-organisation in dynamic networks. <http://www.cs.unibo.it/bison/index.html>.
- [3] Specknet. <http://www.specknet.org>.
- [4] *Autopoiesis and Cognition*. Reidel, 1980.
- [5] *Immunobiology: the immune system in health and disease*. 3rd Edn. Garland Publishing, New York, 1997.
- [6] *Artificial Immune Systems and their Applications*, chapter The endogenous double plasticity of the immune network and the inspiration to be drawn for engineering artefacts. Springer Verlag, 1999.
- [7] *Symbolic and Neural Net Hybrids*, chapter Life, mind and robotics: The ins and outs of embodied cognition. MIT Press, 2000.
- [8] *Artificial Immune Systems: A New Computational Approach*. Springer, 2002.
- [9] IEEE Standard for Information Technology, IEEE Std 802.15.4-2003. World Wide Web, 2003. Retrieved January 13, 2007 from <http://standards.ieee.org/getieee802/802.15.html>.
- [10] *In Silico Immunology*. Springer, 2007.
- [11] *In Silico Immunology*, chapter Alternative Inspiration for Artificial Immune Systems: Exploiting Cohen's Cognitive Model. Springer, 2007.
- [12] U. Aickelin, J. Greensmith, and J. Twycross. Immune system approaches to intrusion detection: A review.
- [13] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: A survey. *Computer Networks (Elsevier) Journal*, 38(4):393–422, March 2002.
- [14] H. Bersini. Self-assertion vs self-recognition: a tribute to francisco varela. In *Proceedings of 1st International Conference on Artificial Immune Systems*, pages 107–112, 2002.
- [15] I. Carreras, I. Chlamtac, F. De Pellegrini, C. Kiraly, M. Miorandi, and H. Woesner. A biological approach to autonomic communication systems. In *Proc. of Converging Science*, 2004.
- [16] D. Chess, C. Palmer, and C. White. Security in an autonomic computing environment. *IBM Systems Journal*, 2003.
- [17] I. Cohen. The cognitive paradigm and the immunological homunculus. *Immunology Today*, 1992.
- [18] I. Cohen. *Tending Adam's garden: evolving the cognitive immune self*. Elsevier Academic Press, 2000.
- [19] I. Cohen. Real and artificial immune systems: computing the state of the body. *Nature*, 2007.
- [20] L. H. Crockett, N. C. MacEwen, E. Pfann, and R. Stewart. A low power, digital transceiver for wireless sensor networks. In *Proceeding of the 2nd IEEE/Eurasip Conference on DSP Enabled Radio, London, UK*, pages 18/1–18/6, September 2005.
- [21] D. Davoudani, E. Hart, and B. Paechter. An immune-inspired approach to speckled computing. In *Proceedings of ICARIS 2007*. Springer, 2007.
- [22] D. Estrin, L. Girod, G. Pottie, and M. Srivastava. Instrumenting the World With Wireless Sensor Networks. In *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2001), Salt Lake City, UT, May 2001*, volume 4, pages 2033–2036, May 2001.
- [23] S. Forrest, S. Hofmeyr, and A. Somayaji. Computer immunology. *Commun. ACM*, 40(10):88–96, 1997.
- [24] S. Forrest, A. Perelson, L. Allen, and R. Cherukuri. Self-nonsel self discrimination in a computer. In *Proceedings of the IEEE Symposium on research, security and privacy*, pages 202–12, 1994.
- [25] E. Hart, H. Bersini, and F. Santos. How affinity influences tolerance in an idiotypic network. *J. Theor. Biology*, 2007.
- [26] J. Heidemann and R. Govindan. An overview of embedded sensor networks. Technical Report ISI-TR-2004-594, USC/Information Sciences Institute, November 2004.
- [27] N. Jerne. Towards a network theory of the immune system. *Annals of Immunology (Inst. Pasteur)*, 125:373–389, 1974.
- [28] J. Kephart and D. Chess. The vision of autonomic computing. *Computer*, 2003.
- [29] J. Kim, P. Bentley, U. Aickelin, J. Greensmith, G. Tedesco, and J. Twycross. Immune system approaches to intrusion detection. *Natural Computing*, 2007.
- [30] R. Langman and M. Cohn. Editorial introduction. *Seminars in Immunology*, 12(3):159–162, 2000.
- [31] P. Levis and D. Culler. Maté: A tiny virtual machine for sensor networks. In *ASPLOS-X: Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 85–95, New York, NY, USA, 2002. ACM Press.
- [32] P. Matzinger. The danger model: a renewed sense of self. *Science*, 2002.
- [33] P. Matzinger. Tolerance, danger and the extended family. *Ann. Rev. Immunology*, 1994.
- [34] R. McNally, K. J. Wong, and D. K. Arvind. A distributed algorithm for logical location estimation in speckled computing. In *In Proceedings of the IEEE Wireless Communications & Networking Conference 2005, USA*, March 2005.
- [35] Y. Neuman. the immune self, the sign and the testes. *Semiotics, Evolution, Energy, Development*, pages 85–109, 2005.
- [36] S. Stepney, R. Smith, J. Timmis, A. Tyrrell, M. Neal, and A. Hone. Conceptual frameworks for artificial immune systems. *International Journal of Unconventional Computing*, 1(3):315–338, 2005.
- [37] J. Stewart. Cognition without neurons: adaptation, learning and memory in the immune system. *CC AI*.
- [38] A. Tauber. The biological notion of self and non-self. In *Stanford Encyclopedia of Philosophy*.
- [39] J. Timmis. Artificial immune systems - today and tomorrow. *Natural Computing*, 6:1–18, 2007.

- [40] F. Varela, A. Coutinho, B. Dupire, and n. Vaz. Cognitive networks: Immune, neural and otherwise. *J. Theoretical Immunology*, 1988.
- [41] G. W. M. Whyte, N. Buchanan, and I. Thayne. An omnidirectional, low cost, low profile, 2.45 ghz microstrip fed rectaxial antenna for wireless sensor network applications. In *IEE and IEEE conference, Loughborough Antennas and Propagation Conference (LAPC)*, 2006.
- [42] K. Wong and D. Arvind. Speckled computing: Disruptive technology for networked information appliances. In *Proceedings of the IEEE International Symposium on Consumer Electronics (ISCE'04) (UK)*, pages 219–223, September 2004.
- [43] K. Wong and D. K. Arvind. Specknets: New challenges for wireless communication protocols. *Third International Conference on Information Technology and Applications, ICITA 2005*, 2:728–733, July 2005.
- [44] A. Young, M. Ling, and D. Arvind. Orient-2: A realtime wireless posture tracking system using local orientation estimation.
- [45] F. Zambonelli and H. V. D. Parunak. Signs of a revolution in computer science and software engineering. In *Engineering Societies in the Agents World III: Third International Workshop, ESAW 2002, Madrid, Spain, September 16-17, 2002. Revised Papers*, pages 13–28, 2002.