Konrad Kulesza (1)
Zbigniew Kotulski (2)
(1) Rhodes University, Grahamstown South Africa, e-mail: g98k6080@campus.ru.ac.za
(2) Institute of Fundamental Technological Research, Polish Academy of Sciences ul.Świętokrzyska 21, 00-049, Warsaw Poland, e-mail: zkotulski@ippt.gov.pl

Decision Systems in Distributed Environments: Mobile Agents and Their Role in Modern E-Commerce

Streszczenie.

Praca wskazuje Mobilnych Agentów jako znakomite narzędzie w nowoczesnych systemach decyzyjnych. Wszelkie zmiany systemów nie są wynikiem przypadku, lecz rezultatem zmian otoczenia, to zaś może podlegać zmianom przypadkowym. W tej pracy autorzy przedstawią agentów jako najbardziej odpowiednie przy obecnym stanie wiedzy rozwiązanie do podejmowania decyzji we współczesnym, zmiennym, a więc wymagającym otoczeniu. Podane są przykłady zagadnień możliwych do rozwiązania metodą mobilnych agentów i alternatywne sposoby ich rozwiązania. Podstawowym przykładem ilustracyjnym w pracy jest globalny rynek giełdowy.

Abstract
This paper focuses on Mobile Agents as the fine vehicle of modern decision systems. Changes don't just happen but are the result of changes in the environment, which itself is changing as a result of the changes mentioned above. In this paper the authors will look at the agents as a state of the art solution to decision-making problems of today's highly demanding environment. Potential problems will be illustrated and competing technologies discussed. The stock market will be used as a case study and referred to on numerous occasions.
Słowa kluczowe: Moblini agenci, systemy decyzyjne w układach rozproszonych, rynki elektroniczne.
Keywords: Mobile Agents, decision systems in distributed environment, e-commerce.

Introduction
Before discussing mobile agents we need to define and understand the process of decision making; then take a brief look at the evolution and history of decision-making systems, which is going to lead us to a discussion of mobile agents. The authors believe this approach to be the best.

Any author of an article faces a challenge between his article being either too specific or too general. In other words saying too much about nothing and being far too narrow to lead the reader to any long-term conclusion, or saying nothing about everything and simply not addressing anything at all. Keeping this challenge in mind the authors will try to establish an optimum balance. The authors think it is fundamentally necessary to lay down the structure in this way:
Future Applications -- Agents --- Evolution --- Concept ---
Only then we can get the most out of this paper.

Chapter 1 --- Decision Making as an Activity, Art and Science
"We are what we decide" -Konrad Kulesza
Once upon a time I was puzzled by a problem of a religious and philosophical nature. How, if God can foresee everything, can we still have free will? If he can see things before

they happen, then we haven't got freedom to decide. I discussed the problem with an elderly woman and she gave me a very interesting answer. Imagine you are in a train, you only see things passing you by. But if you stand on top of the train you can see much more, including what's coming up ahead. Now imagine someone high above the train, he can see what's ahead of the train as well as the pathways it can take. It really made me think…

Decision-making is an issue that is fundamental and essential to the life of every human being on this (and other) planet(s). Every moment of our lives we make decisions; even when we are not aware of doing so, our bodies make them for us. Decision-making is something that determines us, that determines our entire existence, and, one might say, is an ultimate purpose of our existence. I don't think it is short of the truth to say that we are what we decide.

Whatever our personal beliefs are and whatever we think decision-making is, it is what we all do every day of our lives. Decision-making influences every single activity in our lives. Our lives actually consist of the decisions that we make. After we die everything that we have done in our lives can be seen as a path through a giant maze or a huge graph showing our decisions at every intersection and the resulting consequences.

Modellers and especially actuaries, stockbrokers, and scientists are trying to predict the future. Most people would like to know the future so that they can make the best decisions. Unfortunately (or fortunately) we can't always accurately forecast the future. (How uninteresting life would be if we could!) But we can't - we don't even know if life system + other systems is a finite, discrete and countable system, which can be mapped out by a giant maze or graph (like a chess game); or an infinite, continuous and uncountable one.

But philosophy aside, let's get to the crux of the matter.

Chapter 2: The Evolution of Decision Systems – Road Toward the Agents

"Our environment shapes us – Byt kształtuje swiadomość" Karol Marx

Or: *"Our characteristics are shaped by the demands of our environment" Konrad Kulesza*

Decision strategies have been evolving through the entire history of mankind, but the evolution of decision systems has only occurred over the last 40 years, closely linked to the evolution of the computer itself. For a long time decision systems have (quite rightly) been closely linked to the evolution of Operating Systems; but let's take it step by step. The authors would like to divide decision systems into 3 categories. These three types, which still exist today, are obviously highly interlinked, and there is no mutual exclusion between them.

It is incorrect to say that a specific system will be newer or older just because it belongs to a specific category. Nevertheless the authors decided to divide them in this way.

The first category is what the authors have called Discrete Systems. They encompass systems built on the principals of logic and discrete mathematics. Examples are numerous. The simplest example would be tree traversing or a maze searching system. They are very commonly used in medical diagnosis, search engines, and engineering problems; you simply decompose the problem and follow a specific route. (Cormen, Leiseron, Rivest 1990.)

These were the first systems to come around, with beginnings in the 1960 and 70's.

Disadvantages are obvious; systems are often large, inflexible and very specific. Although they have grown since their beginnings and plenty of research has gone into them, they still leave a lot to be desired. (Cichosz 2000)

The next category is Neural Network Systems. They are a true attempt at creating Artificial Intelligence. It is no coincidence that the name 'neural' refers to the brain in everyday use. The structure of computer neural networks is analogous to the network naturally occurring in our brains.

What is a neural network? It is a loose combination of knots – neurons (processing + storage points) that combine within the network into a structure resembling a graph. The major advantage of a neural network is that it is not merely a static, unchanging system; instead it is constantly learning. It fully bonds with its environment and makes future decisions based on past experiences.

But even neural networks, promising inventions though they were, have their shortcomings.

The most obvious shortcoming is that they always make decisions based on what they experience. To give an example in terms of stock markets, the subject of my case study, assume we have a neural network making decisions about sales and purchases on the stock exchange. Practical evidence shows that if such a neural network had been trained during a boom, it will continue to play as if the boom hasn't ended, even during times of crisis. There are more problems with neural networks but the authors think this is the most important.

Before we get to the third category we must remember that the arrival of new decision systems was concurrent with the changes taking place in the environment (namely advances in computer technology, particularly computer hardware and software in particular operating systems). As computers were getting more powerful, software more sophisticated and the public more demanding, so were decision systems getting more sophisticated.

At the beginning of the 90's, the arrival of the internet and large scale networking caused a revolution within computing itself and lead to a telecommunications revolution, which has since been gathering in speed and which will definitely be a dominant phenomenon of the early stages of the present century. Although the revolution in computing is already 40 years old, the telecommunications revolution has only just begun and no one knows where it is going to take us. Single machine oriented systems are not enough any more; it is no longer about the computer, but the communication. Sun is right: "Network is a computer" (Blair et.al. 1997; Beauchamp 1987; Seybold 1994)

And at this point we meet Mobile Agents, the third category of decision systems.

Chapter 3: The Agent - What is this Beast? And what can it do for us?

"Being alive is the ability to change in order to keep on living" Konrad Kulesza

Definition:

"Mobile agents are agents that can physically travel across a network, and perform tasks on machines that provide agent hosting capability. This allows processes to migrate from computer to computer, for processes to split into multiple instances that execute on different machines, and to return to their point of origin. Unlike remote procedure calls, where a process invokes procedures of a remote host, process migration allows executable code to travel and interact with databases, file systems, information services and other agents." (Reilly 1998)

The following quote refers to Java mobile agents (aglets) but the characteristics mentioned are common to all mobile agents:

- "Mobility: Aglets can carry their code or data execution state with them from one computer to another across the network.
- Autonomy: Algorithms implemented in the code of aglets enable them to make local decisions on what to do, where to go and when to go.
- Concurrency: Multiple aglets can be dispatched simultaneously to accomplish various parts of a task in parallel.
- Local interaction: Mobile aglets interact with local entities, such as databases, file servers and stationary aglets through method invocation, while interaction with remote entities is by message passing.

- Flexible routing: The route traversed by an aglet can be predetermined, but it can also be modified dynamically by the aglet as it discovers additional information during its journey.
- Rapid response: An aglet can visit several sites, negotiating with local software at each site, and can return to its home base in only a few seconds" (Dasgupta, Narasimahan, Moser, Melliar-Smith 1999)

As we can see from the quote it is already implied that mobile agents posses some degree of operational freedom and wisdom, but one important characteristic should be added, namely:

- Intelligence: Ability to make decisions driven by the concern of own survival. Powered by genetic algorithms and the will to survive.

Adding this characteristic, we arrive at Intelligent Mobile Agents. The development of these agents has been a breakthrough for two main reasons:

Firstly, mobility is something that suits the network environment. Whilst all previous systems have been built with stand-alone machines in mind, agents are built with networks in mind. The development of agents responds to an ever-increasing need to modernize the way in which we think of solving problems on the network.

The second reason refers to their intelligence. Intelligent mobile agents are the most refined form of decision systems that we have yet created. Their main goal is to collect information for their owner; their very survival depends on their ability to do so. This philosophy closely mirrors the real world. The agent, unlike the neural network or discrete system, has a personal interest in making things happen. He will do anything to survive: he will evolve, find the shortest path through the network, lie to you or fight his way with other agents.

Agents are the future of networking and Artificial Intelligence, which over the last 30 years has been disappointing as a domain. The authors see their survival instinct as the breath of life inside them.

It is amazing how analogical the world is. If we look at the evolution of computers we see amazing similarities to the evolution of life on this planet. Life began with small primitive organisms that in software development can be compared to the elementary software of the 70's. As organisms evolved they became more complex, first reaching monstrous sizes, for example the dinosaurs; later becoming smaller in size but more intelligent and refined. A similar evolution has occurred with computers, from huge, clumsy structures to the tiny, refined, more intelligent devices of today: networking, open communications and distributed systems seem to demand that of us today and agents can be the answer.

Referring back to my case study, as we all know the global financial system is probably one of the most sophisticated interaction systems ever constructed by human beings. Nowadays it has grown so big and sophisticated that the rich of this world have long since lost control over it; no one seems able to control it. Many people see it as a great danger that human beings cannot control one of their own prime activities. If we push this logic further the clear implication is that agents will introduce more chaos and instability into an environment that is already chaotic and out of our control.

The authors don't think that the environment being unpredictable is such a tragedy; moreover the authors don't think that agents are so terrible either. Just because the environment isn't steered by some central force, that doesn't mean that it is not controlling itself. That the reverse is true has been proven both in practice and by many scientists (including Hayek, whose work is probably one of the strongest defences of the free market system). Evolving self-controlling environments are better then centrally planned ones because units only need very specific knowledge to coexist and they have time to experiment and have evolved in very specific ways. People who think that they can do better than coexisting systems that have evolved over time are doomed.

Agents are useful and friendly. However, there are still many obstacles to overcome. More on this in the next chapter.

Chapter 4: What has been done so far? What remains to be done? What are the challenges?

"You reap what you sow"

Mobile agents have had a varying degree of success. Over the last three years they have been losing some of their popularity to other emerging technologies. However, the authors don't think that this trend will continue, but more about this later.

- Requirements for mobile agents:

"In general, the following things are required to allow agents to migrate across a network:

1. Common execution language
2. Process persistence
3. Communication mechanism between agent hosts
4. Security to protect agents and agent hosts" (Reilly 1998).

The authors believe these requirements are fairly obvious and as the authors don't have the space to discuss them here, feel free to consult the reference with any problems.

- Competing technologies:

1. Message passing systems

"Software agents need not always travel across a network to communicate with information sources, or other agents… Message passing systems like KQML don't require mobility, they can simply pass a message and have it delivered through some transport mechanism." (Reilly 1998)

2. Remote Method Invocation (RMI)

"New objects can be transferred across the network, and RMI is becoming a popular mechanism for agent communication. RMI can be used to facilitate mobile agency (acting as a transport mechanism), or as a replacement that allows agents to invoke methods of other agents." (Reilly 1998)

3. Common Object Request Broker Architecture (CORBA)

"CORBA is a platform and language independent mechanism for invoking remote object methods. Unlike RMI, which is specific only to Java Virtual Machines, CORBA can be used to create distributed systems that execute on many platforms, in many languages. CORBA holds great potential, because of its portability and flexibility. CORBA is a direct threat to mobile agency, and would allow developers to create agents that are capable of complex communication without ever travelling across a network."

- Shortfalls of mobile agency

Since mobile agents have to operate in heterogeneous environments in which many different system architectures are connected, interpreted scripting languages or emulation of a system that is capable of executing machine code are used. Both approaches are relatively slow. The authors proposed solution to this problem is to use mobile agents in conjunction with emerging technologies, let them evolve more and make use of many different solutions in order to optimize compilation, and finally often cut across compilation and do a skip compilation when under pressure (e.g. SPECNAZ method).

The next problem concerns the persistence property, which puts additional weight on the system since its function is to save execution states and all related variables, which then have to be converted into data suitable for transmission over a network. Alternatively a new process has to be created. The authors proposed answer is to create special baby agents and through cooperation between them and their genetic mutation only send data through, as it is required.

Another problem concerns the communication mechanism between agents across the network. It is obvious that such mechanisms must exist to transfer agents across the network. Here a variety of protocols come into play such as: TCP/IP, RMI, IIOP, SMTP or HTTP. According to the authors of the first reference (Reilly 1998), "Mobile agent architectures may even use a variety of transport mechanisms, giving greater flexibility." But flexibility comes at a price…

They go on to say: "An agent's executable code must be transferred, which may consume a large amount of network bandwidth, unless shared code is located at the agent host. Techniques such as shared libraries of code or caching, may be of benefit. In addition, the persistent state of the agent must be transferred". To this list of possible solutions the authors of this paper would add the following: let the agent decide what it wants to do; or at least work with it as a team instead of acting like an inflexible central planner. The network environment is very fluid and it is the agents who live within it. The agent has an incentive to be as efficient as possible; its very survival depends on it.

The next issue is a security consideration. Of all the available technologies, agents are the most prone to damage. Because they are the most sophisticated of the available technologies, they may find themselves being destroyed or interrogated by a third party. Also "Encryption may be of benefit, but the data and code must be decrypted at some point in time for the agent to execute. Once this occurs, the agent becomes vulnerable, and is at the mercy of the agent host…" (Reilly 1998). However, the article referred to also proposes some widely used solutions, too conventional for my liking, but nevertheless a common approach toward security in agent technology, which has simply mirrored security approaches from other types of systems. This practise has to change if agents are to succeed; we must remember that agents are dynamic, interactive and distributed; their very nature requires us to start thinking differently. I think that agents, because of their openness to attack, should themselves be more offensive toward potential interrogators than other systems. They should send their own replicas to deceive potential interrogators and destroy them if necessary. They should even spy if they have to; survival strategies should be built within agents to allow them to combine and fight if they have to. Readers interested in security considerations are encourage too consult one of the authors work (Kotulski 2002).

Another attack on mobile agent technology is the accusation that they were developed to conserve bandwidth yet they have the potential to do the reverse. First, additional bandwidth is required to send agents across the network, and second, many people say it is better to have "a small number of indexing agents collect information for search engines, while millions of queries are made by users. Imagine if the same number of queries were made instead by mobile agents that travelled across the network to sites. Two scenarios are possible. Either a much larger amount of bandwidth will be consumed as users receive more accurate search results because their agents have more control over the search process. Instinct suggests, however, that a simple keyword query entered via a web browser will consume less resources and bandwidth than sending an agent with specialised searching algorithms across the network" (Reilly 1998). Although the authors agree with this sentiment, in they opinion it does not in any way undermine the value of agents. Deep down the authors believe that left to themselves agents would evolve into efficient search machine systems; but the authors don't think that we need to reinvent the wheel. Agents should work in cooperation with existing technologies; there is no one method that is best for solving all problems.

These, in the authors opinion, are the main problems facing mobile agency, and the authors solutions to them. You can read more on this subject in the references. Particularly worth reading are "Mobile Agents – Process migration and its implications," (Reilly 1998) and "Mobile Agents: Are they a good idea?" (Harrison et al. 1995).

The authors don't agree fully with either of them, since their views on mobile agents are pessimistic, while the authors optimistic about their potential application. It has unfortunately been a sad truth of the last four years that agents have been losing popularity to other technologies; however, the authors don't think this will continue to be the case: of all similar technologies, only agents are truly an intelligent (as opposed to manual) tool. In this paper the authors have purposely built the topic up from the base to show how enormous a breakthrough the very idea of an agent has been. There simply is no other system like it. Even neural networks, a huge breakthrough in itself (the first self-learning system) can't be compared to the revolutionary idea of agent survival. Yet the authors do respect other views and only the future will show who is right.

But independent of these discussions mobile agents have already made an impact both in their research and commercial applications; you can read more about this in my second reference. The authors will only quote a few figures and basic requirements from a business perspective (Dasgupta, Narasimahan, Moser, Melliar-Smith. 1999). Some interesting experimental results include the following: it took approximately seven seconds for an agent to complete the trip around the world; a small agent can travel from California to Japan, conduct transactions, and come back in approximately nine seconds. The work also provides an interesting introduction to ecommerce in general. For the purpose of this paper the authors will only quote six stages of the consumer buying model, namely: need identification, product brokering, negotiation, purchase and delivery, and product service and evolution.

Generally people agree with the theory and idea of mobile agents. Differences still exist on implementation and interpretation, which may be a good thing because generally monopolies are unhealthy. There are countless proposed implementations ranging from Functional Programming (which is very useful due to its usefulness in programming communication), Logic Programming, and, of course, Java solutions.

Genetic Algorithms can be very useful when working with mobile agents and here the authors think a book worth mentioning is my six reference, "Genetic Algorithms plus Data Structures = Evolutionary Programs."(Michalewicz 1996). The authors think it is a revolutionary book and the authors think that its future impact could be as great as the famous work from the 60's that it resembles so closely in title. But maybe they are biased because the author is Polish.

When talking about networking there are a few things we must remember. The network system is a lot more complex and governed by different rules than stand-alone machines. On the network it is not simply a case of 1+1=2; it is an environment with many distortions and delays. No one has yet constructed a true distributed operating system, even though many great O.S. minds (e.g. Professor Tranbout Holand) have been working on it. Challenges in networking are huge; nothing is discrete, everything is fuzzy.

My area of interest lies in designing algorithms for network and distributed processing. There are many complex problems in this domain; however, most amazingly, many solutions to these problems already exist and they are all around us. They are most often found in analogies to real life situations. This idea is nothing new – exactly the same principle holds in the domain of operating systems where solutions to difficult tasks like processor time management etc. have been deduced from simple everyday things (e.g. the lift algorithm based on the functioning of a normal lift). The same methodology, in my belief, can be applied to distributed environments. I have already mentioned SPECNAZ translation, which is my term for a very basic skip through compilation instead of phasing the entire code and translating everything, for which we may not have the tools and time.

You only take the things that are most important to you; we can come up with all sorts of back bag algorithms designed to save space and time for a trailing agent.

As I said earlier there are countless strategies offered as solutions to the same problem of implementation of mobile agent technology. My solution is to create a mobile agent by setting up a class called DNA, which will encompass all the information about the agent, its properties and functions. This class will deal with the memory of an agent, its size, allocation of storage space and processing capabilities, and senses. We can actually group together a couple of the most important characteristics that an agent will posses and call them genes. Now the whole idea is to let an agent mutate and optimise it. This process can happen according to some functions (e.g. inc.; dec.) performed on its characteristics or some very advanced functions, which can mutate as well or be nested e.g. congestion increases, size of agent decreases. Also, random mutations should be allowed because otherwise all we get is optimisation, and true advances happen through breakthroughs, not optimisation. A special class of random mutations would be creative mutations, and these are the hardest to implement. We don't know if they even exist in nature; they are not based on any known analogy. What is important to remember is that the stock exchange never goes to sleep; it starts in Tokyo and finishes in New York. Agents have to be ever-living phenomena and when they are not doing anything they are using time and energy to mutate and evolve, trying to improve, especially when they are not performing tasks. It's amazing how unified the world is; problems in one domain have answers in another completely unrelated one. The whole universe blends together in creating this magical world that we see around us.

In the coming century telecommunications and genetics will be the two sciences churning out the most breakthroughs; they can definitely learn a lot from each other. In an episode of "Beyond 2000" a scientist was using computers to simulate the long process of evolution. He was playing with his creations. I found that very interesting. Modern computers could learn a lot from genetics and vice-versa.

Agents should be able to split and get together and make use of other technologies. Moreover, strategies can be built and nested inside the DNA (business strategies, math strategies, even strategies from logic games like chess). Agents will get bigger or smaller depending on the demands of the environment. The function of changes could also be influenced by the environment e.g. two types of decision functions: make decisions quickly, or slowly but more accurate, with all phases in between the two extremes defined by fuzzy logic D->(0; 1).

The sky is the limit.

Chapter 5: Conclusion -What does the future hold? Ethical issues. How much do we want to make things happen?

"May all your wishes come true"- Chinese curse

"*We could well become the victims of our own success, as has been the case in the history of human beings till now*"- Konrad Kulesza

There are huge opportunities for agents. There are also many unknowns: how good is this invention and where it will lead us? The most important ethical questions are: How many restrictions should be placed on agents? As we know, too much regulation kills initiative, whilst too little creates dangerous situations. An example is the cow. Cows can't really create anything, but they can't be evil either. Humans are the only species that can damage the earth, commit mass genocides, or suicide; but they are also the only species that can produce their own food. The second issue is whether agents should be allowed to breed uncontrollably or should there be some limitations; or should we let them control themselves. There is a phrase that comes to mind from the movie, 'The Matrix:' "Human beings are not mammals, because every mammal develops a natural equilibrium with its environment, but humans just spread and spread….”

Agents are a promising technology and no one knows where they can take us, but as scientists it is our responsibility to think of the consequences.

Acknowledgment

The authors would like to thank Anton Nahman for his editorial input. The authors would also like to thank Kamil Kulesza for his constant support and advice.

Bibliography

1. BEAUCHAMP, . 1987. Computer Communication. Van Nostrand Reinhod (UK) Co. Ltd; Molly Milkers Lane, Wokingham Bersksihre England Reprinted 1988.
2. BLAIR, G.J., et.al., 1997, Open Distributed Processing and Multimedia. Addison-Wesley Pub. Co.
3. CICHOSZ, P. 2000. Systemy uczące się, WN-T, Warszawa,
4. CORMEN, T.H., LEISERON, C.L., RIVEST, R.C. 1990. Introduction to Algorithms. Massachusetts Institute of Technology,
5. DASGUPTA, P., NARASIMAHAN, N., MOSER, L.E., and MELLIAR-SMITH, P.M. 1999. MAgNET: Mobile Agents for Networked Electronic Trading. Department of Electrical and Computer Engineering University of California, Santa Barbara, CA 93106 **http://alpha.ece.ucsb.edu/~pdg/research/papers/MAgNEThtml/MAgNET.html**
6. HAROLD, E. R. et.al., 1997. Java Network Programming. O’Reilly&Asocc.,
7. HARRISON, C.G, et al. March 28 1995. Mobile Agents: Are they a good idea? [online] at **http://www.research.ibm.com/massive/mobag.ps**
8. KOTULSKI, Z. 2002.Nowoczesne technologie informatyczne: bezpieczeństwo danych. in: M.Kleiber, [ed.], *Nauki techniczne u progu XXI wieku.* IPPT PAN, Warszawa, pp. 181-210.
9. MICHALEWICZ, Z. 1996. Algorytmy genetyczne + struktury danych = programy ewolucyjne. WN-T, Warszawa. (translation from English)
10. REILLY, D. 1998. Mobile Agents – Process migration and its implications. http://**www.davidreilly.com/topics/software_agents/mobile_agents/**
11. SEYBOLD, A.M. 1994. Using Wireless Communication in Business.
12. SUBRAHEMANIAN, V.S. 1998. Principles of Multimedia Database Systems. Morgan Kaufmann. San Francisco.