AGENT SYSTEMS TODAY; METHODOLOGICAL CONSIDERATIONS

MARCIN PAPRZYCKI and AJITH ABRAHAM
Department of Computer Science, Oklahoma State University, Tulsa, USA
email: {marcin, aa}@cs.okstate.edu

Abstract: For many years a message is being perpetuated that software agent technology will become the next revolution with consequences reaching beyond computer science. This change is to occur not only in the ways we construct software but also to have a broad effect on human-computer interaction. Unfortunately, as it is easy to see, the revolution that is prophesized by the agent-believers since at least 1994 does not materialize. The aim of this note is to discuss the reasons for this situation and point out to the necessary steps to achieve substantial progress in the field.

Keywords: Agents, agent systems, e-commerce, agent system development and applications, methodological reflection

1 Introduction

This note is being written approximately 10 years after the birth of one of the most influential and quoted papers in the theory of agent systems, '*Agents that Reduce Work and Information Overload*' by P. Maes [5] and five years after the birth of one of the most prophetical and controversial papers '*A Perspective on Software Agents Research*' by H. Nwana and D. Ndumu [15]. Due to its controversial nature, the later paper belongs to the list of most forgotten papers about practical aspects of the development of agent systems. The former paper represents those who believe that agent technology is about to become an extreme event that will introduce changes to many facets of our lives. These expectations have propagated far outside the traditional domain of computer science. For instance, a recent (1999) issue of Nature magazine posed the question, 'Is There an Intelligent Agent in Your Future?' [1]. This change is to occur not only in the ways we construct software [2, 4] but is also to have a much broader impact on the field of human-computer interaction [1, 3, 5]. Some of the principle areas software agent technology is expected to have an impact are [1, 2, 3, 4, 5, 23]:

- development and maintenance of complex systems,
- resource management,
- delivery of personalized content,
- e-commerce on a large and small scale.

Obviously, this list is far from exhaustive, and there exists considerable overlap between these application and research areas as well (for instance, fully developed e-commerce systems not only have to belong to the category of complex software systems [21], but also have to contain a robust system for delivery of personalized content [22]). Obviously, the breadth and depth of these areas supports the claim that software agent technology, if successful, can become the next extreme event, leading to breakthroughs in a wide variety of fields, including those especially relevant to the corporate world, such as management and economics (and in particular, all areas related to the Internet-based business and economics). The agent paradigm also promises to have a profound effect on the social fabric, adding a new dimension to our perception of and interaction with computers.

Unfortunately, as it is easy to see, the revolution prophesized by the agent-visionaries (work of P. Maes could be used as a paradigmatical example) thus far did not materialize (regardless of the rapidly increasing number of conferences, workshops, special sessions, publications, etc). It is not the case that when we turn the computer on in the morning, we contact (or are contacted by) "our personal agent" and receive a personalized newscast and our day-plan. In addition, on the basis of that plan as well as the weather forecast and knowledge of our dressing-preferences, our agent gives us a friendly advice what to wear (agent ideal servant). Similarly, when developing software for an e-store we do not utilize pre-existing agent-modules (e.g. marketing agents, price negotiators, inventory managers etc.) that would provide us with a natural decomposition of the task into functional-components. To the contrary, it is difficult to point to a successful large-scale commercial implementation of an agent system. In particular, the three examples of agent systems described in [5]: the electronic mail agent, the meeting scheduling agent and the news filtering agent, have never become popular outside of the MIT Media Laboratory (and it would be really interesting to learn if they are still being used inside of the MIT). This could indicate that the strength of pragmatic arguments presented in [15] by Nwana and Ndumu has advantage over the optimism of agent-believers. To find out if it is really the case, it is time to have a close look at the field of development of agent systems and see what has changed over the last five years and see for how much longer arguments put forth in [15] will remain valid.

2 Problems

2.1 Definition

Before we turn our attention to the work of Nwana and Ndumu, let us note that there exist a very simple way in which one may notice an existence of a problem in the "world of agents." It is by trying to pinpoint a definition of a software agent. When browsing literature one finds that the following four definitions are very often quoted by a number of authors (while other definitions are also used):

Definition 1

Anything that can be viewed as perceiving its environment through sensors and acting upon that environment through effectors (Russell and Norvig, 1995).

Definition 2

Software entities that carry out some set of operations on behalf of a user or another program with some degree of independence or autonomy, and in so doing, employ some knowledge or representation of the user's goals or desires (IBM, 1997).

Definition 3

Autonomous system situated in a dynamical environment acting independently of its restrictions and fulfilling in it a set of goals or directives for which it was created (Maes, 1998).

Definition 4

An encapsulated computer system that is situated in some environment and that is capable of flexible action in that environment in order to meet its design objectives (Jennings et al, 2000).

While there are some similarities between these definitions (a somewhat autonomous software artifact interacting with its surroundings), it is rather difficult to make the case that they are defining exactly the same entity and can be substituted for each other. While this argument is relatively weak (as apparently there exist more than 400 definitions of information and this fact does not have an adverse effect on development of "information processing sciences and technologies"), the very fact that an agent is not a well-defined entity raises some concern. It also has some peculiar consequences, as further study of literature reveals that there exists a widely-varied array of features which are used to describe software agents [6] (and again, this list is not exhaustive):

- autonomy,
- reactivity,
- ability to communicate,
- capacity for cooperation,
- capacity for reasoning,
- capacity for learning,
- adaptivity,
- intelligence,
- goal orientation,
- proactivity,
- mobility,
- robustness,
- reliability,
- scalability,
- flexibility,
- reusability.

It could be observed that these characteristics overlap, and their own definitions could lead to serious conceptual disagreements between scientists representing various disciplines. For instance, the question what exactly is the relationship between adaptivity, capacity for learning and intelligence remains difficult to delineate. In practice, different researchers take a separate subset of the above characteristics as their core definition and study agent systems satisfying these requirements [6]. On the one hand, this further raises the questions about dealing with an entity without a clear definition. However, it also facilitates a broad spectrum of research projects involving software agents. Even though sometimes agent systems could be utilized in a very limited scope, an important aspect here is their practical importance (e.g. developing software agents that traverse Internet in search of the best price [7]), on the other hand theoretical research utilizes agent-metaphor to study human / social behavior / interactions (multi-agent systems [8]).

In our research, we are interested in developing agent-based systems that support Internet-related activities (e.g. e-commerce systems, personalized information delivery, etc.). In particular, we are interested in developing an agent based e-travel agency [9, 10, 11, 12]. While the description of this project is outside of the scope of this note, in approaching our research, we have noted that a very large number of agent system development projects have been initiated but never completed (implemented). For instance, in the area of agent-based travel support alone, we have located eight web sites of unfinished projects, and two limited in scope demonstrations [13, 14] (which also have been later abandoned). We therefore pose a question, "why agent systems have not been successfully developed/implemented?" To simplify the answer it is better to go back to the arguments presented in [15].

2.2 Six major issues of agent based technology

Since Nwana and Ndumu were co-authors of two early demonstrator systems applying agent technology to develop an Internet travel agency [13, 14], they used this very area to launch their criticism of the state of the art of agent system research (they also point out, that a travel agent is a perfect metaphor and a paradigmatic case for what we should strive for in developing software agents). They cite the following primary reasons for the failure of agent technology (these most important problems are followed by a number of slightly lesser important and/or more detailed ones):

(i) *information discovery problem* (how to extract information in the situation, when the web sites are constantly changing their structure and pages appear, disappear or change their location; where the design of pages is also highly dynamic and an agent crafted to be capable of extracting useful information from a web site on a given day may fail completely the next day; when the total amount of information is increasing exponentially and includes both the reliable and unreliable information),

(ii) *communication problem* (how to assure communication between agents, in particular mobile agents, as well as communication between agents and other entities of the Internet; what language should all these entities speak to be able to successfully communicate with each other),

(iii) *ontology problem* (how to assure that agents have at least a minimal "understanding" of the meaning of information they are accessing and processing; here the question is not only how to introduce a semantic description into a particular limited sub-area of knowledge, but also how to develop a general support for the semantic demarcation of the Internet content so that agents will be able to consume and utilize it efficiently and effortlessly),

(iv) *legacy software integration problem* (how to deal with the enormous amounts of legacy software that is not agent-ready even it is already accessible through the Internet),

(v) *reasoning and co-ordination problem* (how to reason about the extracted / located information and, in the case of a travel system how to combine this with an ability to coordinate activities by adding to the system's ability of temporal reasoning),

(vi) *monitoring problem* (a problem that is very specific to the travel support system, and involves constantly monitoring the status of transactions to assure their successful fruition, i.e. avoiding pitfalls related to weather-related travel delays etc.).

Five years later, most of these problems have not been resolved. Nonetheless, interest in agent technology and its potential has grown considerably. We have reached the point when, in a recently (2000) published paper, Jennings argues that agent technology is one of the more promising approaches to the development of large complex systems [8]. In addition, we observe a constant stream of new agent toolkits and platforms being developed. For instance, in [9] authors evaluate 30 most popular agent platforms, while the total number of existing ones already exceeds 90. Finally, the exponential growth of the amount of data on the Internet brings back arguments put forward by P. Maes in [5]. If we do not find a new way of dealing with information overload, the Internet, instead of fulfilling its promises of becoming the ultimate source of useful (and complete) information and the ultimate equalizer of access, will become one more source of pollution, the cesspool of data collapsing under its own weight and informational entropy. We may therefore need agents after all. Let us therefore reflect what is the state of the research in solving the above-described problems of agent technology. In other words, what has changed in the agent system research over the last five years?

3 Responses and possible solutions

3.1 Current hot topics

It is clear that progress has been made in addressing all problems, as they became one of the hottest and most fashionable directions of research.

(a) A number of technologies for extraction of information from both structured and unstructured web resources has been proposed and experimented with. While some of them are general and would have to be modified to be included in the agent's domain [25, 26, 27, 29], others are already described in the context of agent systems [24, 28, 30]. These advances, combined with growing sophistication of search engines lead us to believe that the information discovery problem has been solved to the degree that a new generation of search agents is possible.

(b) A substantial amount of work has been done by the Foundation for Intelligent Physical Agents

(FIPA) [31] in proposing and promoting standards related to various aspects of agent system development and implementation; including a standard for agent communication (the ACL, agent communication language). This standard, as well as others have been implemented in a number of agent platforms (e.g. JADE, Grashopper, LEAP etc.). It can be thus be stated, that also the communication problem has been solved to a relatively satisfactory degree.

(c) With the enormous success of Java it is easy to see how it is possible to wrap the existing legacy software in a thin Java interface and this way claim a relatively natural way of handling legacy software (and such a solution has been proposed in many occasions). While this may be a long-term process, but it is clear that a path is set and there exists more and more tools to support it.

(d) Nwana and Ndumu have already acknowledged the relatively positive state of the art of reasoning and coordination methods and existing software. Thus any further improvements in our understanding of these aspects of intelligent system design and application further diminish the importance of this problem.

(e) Finally, we need to understand that the monitoring problem is somewhat less important in the context of general agent system development, as it is directly related to e-travel systems discussed in [13, 14, 15]. At the same time, the research in temporal reasoning is also steadily progressing slowly reducing its impact.

A slightly more complicated is the ontology problem (Section 2.2, item iii). While this is also a very hot research area, the progress is limited not only by the lack of tools, but also by the lack of "material" to work with. In other words, the vision of the Semantic Web [17, 18] requires a large number of pages to be "semantically demarcated" for sensible experiments becoming available to establish feasibility of proposed approaches. While the standards for the RDF framework have been proposed and there exists more than 20000 pages denoted using DAML (and some more denoted using some of the older ontology descriptors), this is definitively not enough. Full utilization of the power of the models developed to support the vision of the Semantics Web, can only emerge when the content described in terms of ontologies hits a "critical mass". This is also hinted by our work analyzing the existing "ontologies" in the travel domain [19]. The results have shown, that there exists a number of possible categorizations and only through experimental work with pages denoted using them it will be possible to find the best possible approach (or establish that a multitude of approaches is also useful and tractable). At the same time, a number of approaches have been proposed combining the world of agents and the ontology-oriented research [17, 32]. Summarizing, while the ontology problem seems to be the one with the least progress, it can be claimed that even a new generation of agent systems can be supported by recent developments.

3.2 The integration problem

The question thus arises, why still such a lack of success. Here, we would like to suggest a slightly different way of looking at the problems that agent technology is facing. History of development of agent technology (with its visionaries and unfulfilled promises) is somewhat similar to the history of artificial intelligence: every now and then someone promises that the final answer is just around the corner, and only 10 years later we find out that even if this was not a "dead end," we are almost as far from the ultimate goal as before. Moreover, these two histories are very much entwined as if software agents are to be intelligent, then what is needed is working artificial intelligence. Such intelligence could, for instance, be embodied in a classic expert system (with knowledge base and an inference engine [20]), but it is known that such systems work well in restricted domains and typically do not generalize beyond them. Thus, how are we going to build intelligent agent systems capable of acting in such a heterogeneous domain as the Internet?

When pursuing the comparisons between agent systems and artificial intelligence it is possible to note that there is one more area of similarities. During the course of its history, artificial intelligence has progressed rapidly in a number of relatively independent areas, but there have been very few serious attempts at reconciling them. For instance, separate groups have spent their research efforts to develop theories leading to creation of very useful and practically applicable technologies: heuristics, probabilistic computing, rule-based expert systems, fuzzy inference systems and evolutionary systems, among many others. Until only recently, these theories and technologies have been developed in a relative separation and even today, attempts at combining them and developing hybrid systems are few and far between. Moreover, the divisions do not stop at this level, e.g. in the area of neural networks, researchers have been perfecting underlying theory and practical tools for different approaches, e.g. multilayer perceptrons and self-organizing maps; without much research devoted to the possibility of combining them (for instance, where a self organizing map is used to divide the data into broad “categories” a feedforward multi-layered perceptron is often used inside each category to make fine-tuned predictions).

In this context, the computational practice of today tells us that, although there are classes of problems that are best solved by one technology or another, while most of large scale real-life problems will require a combination of one or more of these systems (as they cannot be successfully tackled by either one of these approaches alone). Therefore, in the case of artificial intelligence research, hybridization of intelligent systems is touted as a promising approach that may result in development of the next generation of intelligent systems. Here, a fundamental stimulus to the investigations of hybrid intelligent systems is the awareness that combined approaches will be necessary if the tough problems in artificial/computational intelligence are to be solved.

These considerations are strikingly similar to those confronted by many other areas of information technology in general and agent systems in particular. All of them can be conceptualized as a *lack of theory and technology integration*. In the case of our main interest: agent systems, a very large number of theories and technologies originating from multiple research areas (each of them developed separately, with their own particular applications in mind) must now come together to form a cohesive system. In this way, our analysis of the six problems of Nwana and Ndumu is pointing to a completely different direction as to where the answer can be. Instead of looking for perfect solutions to each separate area, we believe that it is the problem of **combining currently existing answers** into a cohesive unit that is the most striking problem. As an example, let us briefly survey some of the research areas that provide theoretical and practical research results, technologies and tools involved in the design of a successful agent-based e-commerce system [21] (this list is not exhaustive, but its breadth already illustrates our main point):

- marketing and personalization,
- knowledge management,
- real-time systems,
- distributed resource management, scheduling,
- expert systems,
- human-computer interaction,
- ontology integration,
- machine learning,
- large scale distributed computing
- knowledge discovery,
- network management.

It is only when the “best of the breed” in each of these areas is combined together, then the next generation of agent systems will emerge and we will be finally able to evaluate the ultimate truth of the positive program of P. Maes [5] and the pragmatical criticism of Nwana and Ndumu [15].

4 Concluding remarks

In this note we have attempted to address a simple question: why, regardless of years of intensive research and increasing interest in the field, agent technology is still at its infancy? Or, to put it in a different way: why (ten years after such an optimistic paper by P. Maes) do we have to write this note to ourselves, instead of giving the general parameters and most important points to our personal research agents and leaving them to hammer the details and provide us with the initial draft of the manuscript?

To find our answer to this question, we have first depicted the problem of defining the very idea of a software agent. Then we have proceeded to analyze the current state of finding solutions to the six fundamental problems posed by Nwana and Ndumu. Except of the ontology problem, which is one of the hottest research areas today and should deliver initial solutions shortly, we have found that most of their problems have been addressed to the extent warranting a new generation of agent systems to be developed. Thus the question remained open: why this is not materializing? Our answer is: while it is relatively easy to find answers to separate questions, it is much more difficult to make these answers working together.

This being the case, the necessity of combining technologies and tools originating from various disciplines should define the approach to the development of agent system. *We believe that agent technology is the means to this end, and that, by delivering on its promises, it has a definite chance of becoming the next extreme event.* In order for this prospect to materialize, we need to commence developing and implementing such systems. We should start from small demonstrators combining basic technologies, and proceed toward more sophisticated systems (it is worth nothing that this approach is pursuing the positive suggestions put forth by Nwana and Ndumu in [15] and specified by them as the *only way* in which a realistic progress can be achieved).

Obviously, we cannot forget, that an agent system is ultimately an intelligent system. For such a system to be truly intelligent (however this property is to be ultimately defined) in a distributed environment the basic questions like: how should agents interact, cooperate and compete to meet the overall objective all need proper answers. Coordinating specific agent capabilities to group level rules, principles is always a challenging task. Decomposition of goals/subtasks, communication, ability to reason/action, learning from mistakes etc. are a few to mention. Once we are able to find good frameworks to these challenging issues, we will be able to simulate meta-level intelligence. However, it has to be stressed again, all the theoretical work will not be able to reach its goals until a true experimental work commences. Until we show that we can make all the necessary technologies work together it will be hard to visualize wide usage of agent based systems for solving complex real world problems.

Results of these investigations may naturally be applied to other areas of the information technology field, using these same building blocks as a foundation for the next generation of human-computer systems.

References

[1] J. Hendler, Is There an Intelligent Agent in Your Future?, Nature, http://www.nature.com, March 11th, 1999.

[2] M. L. Griss, My Agent Will Call Your Agent … But Will It Respond?, Technical Report, Hewlett Packard, 1999, http://www.hpl.hp.com/techreports/1999/HPL-1999-159.pdf.

[3] N. R. Jennings, An agent-based approach for building complex software systems, *Communications of the ACM*, 44 (4), 2001, 35-41

[4] T. Berners-Lee, J. Hendler, O. Lassila, The Semantic Web, *Scientific American*, May, 2001, http://www.sciam.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21

[5] P. Maes, Agents that Reduce Work and Information Overload, *Communications of the ACM*, 37(7), 1994, 31-40

[6] V. Galant, J. Tyburcy, Intelligentny Agent Programowy, *Prace Naukowe AE Wrocław*, Nr 891, Wrocław, 2001, 46 – 57, in Polish.

[7] http://www.metacrawler.com

[8] J. Ferber, *Multi-Agent Systems*, Addison-Wesley, London, 1999

[9] R. Angryk, V. Galant, M. Gordon, M. Paprzycki, Travel Support System – an Agent-Based Framework, Proceedings of the International Conference on Internet Computing (IC'02), CSREA Press, 2002, 719-725.

[10] M. Paprzycki, R. Angryk, K. Kołodziej, I. Fiedorowicz, M. Cobb, D. Ali, S. Rahimi, Development of a Travel Support System Based on Intelligent Agent Technology, in: S. Niwiński (ed.), *Proceedings of the PIONIER 2001 Conference*, Technical University of Poznań Press, 2001, 243-255.

[11] M. Paprzycki, M. Gordon, P. Harrington, A. Nauli, S. Williams, J. Wright, Using Software Agents to Index Data for an E-Travel System, *Proceedings of the ABC Symposium*, Orlando, FL, 2003, CD, file: 001428.pdf, 6 pages.

[12] M. Paprzycki, P. Kalczyński, I. Fiedorowicz, W. Abramowicz, M. Cobb, Personalized Traveler Information System, in: B. Kubiak, A. Korowicki (eds.), *Proceedings of the 5th International Conference Human-Computer Interaction*, Akwila Press, 2001, 445-456.

[13] D. Ndumu, J. Collins, H. Nwana, Towards Desktop Personal Travel Agents, *BT Technological Journal*, 16 (3), 1998, 69-78.

[14] J. N. Suarez, D. O'Sullivan, H. Brouchoud, P. Cros, Personal Travel Market: Real-Life Application of the FIPA Standards, *Technical Report*, BT, Project AC317, 1999.

[15] H. Nwana, D. Ndumu, A perspective on software agents research, *The Knowledge Engineering Review*, 14 (2), 1999, 1-18.

[16] J. Altmann, F. Gruber, L. Klug, W. Stockner, E. Weippl, Using Mobile Agents in real World: A Survey and Evaluation of Agent Platforms, *Proceedings of the Second International Workshop on "Infrastructure for MAS, and Scalable MAS,"* Montreal, Canada, May 28 – June 01, 2001.

[17] J. Hendler, Agents and Semantic Web, *IEEE Intelligent Systems Journal*, 16(2), 2001, 30-37.

[18] http://www.semanticweb.org/.

[19] A. Gilbert, M. Gordon, M. Paprzycki, J. Wright, The World of Travel: A Comparative Analysis of Classification Methods, in press

[20] W. Brenner, R. Zarnekow, H. Wittig, C. Schubert, *Intelligent Software Agents*, Springer-Verlag, 1998.

[21] V. Galant, J. Jakubczyc, M. Paprzycki, Infrastructure for E-Commerce, in: Nycz M, Owoc ML (eds.), *Proceedings of the 10th Conference on Knowledge Extraction from Databases*, Wrocław University of Economics Press, 2002, 32-47.

[22] V. Galant, M. Paprzycki, Information Personalization in an Internet Based Travel Support System, in: Abramowicz (ed.), *Proceedings of the BIS'2002 Conference*, Poznań University of Economics Press, Poznań, Poland, 2002, 191-202.

[23] Hendler J, Agents and Semantic Web, *IEEE Intelligent Systems Journal*, 16, 2, (2001), 30-37.

[24] L. Arjona, R. Corchuelo, A. Ruiz, M. Toro, A Practical Agent-Based Method to Extract Semantic Information from the Web, *Lecture Notes in Computer Science*, 2348, 2002, 697-702.

[25] C. A. Knoblock, K. Lerman, S. Minton, I. Muslea, Accurately and Reliably Extracting Data from the Web: A Machine Learning Approach, *IEEE Data Engineering Bulletin*, 23(4), 2000, 33-41.

[26] V. Gudivada, V. Raghavan, W. Grosky, R. Kasanagottu, Information Retrieval on the World Wide Web, *IEEE Internet* Computing, 1997, 57-68.

[27] W. Abramowicz, P. Kalczynski, K. Wecel, Filtering the Web to Feed Data Warehouses, Springer Verlag Publishing, New York, 2002.

[28] J. L. Ambite, C. A. Knoblock, Agents for Information Gathering, *IEEE Expert*, September/October, 1997, 2-4.

[29] A. Kehagias, V. Petridis, V. G. Kaburlasos, P. Fragkou, A Comparison of Word- and Sense-Based Text Categorization Using Several Classification Algorithms, *J. of Intelligent Information Systems*, 21(3), 2003, 227-247.

[30] P. Harrington, M. Gordon, A. Nauli, M. Paprzycki, S. Williams, J. Wright, Using Software Agents to Index Data in an E-Travel System, in: N. Callaos (ed.), *Electronic Proceedings of the 7th SCI Conferenc*e, Orlando, 2003, CD, file: 001428.pdf, 6 pages.

[31] http://www.fipa.org

[32] R. C. Erdur, O. Dikenelli, A Multi-Agent System

Infrastructure for Software Component Market-Place: An Ontological Perspective, *ACM SIGMOD*, 31(1), 2002, 55-60.