

iNet: A Biologically-inspired Adaptation Mechanism for Autonomic Network Applications

Chonho Lee, Paskorn Champrasert and Junichi Suzuki
Department of Computer Science
University of Massachusetts, Boston
{chonho, paskorn, jxs} @ cs.umb.edu

Abstract—This paper addresses an emerging issue in network applications, *autonomous adaptability*, and empirically evaluates a biologically-inspired adaptation mechanism, called iNet. Based on the observation that the immune system has elegantly achieved autonomous adaptation, iNet is designed after the mechanisms behind how the system detects antigens (e.g. viruses) and specifically reacts to them by producing antibodies. iNet models a set of environment conditions (e.g. network traffic) as an antigen, and an agent behavior (e.g. reproduction and migration) as an antibody. iNet allows each agent to autonomously sense its surrounding environment conditions (i.e. an antigen) to evaluate whether it adapts well to the sensed conditions, and if it does not, adaptively perform a behavior (i.e. an antibody) suitable for the conditions. Empirical evaluation results show that iNet works efficiently makes network applications adaptive to dynamic changes in the network.

1. Introduction

Network applications are expected to be more *autonomous* and *adaptive* to dynamic changes in the network (e.g. changes in network traffic and resource availability) in order to improve user experience, expand application's operational longevity and reduce maintenance cost [1, 2]. As inspiration for a new design paradigm for network applications, we observe various biological systems have already achieved autonomy and adaptability. We believe if network applications are designed after certain biological concepts and mechanisms, they may be able to increase their autonomy and adaptability.

The NetSphere architecture applies key biological concepts and mechanisms to design network applications. A network application is modeled as a decentralized group of autonomous software agents. This is analogous to a bee colony (an application) consisting of multiple bees (agents). Each agent implements a functional service and follows biological behaviors such as migration, replication, energy exchange and death.

This paper addresses autonomous adaptability of network applications (i.e. agents). The proposed adaptation mechanism, iNet, is designed after the mechanisms behind how the immune system detects antigens (e.g. viruses) and produces specific antibodies to kill them. iNet models a set of environment conditions (e.g. network traffic) as an antigen and a behavior of agents as an antibody. iNet allows each agent to autonomously sense its local environment conditions (i.e. an antigen) to evaluate whether it adapts well to the sensed conditions, and if it does not, adaptively perform a behavior (i.e. an antibody) suitable for the conditions.

For example, agents may invoke migration behavior for moving to network hosts that accept a large number of user requests for their services. This leads to the adaptation of agent population and locations, then agents can reduce re-

sponse time for users. Empirical evaluation results show iNet works efficiently in high degree of accuracy and makes agents adaptive to dynamic network environment.

2. NetSphere architecture

In the NetSphere architecture, agents are designed based on the three principles described below [3].

Decentralization: Agents are decentralized. There are no central entities to control and coordinate agents (no directory servers). Decentralization allows network applications to be scalable and simple by avoiding performance bottleneck and any central coordination in deploying them [4].

Autonomy: Agents are autonomous. Agents monitor their local network environments and autonomously behave and interact without any interventions from/to other agents, platforms and human users.

Adaptability: Agents are adaptive to dynamic environment conditions (e.g. user demands and resource availability). Each agent contains iNet, which allows it to adaptively behave to the current environment.

Each agent is implemented as a Java object and runs on a NetSphere platform. The platform is also implemented in Java and runs atop a Java VM on a network host [3]. Each agent consists of *attributes*, *body* and *behaviors* [3]. Attributes carry descriptive information regarding the agent (e.g. agent ID). The body implements a service the agent provides. For example, an agent may implement a genetic algorithm for an optimization problem. Behaviors implement non-service related actions that are inherent to all agents. Although NetSphere defines nine behaviors, this paper focuses on four of them [4].

Migration: Agents may move between platforms.

Energy exchange: Agents may gain *energy* in exchange for providing services to other agents or users. Agents may also expend energy for services that they receive from other agents, and for resources available on a platform (e.g. memory space).

Communication: Agents may communicate with other agents for the purposes of, for example, requesting services.

Lifecycle regulation: Agents may regulate their lifecycles. They may make their copies (replication) in response to higher energy level. They also may die as a result of energy starvation. If energy expenditure of an agent is not balanced with the energy gain, the agent will not be able to pay for the resources it needs, i.e., it dies from lack of energy.

3. The iNet artificial immune system




Figure 1. Organization of the iNet artificial immune system

This section overviews how the natural immune system works, and describes how the iNet artificial immune system is designed after the natural immune system.

3.1. Natural Immune System

The immune system is an adaptive defense mechanism to regulate the body against dynamic environment changes (e.g. antigen invasions). Through a number of interactions among various white blood cells (e.g. macrophages and lymphocytes) and molecules (e.g. antibodies), the immune system evokes two responses: *innate* and *adaptive* immune response.

In the innate immune response, the immune system performs *self/non-self discrimination* to detect antigens. This response is initiated by macrophages and T-cells, a type of lymphocytes. Macrophages move around the body to ingest antigens and present them to T-cells so that T-cells can recognize them. T-cells are produced in thymus and trained through the *negative selection process*. In this process, thymus removes T-cells that react with the body's own (self) cells. The remaining T-cells are used as detectors to identify non-self cells (i.e. antigens). Detecting non-self cells, T-cells secrete chemical signals to activate adaptive response.

In the adaptive immune response, the immune system produces antibodies that specifically kill an antigen identified by T-cells. Antibodies form a network structure and communicate with each other [5]. This network is formed with stimulation and suppression relationships among antibodies. Thus, the adaptive immune response is offered by multiple types of antibodies, although a single type of antibody (the best matched with an antigen) may play the dominant role. The immune network also helps to keep the quantitative balance of antibodies. Through the stimulation and suppression interactions, the population of specific antibodies rapidly increases following the recognition of an antigen and, after eliminating the antigen, decreases again. Performed based on this self-regulation mechanism, the adaptive immune response is an emergent product from many interactions among antibodies.

3.2. Design and implementation of iNet

The iNet artificial immune system consists of the *environment evaluation* (EE) and *behavior selection* (BS) facility (Figure 1) corresponding to the innate and adaptive immune response, respectively. EE allows each agent to continuously sense a set of current environment conditions as an antigen and examine whether it is self or non-self. A self antigen indicates that the agent adapts to the current environment

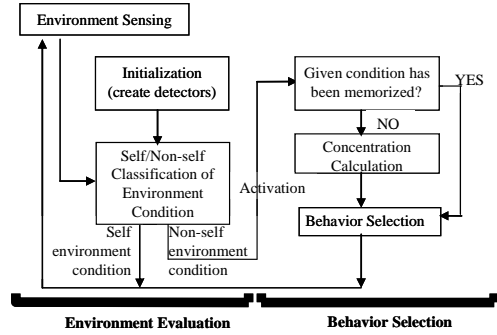


Figure 2. iNet Adaptation Process

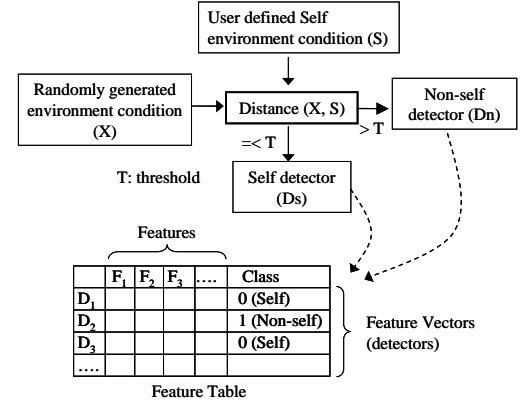


Figure 3. Initialization Step in EE

conditions well, and a non-self antigen indicates it does not. When EE detects a non-self antigen, EE activates BS (Figure 1). BS allows each agent to choose a behavior as an antibody that specifically matches with the detected non-self antigen.

3.2.1. Environment evaluation facility (EE)

EE performs two steps: initialization and self/non-self classification (Figure 2). The initialization step produces detectors that identify self and non-self antigens (i.e. environment conditions). In iNet, an antigen (i.e. a set of environment conditions) is implemented as a *feature vector*. Each feature vector (X) consists of a set of features (F) and a class value (C). F contains a series of environment conditions. If an agent senses agent population on a local platform, resource utilization on a local platform and workload (the number of user requests) on a local platform, a feature vector may be represented such as $X_{current} = ((Low: Agent population, Low: Resource utilization, Heavy: Workload), C)$. C indicates whether a given antigen (i.e. a set of environment conditions) is self (0) or non-self (1).

To evaluate whether an antigen (i.e. feature vector) is self or non-self, the initialization step produces detectors that identify them (Figure 2). This step is designed after the negative selection process in the immune system. In the initialization step, EE first generates feature vectors randomly, and separates them into self detectors, which closely match with self antigens (feature vectors), and non-self detectors¹, which do not closely match with self antigens (feature vectors). This separation is performed via vector matching between

¹ Non-self detectors in iNet are equivalent to T-cells in the immune system.




Figure 5. Antibody Structure

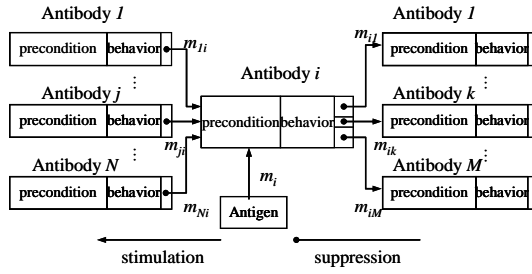


Figure 6. A Generalized Network of Antibodies




Figure 7. Example Network of Antibodies

randomly generated feature vectors and self antigens (feature vectors) that human users supply (Figure 3). Currently, EE uses the Euclidean vector matching algorithm. After vector matching, both self and non-self detectors are stored in a feature table (Figure 3)².

The second step in EE performs self/non-self classification of environment conditions (Figure 2). It uses the detectors in a feature table to classify the current environment conditions into self or non-self. The self/non-self classification step is performed with a decision tree built from detectors in a feature table. Based on the decision tree, EE classifies a current environment conditions, $X_{current}$. Once EE detects a non-self antigen, it activates BS immediately.

The reasons for using decision tree as a classifier are ease of implementation and algorithmic efficiency. Since a decision tree is easy to understand and implement, iNet can maintain a lower barrier for developers to design adaptive network applications. Also, a decision tree performs classification much faster than other algorithms such as clustering, support vector machine and Markov model algorithms [6]. The efficiency of classification is one of the most important requirements in iNet because each agent periodically senses and classifies its surrounding environment conditions.

3.2.2. Behavior selection facility (BS)

Once EE classifies the current environment conditions as a non-self antigen, it activates BS. BS selects an antibody (i.e. agent's behavior) suitable for the detected non-self antigen (i.e. environment conditions). Each antibody is structured as shown in Figure 5. It consists of *Paratope*, precondition under which it is selected (one of environment conditions), *Behavior ID*, one of agent behaviors, and *Idiotope*, relationships to other antibodies (one or more links). Antibodies are linked with each other using stimulation and suppression relationships (see Section 3.1). Each antibody has its own concentration value corresponding to the number of the antibody. The value is used to prioritize antibodies (behaviors) in behavior selection. BS identifies candidate antibodies (behaviors) suitable for a given non-self antigen (environment conditions), prioritizes them based on their concentration values, and selects the most suitable one from the candidates. When prioritizing antigens (behaviors), stimulation relationships between them contribute to increase their

concentration values, and suppression relationships contribute to decrease it. Each relationship has its own strength (affinity), which indicates the degree of stimulation/suppression.

Figure 6 shows a generalized network of antibodies. The antibody i stimulates M antibodies and suppresses N antibodies. m_{ji} and m_{ik} denote affinity values between antibody j and i , and between antibody i and k . m_i is an affinity value between an antigen and antibody i . The concentration of antibody i , denoted by a_i , is calculated as follows.

$$\frac{dA_i(t)}{dt} = \left(\frac{1}{N} \sum_{j=1}^N m_{ji} \cdot a_j(t) - \frac{1}{M} \sum_{k=1}^M m_{ik} \cdot a_k(t) + m_i - k \right) a_i(t) \dots (1)$$

$$a_i(t) = \frac{1}{1 + \exp(0.5 - A_i(t))} \dots (2)$$

In the equation (1), the first and second terms in a big bracket denote the stimulation and suppression from other antibodies. The affinity values between antibodies (i.e. m_{ji} and m_{ik}) are positive between 0 and 1. m_i is 1 when antibody i is stimulated directly by an antigen, otherwise 0. k denotes the dissipation factor representing the natural death of an antibody. This value is 0.1. The initial concentration value for every antibody, $a_i(0)$, is 0.01. The equation (2) is a sigmoid function used to squash the $A_i(t)$ value between 0 and 1.

Every antibody's concentration is calculated 200 times repeatedly. This repeat count is obtained from a previous simulation experience [7]. If no antibody exceeds a predefined threshold (0.7) during the 200 calculation steps, the antibody whose concentration value is the highest is selected (i.e. winner-takes-all selection). If one or more antibodies' concentration values exceed the threshold, an antibody is selected based on the probability proportional to the current concentrations (i.e. roulette-wheel selection).

Figure 7 shows an example network of antibodies. It contains four antibodies, which represent communication behavior and migration behavior with two different policies. Antibody 1 represents the migration behavior invoked when resource availability is low on the local platform. Antibody 1 suppresses Antibody 3 when it is stimulated (i.e. when resource availability is low on the local platform). Now, suppose that a (non-self) antigen indicates (1) resource availability is low on the local platform, (2) network traffic is low on the local platform and (3) user location is close. This antigen stimulates Antibodies 1, 2 and 4 simultaneously. Their population increases, and it is likely that Antibody 2's concentration value becomes highest because Antibody 2 sup-

² The immune system removes non-self detectors through negative selection process. However, in iNet, both self and non-self detectors are kept in a feature table to perform self/non-self classification.

Table I. Initialization Time (T_{init})
with $A=3$, DT could not achieve 90% classification accuracy.
N: recommended size of a feature table

		# of features (A)	3	4	5
80%	Tinit (sec)	t1	0.0015	0.0016	0.0034
		t2	2.21	3.913	9.874
		Total	2.212	3.915	9.877
	Size of feature table (N)		20	40	75
85%	Tinit (sec)	t1	0.0015	0.0031	0.0094
		t2	2.767	7.121	16.101
		Total	2.769	7.124	16.111
	Size of feature table (N)		25	70	150
90%	Tinit (sec)	t1	-	0.0034	0.0141
		t2	-	8.748	18.528
		Total	-	8.751	18.542
	Size of feature table (N)		-	90	225

Table II. $T_{classify}$: Overhead of Classification

# of features (A)	3	4	5	6	..	10
$T_{classify}$ (msec)	1.5	3.0	3.0	3.0	..	4.5

presses Antibody 4, which in turn suppresses Antibody 1. As a result, Antibody 2 (migration behavior) would be selected.

4. Empirical evaluation results

This section shows empirical evaluation results to examine the efficiency and accuracy of iNet and the adaptability of network applications (agents) developed with iNet. The efficiency and accuracy of iNet are evaluated with a Java 2 standard edition JVM on a Windows XP PC with a 2.5GHz Intel Celeron CPU and 1GB memory. The adaptability of agents is evaluated with a maximum of nine Windows XP PCs, each PC hosts a NetSphere platform on a Java2 standard edition JVM with 2.0GHz Intel Celeron CPU and 512MB memory. Those PCs are connected through 100Mbps Ethernet.

4.1. Efficiency and accuracy of iNet

The overhead of EE (T_{EE}) includes the initialization time (T_{init}) and classification time ($T_{classify}$). T_{init} consists of the time (t1) to generate self/non-self detectors, and the time (t2) to build a decision tree (i.e. $T_{EE}=T_{init}+T_{classify}$). Table I shows that T_{init} grows in proportion to the values of N (the size of a feature table), A (the # of features in each feature vector) and classification accuracy. For example, if there are four features in each feature vector and iNet is expected to achieve 90% classification accuracy, EE needs to have 90 detectors in a feature table and its initialization time is approximately 8.6 seconds. Please note that agents (network applications) do not incur the initialization overhead at runtime because EE performs initialization before running agents.

Table II shows $T_{classify}$; how long it takes for EE to classify an antigen (a set of environment conditions) into self or non-self. Agents (network applications) incur classification overhead at runtime. However, the overhead is small enough and acceptable in most agents (applications).

In general, EE can achieve 90% classification accuracy with small overhead depicted in Table II. There is a trade-off between efficiency (T_{init} and $T_{classify}$) and classification accuracy. Application developers need to determine the value of




Figure 8. T_{BS} , Overhead of BS facility




Figure 9. Testbed Architecture

N , based on the results shown in Tables I and II, depending on the requirements of their applications.

The overhead of BS (T_{BS}) represents the time to select an antibody (behavior) suitable for a given antigen (a set of environment conditions). In BS, the number of antibodies (Ab) is determined by the number of features in each feature vector (A), the number of distinct values of each feature (V) and the number of behaviors that each agent supports (B). According to the antibody structure (Figure 5), it will contain a value of one of features and one of behaviors. Therefore, the number of all possible types of antibodies would be calculated as $Ab = (A*V)*B$. As shown in Figure 8, the overhead of BS exponentially increases as Ab grows.

Compared with T_{init} and T_{BS} , $T_{classify}$ is very small. For example, when $A=5$, EE requires $T_{init}=9.877$ sec at least for the initialization (Table I), and BS does $T_{BS}=135.41$ msec for behavior selection with a network of 75 antibodies (e.g. $Ab=(5*3)*5$) (Figure 8). T_{BS} will be skipped if the classification spends $T_{classify}=3$ msec (Table II) and says a current environment condition is "Self". EE works well to keep the overhead of iNet lightweight by eliminating T_{BS} when necessary.

4.2. Autonomous adaptability of agents

In this experiment, the testbed network deploys four NetSphere platforms on four PCs (i.e. a platform on each PC) at the beginning of an experiment (0:00). Three hours later (at 3:00), five more PCs are added to the testbed network, and nine NetSphere platforms work on nine PCs from 3:00 to 7:00 (Figure 10). The platforms and PCs are connected with each other based on a grid (2x2 or 3x3) topology (Figure 9).

At the beginning of an experiment, a single web service agent is randomly deployed on a platform. Each agent contains iNet configured with four behaviors (migration, communication, replication and death) and seven environment conditions (resource availability on the local platform and a one-hop away remote platform, workload on the local platform and a one-hop away platform, energy level, user location, the number of agents running on the local platform). A workload generator simulates a user (Figure 10). It generates HTTP request messages and randomly sends them to agents. It keeps track of the locations of agents. When an agent migrates, the agent notifies its new location to the workload generator³. The workload generator pays energy to an agent when it receives an HTTP response message from the agent.

³ In principle, agents are decentralized as described in Section 2. However,




Figure 10. (1) Workload, (2) # of agents, (3) Throughput

Figure 10 shows how agents autonomously adapt their throughput to workload changes by adjusting their population and locations. Figure 10-(1) shows the workload for agents (how many HTTP request messages a user generates and sends to agents). The workload keeps around 600 messages/min from 0:00 to 1:00, spikes to approximately 12,000 messages/min at 1:20, and drops to 600 messages/min at 5:20. Figure 10-(2) shows how agents change their population against workload changes. As agents perform replication behavior in response to enough energy gain from a user, their population gradually grows. This allows agents to process more HTTP request messages from a user. Also, agents migrate to platforms where resource availability is higher. This prevents agents to excessively crowd on a single platform, and contributes to evenly distribute agents over platforms. As shown in Figure 10-(3), agents autonomously increase their throughput as they perform replication and migration behaviors. The agent throughput represents how many HTTP response messages agents send back to a user.

From 2:20 to 3:00, the agent throughput does not increase although a user generates high workload. The agent population does not increase either. In this period, agents cannot replicate themselves because a number of agents (approx 35 agents) work on four platforms and resource availability is very low on the platforms. When five more platforms (PCs) enter to the testbed network at 3:00, agents migrate to the new platforms, where resource availability is higher, and replicate themselves on the new platforms. After that, the agent population and throughput increase again to process HTTP request messages from a user (see Figure 10).

When the workload drops at 5:00, many agents cannot gain enough energy to keep living because they cannot receive enough HTTP request messages from a user. Due to energy starvation, some of agents die, resulting in a drop of agent population. This prevents unnecessary agents from staying in the network and consuming resources.

5. Related work

Artificial immune systems have been investigated in various applications such as anomaly detection [8] and pattern recognition [9]. This paper proposes an artificial immune system to improve autonomous adaptability of network applications. To the best of our knowledge, this work is the first attempt to apply immunological mechanisms to this domain.

for simplicity, the workload generator in this measurement plays a role of a directory that maintains the locations of agents.

Organic Grid investigates a decentralized task scheduling mechanism for large-scale grid computing environments [10]. Similar to iNet, each agent autonomously executes its service and performs replication behavior to complete the service as fast as possible. However, iNet implements more behaviors for adaptation of network applications.

6. Concluding remarks

This paper describes and empirically evaluates a biologically-inspired mechanism that allows network applications to autonomously adapt to dynamic changes in the network. The proposed adaptation mechanism, called the iNet artificial immune system, allows each application component to autonomously sense its local environment conditions to evaluate whether it adapts well to the conditions, and if it does not, adaptively perform a behavior suitable for the conditions. Empirical evaluation results show iNet works efficiently with high degree of accuracy and makes network applications adaptive.

7. References

- [1] P. Dini, W. Gentsch, M. Potts, A. Clemm, M. Yousif and A. Polze, "Internet, Grid, Self-adaptability and Beyond: Are We Ready?," In *Proc. of the IEEE Int'l Workshop on Self-Adaptable and Autonomic Computing Systems*, Aug. 2004.
- [2] R. Sterritt and D. Bustard, "Towards an Autonomic Computing Environment," In *Proc. of 14th IEEE Int'l Workshop on Database and Expert Systems Applications*, September 2003.
- [3] J. Suzuki and T. Suda, "A Middleware Platform for a Biologically-inspired Network Architecture Supporting Autonomous and Adaptive Applications," *IEEE Journal on Selected Areas in Communications*, vol 23, February 2005.
- [4] N. Minar, K. H. Kramer and P. Maes, "Cooperating Mobile Agents for Dynamic Network Routing," In *Software Agents for Future Communications Systems*, Springer, 1999
- [5] N.K.Jerne, "Idiotypic Networks and Other Preconceived Ideas," In *Immunological Review*, vol. 79, 1984.
- [6] P. Berkhin, "Survey of Clustering Data Mining Techniques," Accrue Software, Inc., 2002.
- [7] J. Suzuki and Y. Yamamoto, "iNet: An Extensible Framework for Simulating Immune Network," In *Proc. of IEEE Int'l Conference on Systems, Man, and Cybernetics*, October 2000.
- [8] F. A. González, D. Dasgupta, "Anomaly Detection Using Real-Valued Negative Selection," *Genetic Programming and Evolvable Machines*, 4(4), 2003.
- [9] L. N. de Castro, J. I. Timmis, "Artificial Immune Systems: A Novel Paradigm to Pattern Recognition," In *Artificial Neural Networks in Pattern Recognition*, University of Paisley.
- [10] A. J. Chakravarti, G. Baumgartner, M. Lauria, "The Organic Grid: Self-organizing Computational Biology on Desktop Grid," In *Parallel Computing for Bioinformatics*, Wiley, 2005.