Swarm Intelligence and the Holonic Paradigm: A Promising Symbiosis for a Medical Diagnostic System

Rainer Unland* and Mihaela Ulieru**

*University of Duisburg-Essen
Institute for Computer Science and Business Information Systems (ICB)
45117 Essen, Germany
unlandR@cs.uni-essen.de; http://www.cs.uni-essen.de/dawis/

** Canada Research Chair
The University of New Brunswick,
Faculty of Computer Science
Fredericton, Canada
http://eis.enel.ucalgary.ca

Abstract. A self-organizing medical diagnosis system, mirroring swarm intelligence to structure knowledge in holonic patterns is presented. The system sets up on an alliance of medical experts - realized by agents - that stigmergically self-organize in order to provide a viable medical diagnosis. Starting point is always a flat set of autonomous agents that spontaneously and temporarily organize on a case-based basis into a holarchy (without hierarchical control flow). Despite their sophisticated task the proposed agents, like ants in an ant colony, exhibit a comparatively simple architecture built on reactive behavior. The real power of the system stems from the fact that a large number of those simple agents collaborate in order to come to a reliable diagnosis.

Introduction

Medical diagnostic is fundamentally the process of identifying the disease a patient is suffering from in order to determine how it can be treated best. Although research on medical diagnostic has been a hot topic in computer science for quite a while, especially in the field of artificial intelligence, computer-based medical diagnostic has not yet reached a breakthrough. Reasons for that are manifold. Most proposals published so far are monolithic systems that rely on a sophisticated and ever-growing set of rules whose interplay and “firing“ cannot be controlled easily. For diagnosis in general and medical diagnostic in particular a seamless, fast and easy integration of new and an abortion of outdated knowledge is of the essence. Multi-agent systems are a first major step in this direction since they provide decentralized control and, thus, can be adapted and extended much easier. A next step is to rely on a huge number of relatively unsophisticated (reactive) agents instead of a smaller number of sophisticated (deliberative) ones. This avoids the modeling of complex domain knowledge, thus, avoids sophisticated architectures that can hardly be controlled. To implement such a system we propose to combine swarm intelligence techniques with the holonic paradigm (see [1], [2] for the first papers proposing the idea of a holonic medical diagnostic system). Swarming systems enable a system to satisfy requirements that would be surprising to its original designer as stated in [3]. The holonic paradigm adds the nested hierarchical structure in which the swarms may self-organize [4]. Altogether, such a medical diagnostic system will exhibit all the key properties of these paradigms, most of all self-organization , emergent behavior [5], simplicity, fast and unbounded learning, flexibility, robustness, adaptivity, self-sustainability, and easy extensibility. Nevertheless, the system will still exhibit a relatively simple architecture that can be maintained and controlled comparatively easily. This is because the sophisticated behavior of such a system stems from the fact that a huge number of simple components on the so-called micro-level exhibit a sophisticated behavior at their interface, the so-called macro-level, due to their (implicit) collaboration and interplay. Emergence, self-organization, and the holonic paradigm have some similarities [6]. In fact, holonic systems can be seen as one way to implement emergent systems. Basic building blocks of holonic systems are holons (cf. [7]). On one level of abstraction a holon can be seen as a unit with a unique interface hat provides some services. On the next lower level of abstraction a holon (recursively) consists of a set of lower-level holons that cooperate in order to solve the problem at hand. The unique and sole interface to a holon is realized by a so-called mediator or head agent. All the other agents of a holon are called body agents (cf. [8]). Holons form a nested structure, called holarchy, where each holon is defined on the next lower level of abstraction by a set of body holons. This recursively continues until, finally, the leaf level, consisting of atomic holons, is reached. Holonic multi-agent systems consist of autonomous agents with holonic properties that recursively group together to form holons and, by this, holarchies (cf. [6], [9]).

Functionality of the Diagnostic System

When designing a complex system like the proposed holonic medical diagnostic system (HMDS) a number of questions concerning the architecture of the system have to be answered. After a thorough investigation, where several other alternatives [10] were explored, we concluded that an integrated approach with relevant concepts from the holonic paradigm, multi-agent systems technology and swarm intelligence is the most promising approach. What results is an open, highly adaptable architecture. It consists of many comparatively simple components that collaborate in order to solve complex problems. Basis is the autonomous or flat agents' approach of holonic systems [11] where all agents are initially on the same level. Communication between agents is solely done via blackboards. Each non-leaf agent is assumed to possess such a blackboard on which messages can be published that are of interest for its (potential) body agents. Due to space limitations we can not discuss the system in detail. The interested reader is referred to [4].

The proposed system mirrors stigmergy (for our parallel we will use the ant colony as an example for stigmergy) as follows: When a request for a medical diagnosis is sent to the HMDS it is dealt with by its single entry point, the mediator agent. Based on the classification of the disease (e. g., eye disease, skin disease, inner organ disease, etc.) the request is then assigned to those holons that are assumed to be specialists in the field on this level of expertise. This procedure is recursively repeated until the request reaches the final level, the leaf level of the holarchy. Results that are obtained by individual holons now flow the other way round, from the bottom to the top of the holarchy. On their way up the results are examined, assessed, put in context and, finally, sorted according to their estimated likelihood and quality.

Individual, case-based holarchy formation

In HMDS the nested holarchy will emerge dynamically and individually on a case-based basis. For each diagnosis request a new, individual holarchy evolves. Thus, holarchies are formed in an ad-hoc way and exist only temporarily. As soon as a diagnosis request is sent to the HMDS the holarchy begins to evolve out of the flat set of agents. It emerges in a combination of top-down and bottom-up approach. The general direction is top-down while each holon is formed bottom-up. More specifically, first one or more agents agree to the job of working on the diagnosis request. In order to do so they must have some confidence that they will be able to provide a reasonable diagnosis. Each of them becomes the head of a new holon. In order to find appropriate body agents that are willing to help in the diagnosis finding process each head agent posts its request for support on its blackboard. Other agents can read this request and decide to cooperate. By that they become body agents of the holon - a new holon evolves. From now on body agents of the same holon are called buddy agents. Since communication is solely done via blackboards, requests as well as replies are posted on the appropriate blackboard. Thus, starting with the mediator agent of HMDS the holarchy grows from the top, level by level and holon by holon, till it finally reaches the leaves. A holarchy evolves where each holon can be seen as a swarm or colony of ants, thus, forming a nested ant colony. An ant on one level of abstraction is realized by an ant colony/swarm of ants on the next lower level of abstraction. Note that this temporary holarchy only comprises holons that are actively involved in the diagnosis finding process. Dead branches will not evolve.

Pheromone trails

The main purpose of ants in an ant colony is to gather food to be able to survive. The main goal of an ant in HMDS is to solve a diagnosis request. On this basis there is obvious conformity between the knowledge about the existence of a blackboard and the knowledge of an ant about the existence of a food source due to the pheromone trail. For this reason, this concept was adopted. Whenever there is a relationship between two agents in the sense that one agent reacts to requests posted on the blackboard of a head agent or, more general, uses the blackboard of the head agent, a (the) pheromone trail from the body agent to (the blackboard of) the head agent is established, respectively strengthened. What emerges is a system whose structure is defined by blackboards and pheromone trails. The downward directed relationships from agents with broader knowledge to agents with deeper knowledge is emulated by the blackboard, thus, not encoded anywhere. The reverse relationship, the upward one, is expressed by the pheromone trail between agents. This implied set of acquaintances is not only very simple but also self-maintaining.

Emergence and self-organization by random search and feedback messages

In ant colonies pheromone trails are the means for communication. In HMDS the blackboard plays the role of the pheromone trail and the information on it emulates the scent of the trail.

Scent strength

In ant colonies ants will more likely follow a trail with strong scent. Since the scent evaporates or gets stronger in dependence of the usage of the trail by ants, scent is one of the means for self-organization. In HMDS scent is emulated by the blackboard and the information on it. Each blackboard is partitioned in specific areas that are solely dedicated to those agents that have a pheromone trail to this blackboard, one partition for each agent. Among other information, each partition provides information about the strength of its pheromone trail. Partitions are sorted according to the strength of their pheromone trail. Thus, an agent that is on a quest and accidentally passes by (see the discussion about random search below) can immediately identify the more successful agents with respect to this blackboard and can compare itself to it. In addition to this body agent-specific scent value the blackboard maintains a holon-specific scent value that discloses information about the overall success of the head agent. It is essentially the sum of the individual scent values of all body agents.

Scent type

In ant colonies pheromone trails are open for all agents; i. e., ants will follow it because of its strength but not because of the type of food they can expect at its end. In HMDS the latter plays a role since the ability to provide a successful diagnosis depends on whether the input matches the field of expertise of an agent. If not it is worthless; i. e., the trail is not worth to follow. In comparison to an ant colony this means that each ant can only carry a specific type of food. Thus, only trails to this kind of food source are of interest for them. HMDS emulates this in the following way: for every pheromone trail of an agent information about it is provided in their partition on each blackboard to which they have a pheromone trail. This information comprises the scent value of the trail, recent results from using this trail, the DDPs that were used in solution processes related to this trail, etc. If an agent wants to try new trails/holons it can check this information for compatibility to its own data. If there is a great overlap this indicates that the pheromone trails of this agent are probably of the right type. Thus, it may be worth to test them.

Explorer agents

Ants in an ant colony either are involved in food carrying or are on a quest to find new food sources. Similarly, agents in HMDS either are engaged in the solution process of a request or are on a quest through the information space of the whole system (random search) in order to find better holons. Agents on a quest are called explorer agents. Such agents are very important for an uncontrolled and unpredictable, however, positive development. Due to the randomization accidental hits are possible which may lead to developments that could not be anticipated beforehand. An agent that wants to explore new possibilities has several ways to do so. The easiest way is to check the blackboards of all holons it is member of for success stories about involvement from other body agents in other holons, called buddy check in the following. As discussed above this is not exploring in its original sense but following an already existing, however, from the point of view of the agent, new pheromone trail that was laid by a buddy agent. Seen from the likelihood of success such a that was laid by a buddy agent. Seen from the likelihood of success such a proceeding seems to be the most promising since it can be assumed that there is a lot of overlap between buddy agents. An explorer agent explores its environment by walking around randomly.

While explorer agents will walk randomly, it seems to be more logical that they will find new promising holons especially in their direct neighborhood. By direct neighborhood those holons are meant that share the same higher level (parent) holon than the original holon of the explorer agent. This guarantees that the new holon is at least a representative of a (class of) disease(s) from the same super-class of diseases than the original holon. However, in order to make accidental positive development happen a random walk is a prerequisite.

Pheromone trail selection

As long as agents are in a prospering symbiosis with one or several holons they probably will decide not to worry about other success stories on blackboards they have access to. However, especially if all pheromone trail of an agent are on a comparatively low level then the likelihood that the agent may try new possibilities increases substantially. As discussed above there are two alternatives for a change. One is to read success stories on accessible blackboards and the other is to start a quest. In both cases a decision has to be made whether to join a visited (quest) or proposed holon (buddy check). In both cases the information about agents has to be analyzed. Decisive factors can be, on the one hand, the scent strength and, on the other hand, the overlap between the set of symptoms that was used by the "advertising" agent and the set of symptoms that is used by the searching agent.

Reinforcement learning as basis for self-organization

Knowledge about pheromone trails can be seen as the life elixir that allows the agents to stay alive. Thus, it has to be their main goal to strengthen their pheromone trails as much as possible, respectively to add new, successful relationships to their list of acquaintances/holons.

Strengthening of pheromone trails by providing sound solutions

An agent that provided something positive to the solution process of a request (a diagnosis that turned out to be valid) will strengthen its pheromone trail to the according blackboard.

Evaporation of pheromone trails by negligence

Like with real ant colonies, a pheromone trail will evaporate over time if it is not replenished. Thus, if an agent does not want to work for a holon any longer (due to more attractive possibilities with other holons) the trail to it will slowly evaporate and finally diminish. Relationships that are cultivated will become stronger and stronger, especially, if the contribution of the delivering agent is of high quality and the competition is not too high; i. e., not too many other body agents deliver the same diagnosis to a request.

General evaporation of pheromone trails

As with real ant colonies as well a pheromone trail that is not used enough any longer will weaken its scent and finally diminish. In order to prevent successful agents from losing their knowledge about blackboards just because their advice has not been

asked for for a while (because, e. g., they represent a very seldom disease) the evaporation of the pheromone trail depends on the frequency in which requests are posted on a blackboard to which a pheromone trail exists and not on the time. Whenever a request is posted on the blackboard of a head agent of a holon the pheromone trail of all body agents will be evaporate a little bit. If an agent decides to contribute to the solution process, the delivery of a response will once again weaken the trail. This ensures that a trail will evaporate if a body agent does not use it any more. It also ensures that it will evaporate if the body agent delivers a poor response. By how much the trail is weakened overall depends on the influence of the response on the diagnosis process. Exceptionally poor performances of agents lead to a faster evaporation of their pheromone trails since they have to pay the entrance fee, however, will not get a reward. In general, each agent that reacts to a diagnosis finding request will first weaken its pheromone trail. Thus, only when its contribution was strong enough the difference between this “entrance fee” and the reward will be positive, thus, strengthening its pheromone trail. This proceeding will filter out all those agents that only submit average or below average diagnoses in comparison to other agents.
When a new pheromone trail is created, i. e., an agent has joined a new holon, this trail needs some “seed capital” to allow it to participate in solution processes. This capital/scent can be withdrawn from existing pheromone trails, e. g., from those ones that will most likely not be used any longer. They will evaporate faster or even diminish immediately, however, finance the exploration of new trails.

Emergence through birth and death of agents
Less successful agents will lose knowledge about their pheromone trails. They may have to search for better holons in order to survive. If they do not find at least one proper holon, they will finally die. The other way round, agents that are very successful, what implies that there is a high demand of their services, can clone themselves. The “value” of an agent is decided by the number of pheromone trail it has knowledge of and the strength of each of these trails. If a certain threshold is exceeded an agent clones itself. This supports the natural selection process in the system. Successful agents will survive and proliferate while luckless ones finally will die (principle of survival of the fittest). Since new agents first need to find their place in the system, they will be equipped with an asset high enough to allow them to survive for some time even if they are not successful. They will not only inherit the knowledge of their generator about blackboards but also the strength of the scent that is related to each of these pheromone trails. However, they also exhibit a strong impulse to explore new possibilities. More specifically, the threshold for abandoning an existing pheromone trail and searching for new ones is much higher than that of their generators. This guarantees that successful agents do not flood their environment with, in principal, identical buddies that do not contribute to the quality of the solution process at all.

Conclusions

In today’s ever-changing world that grows together ever faster, a quick, ubiquitous, and reliable diagnosis generation is of eminent importance as can, for example, be seen from the recent problems with SARS. In addition, the cost explosion problem in the health care sector asks for a highly efficient, cost-effective and reliable health care system. In this paper a holonic medical diagnostic system was presented that unifies the advantages of multi-agent system technology with the self-organizing, emergent behavior, the simplicity, the capability of fast and unbounded learning, the flexibility, the robustness, the adaptivity, and the reliability of swarm intelligence and holonic systems. The resulting system supports the medical institution in its diagnosis generation, especially in cases in which the kind of disease the patient is suffering from can only be assumed.

REFERENCES

[1 [1] Ulieru, M. and Geras, A.: Emergent Holarchies for e-Health Applications: A Case in Glaucoma Diagnosis; Proc. IECON '02; HOL - Agent-based Intelligent Automation and Holonic Control Systems; Sevilla, Nov. 2002

[2] Ulieru, M.: Internet-Enabled Soft Computing Holarchies for e-Health Applications, Invited Chapter in: Fuzzy Logic and the Internet; Zadeh, L.A.; M. Nikravesh (Eds.), Physica Verlag; 2003

[3] Van Dyke Parunak, H.; and Brueckner, Sven A.: Engineering Swarming Systems; in F. Bergenti, M.-P. Gleizes, F. Zambonelli (eds.): Methodologies and Software Engineering for Agent Systems; Kluwer Publishing Company; 2004

[4] Mihaela Ulieru, Maja Hadzic and Elizabeth Chang, "Soft Computing Agents for e-Health Applied to the Research and Control of Unknown Diseases", Information Science (in print – to appear 2005).

[5] De Wolf, T., Holvoet, T.: *Emergence and Self-Organisation: A Statement of Similarities and Differences*; Workshop on Engineering Self-Organizing Applications (ESOA); July 21-24, New York, USA; 2004

[6] Ulieru, M: Emergence of Holonic Organizations from Multi-Agent Systems: A Fuzzy-Evolutionary Approach, Invited Chapter in: Loia, V. (ed.): Soft Computing Agents, IOS Press; 2002

[7] Koestler, A: The Ghost In The Machine, Arkana Press; 1967

[8] Schillo, M: Self-Organization and Adjustable Autonomy: Two Sides of the Same Coin?; In Connection Science, 14 (4), pp. 345-359, 2003

[9] Ulieru, Mihaela: The Holonic Enterprise: Modeling Holarchies as MAS to Enable Global Collaboration. Third Int. Workshop on Industrial Applications of Holonic and Multi-agent Systems – HoloMAS, 2-6 Sept. 2002, Aix-en-Provence, France. IEEE Computer Society, ISBN 0-7695-1668-8, pp 603-607; 2002

[10] Ulieru, M.; Stoica, C.;, Klueüer; J.; Unland, R.: A Holonic Medical Diagnosis System Based on MAS Technology and Neural Networks; Journal Special Issue on Computational Intelligence for Bioinformatics; Special Issue Eds: N. Kasabov, M. Nikravesh and L. A. Zadeh; to appear 2005

[11] Fischer, K.; Funk, P.; Ruß, C.: Specialized Agent Applications; in [36]; pp.365-382; 2002.