Objective:

In this project, the PI explores an innovative idea of applying key concepts and mechanisms from the biological world onto network application designs to enable construction of large scale network applications. The Bio-Networking Architecture that the PI proposes is inspired by the observation that the biological world has already developed the mechanisms necessary to achieve the key requirements for the Next Generation Internet (NGI), namely, scalability, adaptability to heterogeneous and dynamic conditions, survivability, and simplicity.

In the biological world, each individual entity (e.g. a bee in a bee colony) follows a simple set of behavior rules (e.g. migration, reproduction, energy exchange, mutation, and death), yet a group of entities (e.g. a bee colony) exhibits complex, emergent behavior (e.g. scalability, adaptation, and survivability). Therefore, if services and applications adopt biological concepts and mechanisms, they too may be able to achieve the key requirements of NGI.

In the Bio-Networking Architecture, a network application is implemented as a decentralized collection of autonomous objects called *cyber-entities*. This is analogous to a bee colony (a network application) consisting of multiple bees (cyber-entities). Each cyber-entity implements a functional service related to an application and follows simple behavior rules (e.g. replication, reproduction and migration) similar to biological entities. Cyber-entities store and expend *energy*. They gain energy in exchange for providing their services, and expend energy for receiving other cyber-entities' services, utilizing resources (e.g. CPU cycles and memory space), and performing their biological behaviors. In the Bio-Networking Architecture, useful applications emerge from the autonomous interaction of cyber-entities, and useful system behaviors and characteristics (e.g. scalability, adaptation, and survivability) arise from the simple behaviors and interaction of individual cyber-entities. .

Approach:

The PI applies key concepts and mechanisms from the biological world and designs the proposed Bio-Networking Architecture. The Bio-Networking Architecture is the first attempt to apply the biological concepts of emergent behavior, autonomous control, and adaptation and evolution to a broad and general class of network applications.

The PI's approaches to the Bio-Networking Architecture are to investigate the feasibility through simulations and to empirically evaluate the Bio-Networking Architecture through prototype design and implementation. The progress in both approaches is summarized below.

The PI has developed a simulator for the Bio-Networking Architecture in Java. The simulator can simulate cyber-entities with different behavior policies under a wide variety of network topologies and user demand workloads. It can also simulate energy exchange mechanism and evolution of cyber-entities. **The PI has recently added to the**

simulator functionality necessary to simulate peer-to-peer discovery mechanisms in order for cyber-entities to locate other cyber-entities in the Bio-Networking Architecture. Various adaptation and evolution scenarios, as well as discovery algorithm developed in the project, are currently being simulated to evaluate the feasibility of the Bio-Networking Architecture.

In order to empirically evaluate the Bio-Networking Architecture through prototype design and implementation, the PI has designed the Bio-Networking Architecture. Design for the Bio-Networking Architecture is described in the technical report submitted in December 1999. In addition, **the PI has also designed and is currently implementing the Bio-Networking platform, a software that runs on a computer to support cyber-entities, in order to empirically evaluate the Bio-Networking Architecture.**

The detailed information (including design documents and source code) on the simulator, Bio-Networking Architecture designs, and the Bio-Networking platform (including design documents and source code) is available at netresearch.ics.uci.edu/bionet/resources/.

Recent Accomplishments:

With the funding from DARPA, the PI has produced significant research results. **The PI has published a number of papers: two journal papers, nine conference papers, eighteen workshop papers, two technical papers, two standardization documents, and four papers in submission. All the papers are available at netresearch.ics.uci.edu/bionet/publications/. One of the papers [netresearch.ics.uci.edu/bionet/publications/saint2002.pdf] received the best paper award at the 2002 IEEE SAINT (Symposium on Applications and the Internet) conference. This is a major IEEE conference specializing in Internet applications with ten technical paper sessions, three panel sessions and four workshops.**

The PI has been working with the super distributed object (SDO) working group of the Object Management Group (OMG), the largest standard making body for object oriented software technologies. **The key concepts and mechanisms of the Bio-Networking Architecture are reflected in the reference architecture developed by the OMG SDO working group [netresearch.ics.uci.edu/bionet/publications/SDO-WhitepaperV100.doc]. The working group also published a request for specification proposals under the initiative of the PI.**

The accomplishments reported in the past reports are summarized first.

- Feasibility study of the Bio-Networking Architecture through simulations of a small-scale web content distribution application. The results are described in the technical reports submitted in December 1999 and July 2000, and also published in a paper [netresearch.ics.uci.edu/bionet/publications/mwang-saint2001.doc].

- Design of a secure mechanism for exchanging energy: This work is described in the technical report submitted in December 2000, and published in a paper [netresearch.ics.uci.edu/bionet/publications/icact2001.doc].
- Stability analysis of the Bio-Networking Architecture. The results are described in the technical report submitted in July 2001, and published in a paper [netresearch.ics.uci.edu/bionet/publications/miyamoto.pdf].
- Preliminary design of the Bio-Networking platform. This work is described in the technical report submitted in July 2001, and presented in a workshop [netresearch.ics.uci.edu/bionet/publications/suzuki_jwaits01.ppt].
- Design of an e-commerce type application (automated ticket sales application) using the Bio-Networking Architecture. This work is described in the technical paper submitted in July 2001, and published in [netresearch.ics.uci.edu/bionet/publications/dicomo2001.pdf].

In the following sections, the new research results obtained to date since the last technical report submitted to DARPA in July 2001, as well as various PI's activities to gain research community's support, are summarized.

Recent Accomplishment: Distributed Discovery

As with objects in distributed peer-to-peer applications, cyber-entities in the Bio-Networking Architecture often need to discovery and locate other cyber-entities. In the discovery mechanism that the PI developed, cyber-entities contain keywords and also contain a limited set of relationships (links that include information about other cyber-entities) to other cyber-entities. These relationships between cyber-entities form a network on which discovery queries are forwarded.

In this first discovery mechanism that the PI developed, discovery queries are forwarded based on keyword similarity and discovery history of relationships between cyber-entities.

Keyword similarity is a measure to evaluate the degree of similarity between two cyber-entities. Keyword similarity is defined as the ratio of keywords that are in common between a cyber-entity and its relationship partner. Relationship history summarizes information on how a relationship partner performed in discoveries in the past. History of a relationship is defined as the relative ratio of successful discovery queries against all the discovery queries forwarded on the relationship. A cyber-entity that has performed well at satisfying past is more likely to perform well in future discoveries.

Keyword similarity and relationship history are used at each cyber-entity to determine which relationships have priority in forwarding discovery queries. A discovery query is forwarded with greater priority to cyber-entities that are most similar to a discovery query and for relationships that are equally similar, history is used as a secondary priority.

In order to evaluate the discovery mechanism, simulations have been performed, comparing: cyber-entities with random relationships, cyber-entity relationships based only on keyword similarity (no discovery history), and cyber-entity relationships based on both similarity and history. Simulations have shown that cyber-entities that have relationships including both keyword similarity and discovery history improve discovery performance.

The following figure shows how adapting the relationships from a random state to a state with a stronger keyword-similarity affects average discovery search time. Initially, relationships are random (Figure 1 (b)), and the average discovery search time is large (Figure 1 (a)); however, as average relationship similarity increases, the average time per search decreases. [AINS Symposium: netresearch.ics.uci.edu/bionet/publications/discovery_ains.pdf, DARPA PI Meeting: [netresearch.ics.uci.edu/bionet/publications/darpa_discovery_michael]

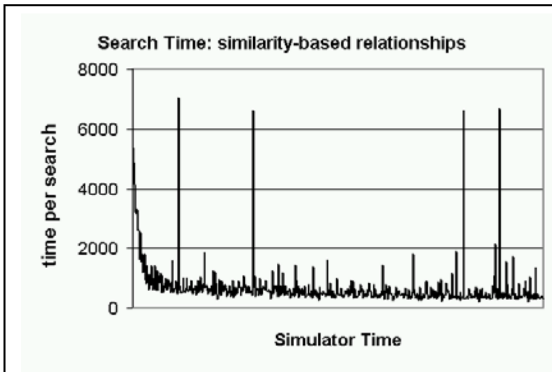

Figure 1 (a)

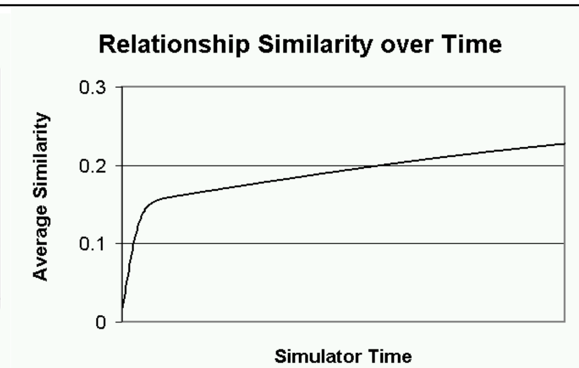

Figure 1 (b)

In addition to the discovery mechanism based on keyword similarity between cyber-entities and relationship history of cyber-entities, the PI has developed another discovery mechanism that has the capability to reflect user's (i.e., discovery originator's) evaluation of the received discovery hits. (For instance, a user may prefer router information on a reputable web site to router information on a unfamiliar web site.) This second proposed discovery mechanism allows a user to express the degree of satisfaction by sending a reward message after receiving discovery hits. The degree of satisfaction, or reward, is used in later discoveries to obtain better results (hits) and to achieve better performance and resource usage. [netresearch.ics.uci.edu/bionet/publications/enomoto_master_thesis.pdf] [netresearch.ics.uci.edu/bionet/publications/darpa_discovery_enomoto].

In this second proposed discovery mechanism, each relationship is associated with one or more keywords, and for each keyword, a strength value is also associated. Keyword strength represents the usefulness of the relationship in discovering a cyber-entity that contains the given keyword. In this discovery mechanism, both the new keywords added to relationships and the strength values of the keywords are adjusted based on user evaluation in the rewarding phase in the following manner.

Upon receiving a discovery hit, a user (i.e., discovery originator) sends a reward message to indicate the degree of his/her satisfaction with the received discovery hit (and the keywords for which the discovery hit was returned). A reward message is propagated to the cyber-entity that returned the discovery hit through the path that the discovery hit has traversed before. Over time, the keyword strength of a relationship represents how useful the relationship was at discovering a cyber-entity that contained the given keyword and also satisfied many users.

Recent Accomplishment: Implementation and Measurements of the Bio-Networking Platform

The Bio-Networking platform is a middleware environment for deploying and executing cyber-entities. It runs on a network node and provides execution environment for cyber-entities. It consists of three components; bionet services, bionet message transport, and bionet container. Bionet services provide runtime services that cyber-entities frequently use, and major bionet services are listed below.

- Bionet lifecycle service
- Bionet migration service
- Bionet relationship service
- Bionet discovery service
- Bionet cyber-entity sensing service
- Bionet topology sensing service
- Bionet resource sensing service
- Bionet pheromone emission service
- Bionet energy management service

Bionet message transport and bionet container abstract low-level operating and networking details such as I/O, concurrency, messaging, and network connection management.

The design of the Bio-Networking platform was presented at the DARPA PI meeting at January 2002 [netresearch.ics.uci.edu/bionet/publications/darpa_platform.ppt]. It is also described in a paper that is currently under submission [netresearch.ics.uci.edu/bionet/publications/doa02.pdf]. Recent accomplishment includes (1) implementaation of bionet message transport, bionet container and four bionet services (lifecycle, relationship management, energy and resource sensing), and (2) preliminary measurements to examine efficiency and scalability of the Bio-Networking platform [netresearch.ics.uci.edu/bionet/publications/doa02.pdf].

Figure 2 shows a result of preliminary measurements of the time required for cyber-entities to exchange messages using the Bio-Networking platform implemented by the PI. In the measurement, the Bio-Networking platform runs on a Java 2 virtual machine (version 1.4) atop of a Windows XP with an Intel Pentium 3 processor (1 GHz) and 256 MB RAM. As shown in Figure 2, the latency for message transmission in the Bio-Networking platform is comparable with existing distributed object platforms (Jac ORB, Java IDL, Zen). Figure 2 also shows that the message exchange latency using the Bio-Networking platform remain relatively

constant as the number of cyber-entities grows (0.18 msec on the average).

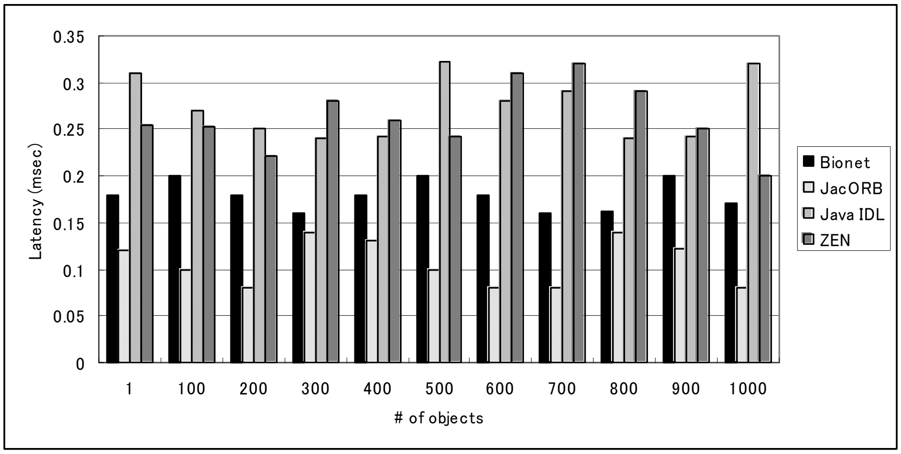

Figure 2. Latency for message transmission between two cyber-entities

Recent Accomplishment: Service Composition

In the Bio-Networking Architecture, complex application emerges through an interaction of multiple cyber-entities. Since cyber-entities may be designed and developed independently by different designers, it is important to ensure that such independently developed cyber-entities can communicate and interact to collectively provide a network application. The PI developed a set of specifications (called Loose Interface Definition) that defines cyber-entity interface in a loose manner so that cyber-entities with different interfaces can communicate and interact.

In the proposed Loose Interface Definition, cyber-entity interface is defined as a set of elements, namely, a name, a type and a value. Unlike existing interface definition languages, Loose Interface Definition allows each element to have multiple names so that it can accept synonyms. Loose Interface Definition also supports optional elements so that two cyber-entities having similar, but not exactly the same, interfaces can communicate. Loose Interface Definition also disregards the order of the elements in the interface so that two cyber-entities that have interface with identical set of elements can communicate and interact.

PI's recent accomplishments in this work include design and preliminary implementation of Loose Interface Definition. The PI has also shown that Loose Interface Definition increases the number of web services that can exchange their information. The result was presented at the DARPA PI meeting in January 2002 at the UCLA AINS workshop in May 2002 [http://netresearch.ics.uci.edu/bionet/publications/darpa_service.ppt, http://netresearch.ics.uci.edu/bionet/publications/loose_interface.pdf]

Recent Accomplishment: Adaptation and Evolution

Adaptation and evolution by natural selection is one of key biological concepts that the PI applies in the Bio-Networking Architecture. Analogous to a biological entity in the real world biological systems, cyber-entities in the Bio-Networking Architecture have simple behavior rules (e.g., replication, reproduction and migration) and invoke their behavior based on information locally available to them.

A cyber-entity evolves in its behaviors. When a cyber-entity replicates or reproduces with another cyber-entity, diverse behavior policies are created through mutation and crossover in their behavior policies. Since an inefficient cyber-entity exhausts energy quickly and eventually dies, only beneficial cyber-entities are retained in the system, enabling network applications to adapt to changing environments.

The PI is currently investigating adaptation and evolution aspects of the Bio-Networking Architecture using the simulator that the PI developed. In the preliminary simulations, a network is configured as a 6×6 grid topology, and 10 users with different request rates ranging from 2 to 12 requests per second are randomly distributed over the nodes in the network. A single cyber-entity is initially placed in the network, and it processes 5 requests per second. It receives 10 energy units from a user in return for providing a service, while paying energy to a platform to use computing resources (1 energy unit per second) and platform services (100 energy units per behavior). In reproduction a partner cyber-entity is selected from those on the adjacent nodes, based on a fitness value of a possible mating partner cyber-entities. A fitness value of a cyber-entity is calculated as (acquired energy)/(consumed energy). The cyber-entity who has the highest fitness value is selected as a mating partner, and a new offspring with diverse behavioral policies is created through mutation and crossover.

Figure 3 shows preliminary simulation results. In this figure, the vertical axis indicates the simulation time (in seconds), and the horizontal axis indicates the energy gain (= acquired energy – consumed energy) per simulation cycle. This figure shows that, as time progresses, the cyber-entities with an evolutionary mechanism gradually reach higher energy gain than those without an evolutionary mechanism. This preliminary simulation result reveals that cyber-entities in the Bio-Networking Architecture evolves to optimize their behaviors to minimize the energy consumption while providing services with users.

Accomplishments in this work include the simulator (available at [netresearch.ics.uci.edu/bionet/resources]) and theoretical understanding of adaptation and evolution aspects of the Bio-Networking Architecture. The simulation results are published in conference proceedings [netresearch.ics.uci.edu/bionet/publications/mwang-saint2001.doc], and presented at the IEICEJ Next Generation Networks workshop, December 2001, and the DARPA PI meeting, January 2002 [netresearch.ics.uci.edu/bionet/publications/darpa_evolution_simulation.ppt].

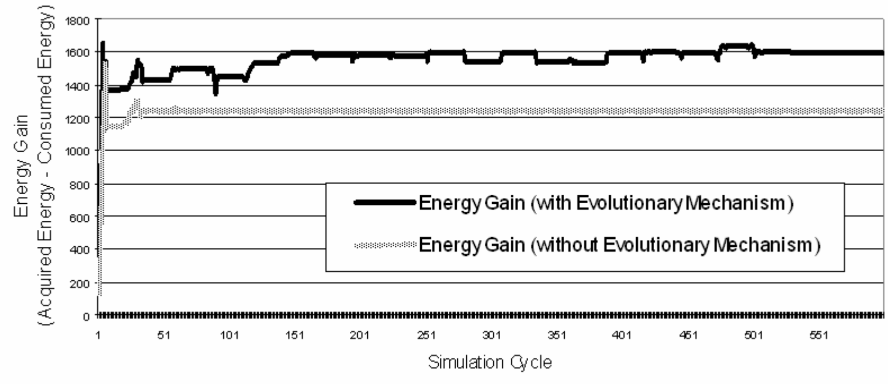

Figure 3. Energy Gain with Evolution

Other Accomplishments

In addition to the research efforts described in the above sections, the PI has investigated various other aspects of the Bio-Networking Architecture including behavior selection of cyber-entities, development of example applications, resource access control in PDAs, and authentication of cyber-entities.

Recent Accomplishment: Increasing Community Awareness

The PI has taken steps to gain support from the research community for the Bio-Networking Architecture. Since the proposed Bio-Networking Architecture is innovative and new, it is important that the research community recognizes the advantages of the proposed architecture. In order to gain support from the research community on the new Bio-Networking Architecture, the PI has taken the following steps.

- The PI has contacted one of the NTT research laboratories. It has allocated a group of researchers (5 researchers) to investigate the Bio-Networking Architecture with the PI. The PI is contacting other NTT laboratories and a number of university professors (including those at Tokyo Univ. and Osaka Univ.) to join a team to investigate the Bio-Networking Architecture.
- The PI has been working with the super distributed object (SDO) working group of the Object Management Group (OMG), the largest standard making body for object oriented software technologies. **The key concepts and mechanisms of the Bio-Networking Architecture were reflected in the reference architecture developed by the OMG SDO working group [netresearch.ics.uci.edu/bionet/publications/SDO-WhitepaperV100.doc]. The working group also published a request for specification proposals under the initiative of the PI. The PI will continue working with OMG to reflect the key concepts and mechanisms of the Bio-Networking Architecture in OMG standard specifications.**
- The PI organized a workshop on scalable and evolvable distributed systems at the IEEE SAINT conference held in January 2001. The Bio-Networking Architecture was one of the topics discussed at this workshop.

- The PI has given a keynote speech at the IEICEJ (Japanese equivalent to IEEE) meeting in December 2000, regarding the future direction of networking research, and one of the emphasis of the talk was the Bio-Networking Architecture.
- The PI gave a presentation and participated in a panel in conferences and professional meetings (e.g., the IEICEJ (Japanese equivalent to IEEE) meetings in July 2001 and in September 2001). One of the emphasis of the talks was the Bio-Networking Architecture.