

Chapter 3

Application Impact of Multi-agent Systems and Technologies: A Survey

Jörg P. Müller and Klaus Fischer

Abstract While there is ample evidence that Multi-agent Systems and Technologies (MAS&T) are vigorous as a research area, it is unclear what practical application impact this research area has accomplished to date. In this chapter, we describe methods and results of a survey aiming at a comprehensive and up-to-date overview of deployed examples of MAS&T. We collected and analyzed 152 applications, covering important perspectives, such as ownership, maturity, vertical sectors, and usage of programming languages and agent platforms. We conclude that MAS&T have been successfully deployed in a significant number of applications, though mostly in what could be called niche markets. Off the spotlights of mass markets and current funding buzzwords, it appears that MAS&T is useful for various sectors.

Keywords Application impact • Applications of agents and multi-agent systems

1 Introduction

Since its inception in the 1980s, multi-agent systems and technologies (MAS&T) research has established itself as a recognized field within Computer Science, reaching out into other areas including economics, sociology, and psychology. With successful conferences such as AAMAS [3], IAT [17], MATES [23], and PAAMS [29], and with journals such as JAAMAS[19], AAI [1], AAIJ [2], and KER [20], there appears to be a stable community built around the questions

J.P. Müller (✉)

Department of Informatics, TU Clausthal, Clausthal-Zellerfeld, Germany
e-mail: joerg.mueller@tu-clausthal.de

K. Fischer

DFKI GmbH, Saarbrücken, Germany
e-mail: klaus.fischer@dfki.de

of understanding and constructing large-scale open decentralized systems that consist of autonomous components or systems, endowed with properties such as proactiveness, reactiveness, and the ability of flexible social action to achieve their design goals.

In the period from the mid-1990s until the early 2000s, the MAS&T research field went through (we could also say: benefitted, or suffered from...) a phase of hype, characterized by glossy conferences with heavy involvement from companies ranging from Pixar and Disney to Siemens, Daimler, Motorola, and British Telecom, and by ample public funding both in America (most notably the DARPA programmes in agent communication (leading to KQML) and CoABS (Control of Agent Based System)) and in Europe (exemplified by the AgentLink network¹) in Europe. In particular, AgentLink acted as an important dissemination channel toward industry, pushing the perception of *application impact* through the Agent Technology Conference series (held annually from 2002 until 2005) and through the AgentLink case studies featuring “real” applications of MAS&T.

However, since the mid-2000s, public perception of our research community appears to have become less prominent: Information and Communication Technologies (ICT) funding programmes were focusing on other labels, such as Service-Oriented Computing, Grid Computing, Autonomic Computing, or The Cloud; success stories in the software business, be it Apple, Facebook, Google, or SAP, have not been associated with MAS&T—at least not in the public perception.

Before this background, recently a perception among some MAS&T researchers appears to have formed—a perception that the field is lacking practical impact outside our own research community. The question driving this research activity has been to gain information and insights as to the extent to which this is true or not. We wish to study the application-oriented impact our research area has reached today.

There are some studies investigating agent applications, but we did not find any up-to-date work on the application *impact* of MAS&T, that is, the force of impression in the sense of being routinely and productively used in industrial, commercial, or public contexts. An obvious starting point for related work are the AgentLink case studies [4, 6]. While they did investigate a good selection of promising applications, they came too early to produce results related to impact—indeed, by the time of publication, the case studies were still prototypes or at the beginning of commercial use. After the end of AgentLink, the further impact of the applications described in the case studies was not systematically followed up.

In a study carried out in 2008 [9], Dignum and Dignum have collected and systematically analyzed agent applications. Their survey revealed a rather limited coverage in terms of replies, with very little industry participation. The focus of their study was on the characteristics of the applications rather than on their impact. [21] present a collection of industrial (manufacturing, logistics) applications; they do not aim at a systematic discussion or comparison of impact.

¹www.agentlink.org.

In a recent paper [5], Balke et al. analyze how implementations of software systems employing agent technology are represented in research-oriented publications, both at conferences (AAMAS, PAAMS, ICAART [18], MATES) and in journals (e.g., JAAMAS, AAIJ, KER). Their focus is on the impact of applications-related work in general on scientific publication venues. The focus is neither on practical application impact nor on specific applications. Their paper does not investigate the outreach of agent technology beyond the agents/AI community.

Looking at related studies done for other, similar studies, we came across an assessment project done by the Software Engineering research community: In [28], Osterweil et al. establish a scholarly assessment of the impact that research in software engineering has had so far on the practice in software engineering. In Sect. 2, we will discuss commonalities and differences of their approach and results when compared with our work.

In essence, the absence of systematic studies on the application impact of MAS&T has been the motivation for us to produce the work reported in this chapter. We describe the methods and results of a survey aiming at a comprehensive and up-to-date overview of deployed examples of MAS&T.

In Sect. 2, we define our notion of application impact and review approaches to describing and measuring it. Section 3 underlines method and settings of the survey. Results are described in Sect. 4. In Sect. 5, we provide an analysis of the most important vertical sectors (according to our survey) where MAS&T have developed impact. We discuss our findings and draw conclusions in Sect. 6.

2 Defining Application Impact

Merriam-Webster Online Dictionary [25] defines impact as *the force of impression of one thing on another*, or as *a significant or major effect*. In our work, we are interested in the application impact of MAS&T, that is, the force of impression/the significant effect that MAS&T have had on applications. Capturing this requires us to study related work on impact of ICT. Most literature definitions roughly classify impact into economic, social, and sometimes also environmental impact. *Economic impact* entails decreasing cost or increasing turnover/profits. *Social impact* includes aspects such as supporting human work to make it more satisfying and productive, changing the manner in which human users interact and cooperate, or making work environments safer or healthier. *Environmental impact* means decreasing environmental pollution or increasing sustainability of economic activities.

A prominent approach to defining and measuring ICT impact has been put forward by the Organization of Economic Cooperation and Development (OECD) in [27]; it has been adopted and extended by the United Nations Conference on Trade and Development (UNCTAD) in [24]. The model indicates the web of relationships between impact areas, and with the broader economy, society, and environment. Impacts of ICT arise through ICT supply and ICT demand and are likely to be influenced by the following factors (at a level of individual countries):

1. The existing ICT infrastructure which enables an ICT critical mass that can amplify impacts.
2. The country level of education, skills, and income, which influences both supply and demand.
3. The Government ICT policy and regulation, and the level of e-government.

With respect to measuring ICT impact, the OECD model identifies a number of interrelated segments, including (1) ICT demand (users and usage), (2) ICT supply (being the players of the ICT sector), (3) the level of/investment in the ICT infrastructure, (4) ICT products, information, and electronic content, and (5) ICT in a wider social and political context. The model proposes different types of impacts that address different (positive or negative) impact factors from the perspectives of these segments. These measurable factors appear to be very broad and diverse in terms of intensity, directness, scope, stage, timeframe, and characterization (e.g., economic/social/environmental impact, positive/negative impact, intended/unintended impact, subjective/objective impact).

While the OECD/UNCTAD approach can help us form an understanding of the general nature of ICT impact and its influence factors, we found it to be of limited use for addressing the particular problem of identifying and measuring the impact of the particular ICT research field of multi-agent systems. Firstly, OECD investigates impact by country, whereas we are interested in obtaining results involving a global but still relatively small research community. Secondly, while OECD can build on elaborated statistical data collected from countries and international bodies, no such statistics are available for MAS&T. Third, while (or because) the OECD model is very broad, it is not operational under the limited resources available for this study. Fourth (and maybe most importantly), OECD starts from sectoral and demand sides (e.g., studies the impact of ICT in healthcare, or the impact of ICT to specific user groups), whereas our interest is to measure the impact of a specific bag of models, methods, technologies and tools over a range of sectors and users.

Indeed, the application impact of a research area is hard to capture. While there is a considerable body of work on measuring effects and impact of science and technology [8, 14, 24], they are mostly either domain-specific (e.g., ICT impact for law enforcement [15]) or technology-specific (e.g., RFID technologies [32]). For most relevant MAS applications, a mixture of both domain-specific and technology-specific consideration is required: On the one hand, MAS have been applied in a large number of application domains. On the other hand, the notion of agents and MAS has been very broadly interpreted, making it difficult to subsume the applications under one technology-specific perspective. Also, when MAS researchers talk about impact, they often talk about two very different things: We (as a community) know our academic/scientific impact (measured in citation indices, scientific awards, or acquired research funds) quite well. However, what we know much less well is our *application-related* (i.e., economic, social, environmental) impact.

The aforementioned study by Osterweil et al. [28] (further referred to as *SE study*) investigated determining the impact of Software Engineering research on practice. This study was performed by leading players of the Software Engineering

research area in a funded project, involving different subgroups being responsible for different subareas (such as software configuration management, middleware, or programming languages). The SE study was performed in

the form of a series of studies and briefings, each involving literature searches and, where possible, personal interviews [28, p. 39].

While there are similarities between the SE study and our domain, there are some important differences, too. First, software engineering has a much longer history than MAS&T. A key finding of the above study has been that

[e]xperience, both in software engineering and diverse other disciplinary areas, has indicated that the impact of [...] prototypes might take 20 years to manifest [28, p. 41]

and that

[i]t typically takes at least 10 to 20 years for a good idea to move from initial research to final widespread practice.

Acknowledging that first concepts and prototypes featuring MAS&T are dating back 20 years or less (as opposed to almost 50 years in the case of software engineering), we must also acknowledge that our field is much less mature than software engineering. Second, the perception of the importance of software engineering to industry is very much different from the perception of MAS&T: Software engineering promises to address an urging problem faced by virtually every enterprise in the world, that is, to build robust, safe, efficient, scalable, and sustainable software systems. While we strongly argue that MAS&T can play a similarly important role in supporting future *software-intensive societies* by enabling cooperation, coordination, and evolution of large-scale mixed human-machine systems, our research community has so far been much less successful in attracting funding for a study like the one at hand. This leads to the third difference: Our study is a pure volunteer effort which fully relies on help from within the community. Hence, its scope is limited compared to the SE study. What we can hope for is to take a first step toward better understanding the current level of diffusion of MAS&T in practical applications.

Despite some above-mentioned valuable work done within our community, in studying the application impact of MAS&T we very much start on greenfield, especially concerning the work that has been done from 2004 onwards (i.e., after the end of the AgentLink network activity). While a reasonable number of application-related papers have been published in venues such as the AAMAS industry track, the PAAMS and IAT conferences, many of them are research prototypes, and so far, there have been no systematic efforts to monitor their development over time. Also, apparently there are companies that have been successfully building and improving their businesses using agent (or agent-like) technologies and systems. However, data on and insights into these applications [which are often “non-agent agent systems” (a marvelous phrase coined by Les Gasser in [34])] are often difficult to get from the owners of these applications, as also reported by Dignum and Dignum in their study [9]. These observations lead us to set up the comprehensive study described in this chapter. Our methodology is described in the following section.

3 Survey Methodology

As stated above, the overall objective of this activity is to gain information and insights of the application-oriented impact of MAS&T, and to provide a comprehensive and up-to-date overview of deployed MAS. To this end, we carried out a survey of deployed applications that use/are based on MAS&T, starting from year 2000 onwards. Rather than just providing a list of applications, our approach was to:

1. Classify the systems with respect to their maturity based on a set of indicators. As a baseline for our maturity classification, we use the NASA Technology Readiness Levels (TRL), which are a widely accepted standard [22]. We map the original set of nine TRL levels into the three categories: TRL1 to TRL4 correspond to Maturity Level C (lowest), TRL5 to TRL7 correspond to Maturity Level B, and TRL8/TRL9 correspond to Maturity Level A. We refer to Sect. 4.2 for further information.
2. Provide an at least qualitative characterization of the application impact based on a set of criteria, and identify particularly high-impact application.
3. In particular, follow up the development and impact of previously published application-oriented work (e.g., the AgentLink case studies as well as work presented in the AAMAS industry tracks).

To achieve these goals, we pursued the following activities:

1. An open call for nominations of deployed MAS&T using a web-based online system. This aimed at academic and industrial members of the broad MAS community.
2. A mail-based survey directed toward the authors of papers presented at the AAMAS2005 to AAMAS2012 Industry/Innovative Applications.
3. Direct/personal mails directed to dedicated (industrial but also key academic) players, which would be unlikely to respond to 1 or 2.

In the course of the study, we have been collecting different sets of information: In the first round (web-based survey), we were asking researchers to nominate real-world applications that were deployed in the year 2000 or later, in a corporate, administrative, or public environment. In particular, we were requiring that these applications should have considerably and positively contributed to corporate or administrative value creation, to public/social welfare, or to application-oriented grand technology challenges (such as e.g. RoboCup). To restrict the survey to MAS&T, we further imposed the requirement that the application uses research results (models, methods, architecture, algorithms, technologies and platforms, tools) in the realm of MAS&T at its core—no matter whether the label of MAS&T is actually used or not. In parallel, we carried out a literature research to identify prospective deployed applications based on work published in the AAMAS Industry/Innovative Applications Tracks 2005–2012. Metadata and deployment information was collected from these papers and a consolidated list of candidate applications was created containing the applications from the survey plus the

applications identified from the AAMAS industry track papers. In the second round, a fact sheet form was created and mailed to the developer/owner of each of the applications gathered previously. The objective of this second round was to validate the deployment status based on first-hand, up-to-date information, and to obtain a common level of information for comparing and uniformly presenting the results of the survey, regarding development/deployment timeline, resources spent, and benefits achieved for the different applications.

4 Survey Results

Based on the goals and method of the survey laid out in the previous section, this section reports our results. Advertised using the major agent-related mailing lists, the online survey was performed from July to early October 2012. People were invited to propose either own applications or nominate applications they know about and give a contact person for reference. One hundred and three applications were nominated using the online survey. In parallel, 99 applications were identified as the result of a literature research in the proceedings of the AAMAS Industry and Innovative Applications tracks from 2005 to 2012. The two sets were merged, duplicates and irrelevant entries (e.g., work finished before year 2000, just overview or white paper but no application) were removed. The result was a list of 152 applications, which form the basis of the results presented in this chapter.

For each of the 152 applications, we identified a contact person we approached in order to obtain additional information about the applications. We did so by designing a simple fact sheet template, which we asked the contact person to fill in. We received factsheet information for 89 applications, corresponding to 59 % of all applications. While this appears to be an excellent return, the completeness of input to the different factsheet questions varies. For instance, while over 80 % of the returned factsheet provide information regarding the agent platform used, only 55 % provide information regarding the development resources. Thus, while the fact sheets have fulfilled the purpose to verify the deployment status (maturity) of the applications, they have only to a smaller extent allowed us to gather information about timelines, resources, usage numbers, and quantitative benefits (such as revenue) created through agents and multi-agent systems.

4.1 *Distribution of Applications Across Partner Types*

Based on information collected from the survey, the fact sheets, and additional resources (publications and web pages related to the applications), we classified the applications according to the partner characteristics, making a distinction between applications developed (and owned) by industrial or governmental organizations, applications that were developed in industry-academic cooperations, and applications developed/owned solely by academic partners. Forty-seven applications

(corresponding to 31 %) were exclusively developed and owned by industrial or governmental players, whereas 43 applications (28 %) were built by academic partners, and 62 applications (41 %) were created in industry-academic cooperation. Thus, academic partners were involved in 69 % of the applications. This ratio can be partly explained by the high level of participation of academics compared to industrials in the survey: 61 % of the applications were proposed by academic participants, in 39 % of the cases industrial or governmental players were proposers or otherwise involved in providing the information.

An interesting question is whether there is a correlation between the developer/owner category and the maturity of the applications. One would expect that in general, applications developed by academic partners have lower maturity than applications developed by industrial or governmental organizations. But how about industry-academia co-productions? For this purpose, in the next subsection, we consider the maturity of applications.

4.2 *Maturity of Applications*

Based on information collected from the survey, the fact sheets, and additional resources (publications and web pages related to the applications), we classified the 152 applications in the following three maturity classes (see also Sect. 3):

- Systems that are or have been in operational use in a commercial or public environment (Maturity Level A, corresponds to TRL 8 to 9)
- Industry validated research prototypes (i.e., prototypes that are being validated/piloted in an industrial or public environment with online industrial data under live conditions) (Maturity Level B, corresponds to TRL 5 to 7)
- Research prototypes validated with offline real-world data or in an offline environment (Maturity Level C, corresponds to TRL 1 to 4)

A fourth category contains systems or activities that are not applications in a strict sense, but that have (or have had) some indirect impact via technology challenges, benchmarking activities, or standardization efforts (such as e.g., IEEE-FIPA). Figure 3.1 shows the distribution of the applications in the survey according to their maturity levels. Forty-six applications (31 %) out of those classified as A, B, or C are (or have been) in operational industrial or public deployment. Further 55 applications (37 %) have been validated/piloted in an industrial or public environment with online industrial data under live conditions, whereas other 46 (31 %) are research prototypes that were never (or not yet) deployed. The latter category mainly comprises applications that were described in AAMAS Industry Track papers. We decided to include them in our survey but clearly mark them with respect to their maturity. Using the fact sheets, we tried to confirm the maturity status of all applications in the survey with their developers or owners. We were able to confirm 76 % of maturity class A applications, 59 % of class B, and 44 % of class C. In the remaining cases where no confirmed information about the maturity

Fig. 3.1 Maturity levels of applications


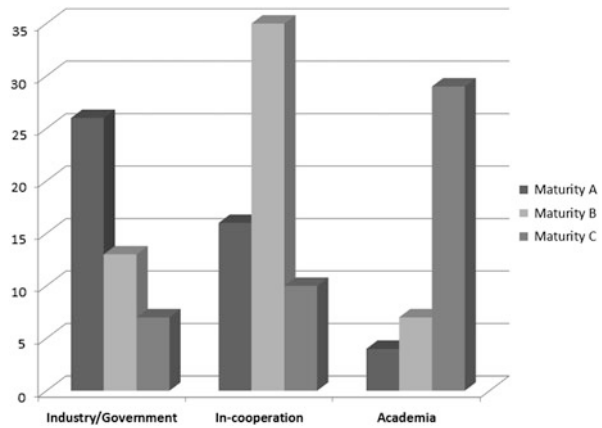


Fig. 3.2 Maturity levels of applications by partner types (y-axis shows absolute numbers of applications)




status was available, we perform the classification based on available information such as publications, product/project websites, and personal communication.

Next let us revisit the correlation between partner types (Sect. 4.1) and maturity of the applications. Figure 3.2 shows the maturity levels of the applications grouped by different partner types. Not surprisingly, applications owned and developed by industrial and governmental players have a considerably higher maturity (26 maturity class A applications, and only 7 maturity class C applications) than applications developed by academics (only 4 maturity A, but 29 maturity C). An interesting result is, however, that applications developed in cooperations of academic institutions with industry or public bodies are performing remarkably well in terms of maturity. This result goes in line with the observation made by [28, p. 41], for software engineering research that

[t]echnology transition is most effective and best expedited when research and commercialization maintain a close synergy over an extended period.

It will be interesting to see how many of the in-cooperation applications currently in maturity level B will ultimately migrate to level A. Our subjective impression

Fig. 3.3 Applications by agent system types (figures are absolute numbers)



based on feedback from the fact sheets is that there are quite a few new and emerging applications “in the pipeline.”

4.3 Agent System Types

It is not only since Franklin and Graesser [12] that we are aware of the heterogeneity of the notion of agents and its interpretation. It would be surprising if a survey on the impact of MAS&T did not reflect this heterogeneity. We have classified the applications in our survey into three categories according to the most well known system types: (1) multi-agent systems focusing on interaction, cooperation, and coordination; (2) intelligent agents focusing on single-agent aspects such as planning or learning; and (3) personal/UI agent focusing on agent-human interaction and assistance. Figure 3.3 shows the distribution of the applications considered in the survey. The large number of applications in the multi-agent systems category certainly reflects the focus towards multi-agent topics in the call for participation rather than a lack of intelligent agent or personal/UI agent applications. Also note that the three categories considered are, while being helpful, not orthogonal and of limited discriminatory power: In many multi-agent systems, single-agent local aspects play an important role. Also, human-agent interaction can be viewed as multi-agent interaction as well depending on the perspective. Also, UI agents should often reveal intelligent behavior. Yet, in most cases, some focus can be observed, which is why we decided to keep these three categories.

4.4 Applications by Country

Next, we consider the distribution of the creators of the 152 applications covered in the survey by country. As a general rule, in the case where an agent’s company




Fig. 3.4 Applications by countries (absolute numbers, total and maturity A only)

located in country A has created an application for a customer located in country B, we allocate this application to country A. In case of a company with multiple locations we use the country of the responsible location in case we know it (e.g., in the case of IBM, two applications were collected from the Haifa Lab, so they count for Israel); otherwise, we count the application for the country where the company headquarter is located. Figure 3.4 illustrates the distribution of applications by the countries of their creators. The 152 applications covered by the survey were created by parties from 21 countries. The USA is by far the country creating the largest number (41, corresponding to 27%) of MAS&T applications, followed by the UK, Germany, the Netherlands, the Czech Republic, and Australia. Also when considering Maturity A applications only, USA takes the clear lead (12, corresponding to 26%), with runners-up being the UK, Spain, Germany, Italy, and the Czech Republic. The strong presence of Spain and Italy for highly mature applications is due to the strong industrial players Telefonica T+D and Telecom Italia.

4.5 Applications by Vertical Sectors

The pie chart in Fig. 3.5 shows the distribution of applications across vertical sectors. Within the 152 applications, 22 sectors are represented. It is noticeable here that eleven sectors cover 86% of all applications, whereas the top six sectors (logistics and manufacturing, aerospace, energy, defense, security and surveillance, and telecommunications) still represent 59% of the applications. The picture changes considerably if we do not only consider the number but also the maturity of the applications in the different sectors. Figure 3.6 illustrates the number of

Fig. 3.5 Number of applications by vertical sectors


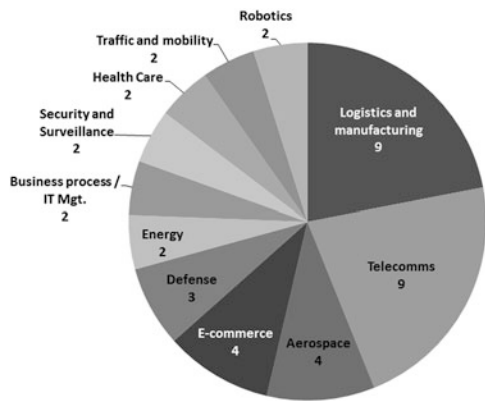



Fig. 3.6 Number of applications by vertical sectors (maturity class A only)



applications with maturity level A by vertical sector. It only displays the sectors for which five or more applications have been recorded. From this view, we see that logistics and manufacturing and telecommunications are the domains with the overall highest number of mature agent applications, whereas energy, security and surveillance, and defense appear to be emergent applications with yet little impact. However, especially in the defense domain, it is well possible that confidentiality issues may distort the picture—we may just not be aware of some successful applications of MAS&T in this domain.

A final consideration reflects on the vertical sectors with a particular high percentage of high-maturity applications. For this, we consider again the vertical sectors for which at least five applications were listed and calculate the percentage of maturity class A applications amongst all applications recorded for this sector. The results are illustrated in Fig. 3.7. We observe that the Telecommunications

Fig. 3.7 Vertical sectors with high-maturity applications



sector is very mature in terms of agent-based solutions, reflecting the historical development with early involvement of telecommunications companies, such as Telecom Italia, Telefonica I+D, British Telecom, Siemens, and Motorola. Logistics and manufacturing, e-Commerce, and robotics follow with 50% to 40% share of maturity A applications. At the lower end of the spectrum in terms of relative maturity, security and surveillance as well as energy sectors feature a large number of applications, but most of them are (still?) of low maturity. Note that the figure for the defense sector are associated with a high degree of uncertainty due to the confidentiality requirements in this domain.

In Sect. 5, we shall review and discuss in more detail the most prominent vertical domains emerging from our survey.

4.6 Programming Languages and Agent Platforms

Seventy-three out of 87 submitted factsheets provide information about the programming languages used in application development. Since we were particularly interested in the usage of dedicated agent platforms and tools, we asked for that information separately. Seventy-five factsheets provide information about agent platforms and tools used (including the rather frequent answer “None”). Java has been by far the most popular programming language, used in 53 applications. It is followed by C/C++/C# (used in 15 applications, including but not restricted to embedded or real-time applications), PHP (seven applications), and Python (four applications). These four groups were used in 75% of the applications for which information was available to us. Note that some applications have used more than one programming language. Figure 3.8 illustrates the coverage of dedicated agent platforms in the applications.

Fig. 3.8 Usage of agent platforms in applications (figures are absolute numbers)


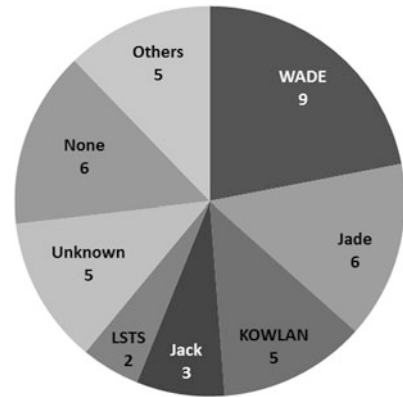


Fig. 3.9 Agent platforms used for highly mature applications (figures are absolute numbers)



We can draw a couple of observations from this chart. First, a large majority of applications (24, corresponding to 32 %) has not used any dedicated agent platform or tool. Second, the most commonly used platforms are JADE (13 applications, $\approx 15\%$), AOS's Jack, CoJack, and C-BDI product family (12, $\approx 16\%$) as well as WADE (11, $\approx 14\%$). Taking into account that WADE is a JADE extension, JADE can be regarded the overall most-used agent platform. These are followed in respectful distance by KOWLAN, and Whitestein's Living Systems platforms (LSTS and LSPS). Third, there is a high fragmentation in the agent platforms landscape in that 20 different platforms and tools were used in a single application only. This fragmentation was already observed in the study by [9].

At the end of this subsection, we shall investigate the question whether the distribution of agent platforms shown in Fig. 3.8 will change if we only consider applications with maturity level A, that is, in operational use. Figure 3.9 shows the absolute usage numbers of agent platforms in these applications. This information was extracted from 30 fact sheets submitted for maturity level A applications, which have provided information about the agent platforms used. What is striking when comparing it to the numbers over all applications considered in the survey

is that the ratio of applications that do NOT use dedicated agent platforms is significantly lower in the case of maturity A applications: only six out of 30 maturity A applications that provide information about agent platforms have NOT used an agent platform, compared to 24 out of 75 applications in total. On the one hand, this reflects the strong role of players such as Telecom Italia, Telefonica I+D, AOS, and Whitestein, who have been applying their agent platforms (JADE/WADE, KOWLAN, JACK/CoJACK, LSTS/LSPS) to build a number of successful deployed applications. On the other hand, we might conclude from this that dedicated agent platforms actually can make a difference regarding business success.

A further interesting observation in this context is that the agent platform landscape is much less fragmented for highly mature applications: While 20 out of 75 applications in our survey are based on proprietary platforms which are not used by any other application, in the case of maturity A applications, only five out of 30 applications are based on “singular” platforms.

For maturity A applications, WADE and JADE are again the most frequently used platforms, followed by KOWLAN, Jack, and the Living Systems Technology Suite (LSTS). Apparently, mature applications, which often have a longer development history, tend to be based on mature platforms; more recent platforms such as CoJack, C-BDI, or Whitesteins Living Systems Process Suite may take some more time to mature with the applications constructed with them.

4.7 AgentLink Case Studies Revisited

Running from 1998 until 2006, The European AgentLink Coordination Action for Agent-Based Computing has gathered important application-oriented work in its case studies. Elaborated in 2004/2005, eight prominent MAS&T applications were described and analyzed [6]. The case studies are still available on the AgentLink website [4]. Eight years after the case studies were written, we have reexamined these eight applications, trying to gather up-to-date information with respect to their further development. Table 3.1 summarizes the key information we were able to obtain. The table provides the maturity level (see Sect. 4.2 above) reached by the systems described in the case studies as well as the information whether the systems are still maintained.

It reveals that, according to our findings, four out of eight case studies (the ones by AOS, Whitestein, Magenta, and Almende) are still maintained and in operational use (maturity A). Three are confirmed to be no longer used, of which only one (EuroBIOS / SCA) had reached maturity A status at some point in time. Combined Systems was a research project, the use case was not taken up after its end. Also, the Chilled Water System Automation case study (Rockwell, CTU Prague) developed a research prototype and was not commercialized. For one case study (Acklin B.V.), no information was available to us.

Beyond the operational status (and connected revenue generation for their owners) reached by specific AgentLink case studies, an important positive effect we can

Table 3.1 What happened to the AgentLink case studies?

App id	Title	Partners	Maturity level	Still maintained?
64	Living Systems/Adaptive Transportation Networks	Whitestein Technologies	A	Yes
69	HV-CGF: Intelligent Human Variability in Computer Generated Forces	AOS, UK Ministry of Defense	A	Yes
71	Agent-based Factory Modelling	EuroBios, SCA Packaging	A	No
73	Intelligent Scheduling of Cargo Fleets	Magenta Technology, Tankers International	A	Yes
75	Software Agents for International vehicle Traffic Insurance	Acklin B.V.	No info	No info
84	Chilled Water System Automation	Rockwell, CTU Prague	B	No
87	Combined Systems	D-CIS, Thales	C	No
88	Agents for Intelligent Communications Systems / Self-organizing systems	Almende, ASK CS	A	Yes

observe from the AgentLink case studies is that these applications have sustainably fertilized products and businesses of the companies involved. For example, the main result of the HV-CGF project was the CoJack reactive architecture, which AOS has been using in further deployed applications. As another example, Whitestein Technologies has not only created additional business in the area of logistics based on the LS/ATN reference application, it has also used the experience with LS/ATN to develop development and execution platforms (in particular, the LSTS and the Living Systems Process Suite, which have been the basis for generating additional business. Also, even if the Chilled Water System Automation case study was not commercialized, results from that project have initiated and driven further successful applications reported by the CTU team.

In summary, we regard the fact that 8–10 years after publication, four out of eight applications are still operational and (apparently) thriving, a positive rather than negative news, given that according to [13], from ten venture-backed startups, three to four fail completely and only two produce substantial returns. To us, this demonstrates that beyond past hype and disillusionment, successful and sustainable businesses can be built on the grounds of MAS&T.

4.8 Agent Companies

A major indication of impact of a specific technology is the number of companies which successfully build business from selling products or services based on this technology. In the following, we give examples of companies whose business is known to build on agent technologies. Our list is exemplary and by no means meant to be exhaustive. It is difficult (and was not in the scope of this survey) to find out details of the business models of individual companies.

We start with a list of companies which successfully managed to establish themselves in the market²:

Whitestein Technologies³ offers agent-based solutions for business process management and execution in the areas of financial services, manufacturing, telecommunications, or logistics. Additionally Whitestein offers solutions for logistics management, optimization, and control.

Agent-Oriented Software (AOS)⁴ claims to be the leading company for providing autonomous and semi-autonomous systems. AOS provides platforms and development tools for the design of agent-based systems (most noticeably Jack and CoJack, the latter of which was a result of the HV-CGF project reported as an AgentLink case study, see above) as well as solutions for dedicated application domains like for example assisting surveillance and intelligence agencies as well as for Oil and Gas industry.

Real Thing⁵ is a rather young company. It is not purely specialized on agent technologies but also offers Apps for smart phones. However, toy robots for kids are clearly agent-related. From the information available in open source, it is difficult to decide which concrete technologies the products of this company build on.

In addition to the above-listed success stories, there are also unavoidably examples of startups who did not manage to establish themselves in the market:

xaitment was a startup founded by members of the multi-agent system group at the German Research Center for Artificial Intelligence (DFKI) GmbH. The company specialized on middleware for the development of computer games. xaitment recently merged with iOPENER.⁶

Agentis Software was a startup founded by members of the Australian Artificial Intelligence Institute (AII). The main objective of Agentis Software was to apply the concepts of BDI agents to business process management and execution.

²Note we have not included players such as Telecom Italia or Telefonica I+D in this list because, while they are using agent technology in their operations, their business does not build on it.

³<http://www.whitestein.com/>.

⁴<http://aosgrp.com/>.

⁵<http://aosgrp.com/>.

⁶<http://www.iopenermedia.com/>.

Nevertheless, the positive examples are a clear indication that it is possible to successfully build business models around agent technologies.

5 Vertical Sector Analysis

In Sect. 4.5, we have identified the vertical domains that, following the results of our survey, appear to be most relevant for agent-related research and application development. In this section, we provide a brief characterization of these sectors.

Aerospace is a very diverse application area with a large number of applications in commerce, industry, and military. Early work in distributed artificial intelligence (DAI) investigated the use of agent and multi-agent system technology for robots designed for exploring the surface of remote planets. The majority of applications collected in our survey investigates agent technologies for unmanned aerial vehicles (UAVs). Additional applications are involved with the support of pilots in military situations.

Defense To identify research on military applications is not really easy because researchers working on such applications are at least sometimes not allowed to publish their results with a direct pointer to the military context. Sometimes a completely artificial application domain is chosen to obfuscate the real application. Nevertheless, our study could identify a significant body of work in this area. An obvious application of agent technologies in a military application is the simulation of a human engaged in a military mission. Simulation of human to human and human to environment interaction are investigated. Simulation of unmanned vehicles or other intelligent and autonomous devices, as well as the use of game-theoretical models for decision support are other relevant MAS&T topics within deployed defense applications.

E-commerce Although online shopping was already invented in the late 1970s, it became in broader use only after the advent of the World Wide Web and its commercialization in the late 1990s. Online shopping forms a major part of e-commerce but any kind of business interactions using the Internet falls into this application area. With this broad scope it has a significant overlap with supply chain management and manufacturing & logistics. The settings e-commerce are inherently multiparty and geographically distributed, see [26].

Energy is a vibrant sector following an important societal theme, which has been providing ample funding opportunities for research over the past few years. Unsurprisingly, it has recently become attractive for agent-related research. With the global change to produce energy (especially electricity) from fossil or nuclear sources to sustainable sources, the control of electricity networks became an even more demanding task than it already used to be. It is very likely that in the near future completely decentralized control mechanisms need to be put in place because individual households are likely to be at the same time source and drain of flow of electricity. Even small electronic devices will get into the position to buy and sell energy when it is economical in a given

market situation. It is therefore not astonishing that agent research tries to adopt well-understood negotiation and market mechanisms for this application domain. Other applications support the retrieval of energy sources [e.g. iWDTM (intelligent Watchdog)].

Health Care Similar to energy, health care is an application area with great importance especially when the demographic change in the western world is taken into account. Related to health care is the application area of ambient assisted living (AAL). With the increase in networking of devices on a wireless basis in the general public or in the users' homes, there is a huge potential of innovative applications for agent technologies in this area. The applications include monitoring of a person's health status, selection of candidates for transplant surgeries, resource management in hospitals, or the supervision and support of people with health problems in their homes.

Manufacturing & Logistics has been an interesting application area for agent-related research from the very beginning in the early to mid-1980s DAI research. Already in the first DAI book by Huhns [16], manufacturing was listed as a major application domain by Parunak [30]. In logistics, early work on transportation scheduling was reported by Sandholm [31] and Fischer and colleagues [11]. Indeed, references collected for this survey reach over the whole period back till 2000, which we used as a cut of year for collecting data. The different applications in the survey collection cover a broad spectrum of topics. Production planning and control, task allocation, product memories, negotiation, and simulation were topics the applications in the survey collection dealt with. Commercialization driven by companies such as Whitestein Technologies and Magenta Technologies are active in this application area.

Robotics A relatively small but notable and successful part of the applications in our survey deal with robotics, and, in particular, multirobot systems. Robotics has always been a natural application area for AI, and such have multirobot systems been a natural application domain for MAS&T. An autonomous robot is a prototype of an autonomous agent which has a physical body. At the same time, coordination and cooperation processes in multirobot systems can be efficiently modeled and implemented using MAS&T, as examples such as Kiva or CogniTAO on the commercial side, and Robocup [33] on the research side show. Robotics itself is a huge application area in which other disciplines (especially engineering) meet with research on pure agent technologies. The Autonomous Agents (AA) conference in 1997 was the first international conference where the two research areas of pure agent technologies and physical robots met in a major international scientific event. Later on the AA, ATAL, and ICMAS conference joined forces to form the AAMAS conference as we know it today.

Security and Surveillance Security is important for basically every application domains. The Internet makes the need for security more than clear to all participants. Agent research offers very interesting settings in which theoretical solutions can be deployed and prove their strength. Surveillance puts the idea of security to another level. In our networked society where wireless networks spread out in an extremely fast manner, surveillance is getting more and more

widely used which in some cases does increase security but also raises issues with respect to privacy. Because the application domain is naturally distributed, it is an ideal setting for the application agent technologies.

Telecommunications Telecommunication companies have been interested and involved in agent-related research from an early stage on. Although the absolute number of applications in our survey was not that high, as already noted, the maturity of the applications has been outstanding (see Fig. 3.7), notably driven by companies such as Telecom Italia and Telefonica I+D, which continue to be active innovators. In present days, where smart phones take over mobile telephony market, mobile devices are ubiquitous which can easily run a MAS application and interact with similar applications running on other devices. It is therefore very likely that there will be a boost of such applications in the near future.

6 Discussion, Conclusion, and Outlook

When we started out with preparations for this survey in Spring 2012, we did so with a certain degree of skepticism and with rather modest expectations. Our (subjective) observation of the multi-agent systems research field as that it was fairly healthy as a research field, but its track record in terms of application impact did not seem to be prominent. Industry participation at conferences (most notably AAMAS) had been considerably decreasing over the past years (see also the analysis by [5]), project funding for research performed under the MAS&T label has become more difficult to obtain, and agent companies and products are rarely a topic in daily technology news. The agent application survey done by Dignum and Dignum in 2008 and published in 2010 [9] enforced our skepticism. They noted that they

were surprised by the small number of responses and by the dominance of academic respondents ([9, p. 231])

and conjectured that

[a]lthough the reasons may be partially related to the announcement medium (the agents mailing lists are mostly used within the academic world), this small number may be an indication that there are indeed not many real applications of MAS around.

Almost 1 year later, at the time of writing this concluding section, we feel we can be somewhat more confident and more optimistic regarding the current and future impact of our research field. Supported by a large number of agent researchers and practitioners (see Acknowledgements), we identified and analyzed 152 applications of MAS&T, out of which 31 % are deployed applications, while additional 38 % were validated/piloted in industrial or public environments with online industrial data under live conditions. So indeed, there are a considerable number of “real applications of MAS around.” Also, an investigation of the destinies of the applications known as the *AgentLink case studies* revealed that half of these are still operational 8–10 years after their inception.

Looking closer at the results of this survey, we can summarize our main conclusions. The main players bringing about significant successful deployed applications over a longer period in time, appear to be (more or less surprisingly) Telecom Italia, Telefonica I+D, Agent-Oriented Software and Whitestein Technologies. On the academic side, various research groups, for example, at CMU (Sandholm), Czech Technical University (Pěchouček), DFKI and, most recently, USC (Tambe) have repeatedly and successfully crossed the prototype-to-deployed-application chasm.

Looking at the main industry players that have been traditionally associated with agent technology (beyond those already mentioned above), some of them seem to have disappeared from the multi-agent business (e.g., Siemens, Motorola); some others market their “agency” solutions under different labels (IBM, Daimler, NASA, Google), some (e.g., British Telecom) keep on developing their successful agent technologies in the rather small scale, some are developing agent technologies for their business using prototypes (e.g., Aerogility, Thales).

Yet, MAS&T seem to thrive best in what we could call niche markets. For instance:

- Multi-agent architecture and distributed management have been successfully applied in telecommunications network management.
- Flexible control, scheduling, planning, and optimization solutions have been successfully applied in manufacturing and logistics.
- Agent-based simulation has become a respectable and respected microsimulation approach for modeling large-scale systems consisting of autonomous entities, so far in specific domains, including crowd, pedestrian, and traffic simulation.
- Applied game theory has grown into an very attractive application area, in particular for security, surveillance, and defense applications.
- Very interesting work is being done in (multi-)robotics by relatively small players (e.g., Kiva Systems, Cognitao).

So one could argue that we have not seen success stories of MAS&T—neither in the large mass markets (such as consumer products) nor in the societal priority areas such as energy and health care yet. Trying to contradict to this argument, in the survey we noted a substantial number of research prototypes in these areas (in particular: energy). As the field matures, many of them may turn into deployed applications.

A second observation we made regarding success stories while doing this survey is somewhat anecdotal. In fact, two applications were proposed, which at first sight beyond doubt qualify for the *success story* predicate: One is the use of proxy bidding agents [10] in Google’s AdWords product, which is allegedly Google’s main source of income.⁷ This work has been nominated by numerous researchers for consideration in the survey. In his nomination, David Parkes wrote:

Google (and other search engines) use a multi-agent architecture to provide automated bidding for their advertisers. An advertiser expresses a high level goal (e.g., maximize my number of clicks without spending more than US\$1,000 a day) and they try to meet that

⁷According to <http://en.wikipedia.org/wiki/AdWords>.

on behalf of the advertiser using an agent that represents the advertiser. This is one of my favorite examples of MAS thinking and a truly agent-based market system.

The second relevant application has been PTIME [7], an application developed by SRI as part of the DARPA-funded CALO research project, an effort to build an adaptive cognitive assistant situated in the office environment. A PTIME agent is an autonomous entity that works with its user, other PTIME agents, and other users, to schedule meetings and commitments in its user's calendar. Indeed, some results developed in CALO were acquired by Apple and formed the technological basis for the Siri assistant today available in Apple mobile phones.

What both high-impact applications seem to have in common is that the owners of these applications do not seem to consider them as applications of MAS&T: Despite numerous attempts, we could not obtain a response from the responsible technical people at Google. When asked about the deployment of results from CALO, we received an email response from a senior scientist at SRI stating that

The [CALO] system as a whole was never deployed externally although several parts of it were deployed in fielded government systems or used as seed technology for startup companies. In none of these cases, however, would I describe the deployed components as multi-agent systems.

In summary, however, we can state that off the spotlights of ICT wonderland, MAS&T has been successfully used in a significant number of applications, and continues to be an increasingly useful and impacting technology in various sectors. Yet, there is no reason for over-enthusiasm: Coming back to the comparison to the Software Engineering study already discussed in Sect. 2, one finding of that study has been that

[c]ontinued support for sustaining a vigorous research community is required ([28, p. 45]).

In our research community, there seems to be selective, but no broad continuous support (public or industrial) over the past few years (in Europe, there has not been much after AgentLink), which is definitely problematic.

In this chapter, we have provided the main results of the impact survey, covering some important perspectives, such as maturity, vertical sectors, and usage of programming languages and platforms. For other aspects, such as the analysis of the system complexity, development effort, timescale, and economic performance of MAS&T, our current data basis for now is insufficient to derive significant results. Therefore, we have not included these aspects in this paper. In future work, we shall attempt to complete our data collection with respect to these aspects, to be able to carry out further analyses.

Additional Information

The following table (Table 3.2) lists the 46 applications contained in the survey, which were classified as maturity level A (TRL 8/9).

Table 3.2 List of maturity A applications contained in the survey

App id	Application name	Sector	Owner/developer
1	Generalized combinatorial multi-attribute auctions/CombineNet for Sourcing	Logistics and manufacturing	CombineNet, CMU, US
2	Live-donor (US-wide) kidney exchange	Health Care	T. Sandholm, CMU, US
3	Poker	Entertainment	T. Sandholm, CMU, US
5	cdmNet	Health Care	Precedence Healthcare, AU
10	GlacsWeb	Geosciences	Southampton University, UK
15	Global package tracking, tracing, recovery	Logistics and manufacturing	DHL, Agentis, US
20	Debatescape	E-commerce	British Telecom, UK
21	Kiva systems	Robotics	Kiva Systems, Peter Wurman, US
23	Proxy bidding agents at Google	E-commerce	Google, US
26	PROTECT	Security and Surveillance	USC Teamcore, US
27	ARMOR	Security and Surveillance	USC Teamcore, US
28	IRIS	Defense	USC Teamcore, US
36	CAST Terminal	Aerospace	Airport Research Center GmbH, DE
38	Catalogue manager and price checker/setter	E-commerce	The Book Depository, UK
51	Wizard	Business process/IT Management	Telecom Italia S.p.A., IT
52	WANTS-Delivery (aka Network Neutral Element Manager)	Telecommunications	Telecom Italia S.p.A., IT
53	WANTS-Assurance	Telecommunications	Telecom Italia S.p.A., IT
54	WeFlow	Telecommunications	Telecom Italia S.p.A., IT
56	Legion	Traffic and mobility	Legion Ltd., UK
57	Steps	Traffic and mobility	Mott McDonald, UK
64	Living Systems/Adaptive Transportation Networks	Logistics and manufacturing	Whitstein Technologies, CH
65	LS/AMC	Logistics and manufacturing	Whitstein Technologies, CH
67	MasDISPO_xt	Logistics and manufacturing	Saarstahl AG, DFKI, DE

(continued)

Table 3.2 (continued)

App id	Application name	Sector	Owner/developer
69	HV-CGF: Intelligent Human Variability in Computer Generated Forces	Defense	Agent-Oriented Software Ltd., UK MoD, AU
71	Agent-based Factory Modelling	Logistics and manufacturing	EuroBIOS, SCA Packaging, SE
73	Intelligent Scheduling of Cargo Fleets	Logistics and manufacturing	Magenta Technology, Tankers International, RU
77	AgentFly	Aerospace	AgentFly Technologies and Agent Technology Center, Czech Technical University, CZ
80	ExPlanTech PPS system	Logistics and manufacturing	Modelarna Liaz, SkodaAuto, CZ
82	Ad-hoc networking in disruptive environments	Defense	CTU Prague, CZ
88	Agents for Intelligent Communications Systems/Self-organizing systems	Telecommunications	Almende, ASK CS, NL
96	DHS Control	Energy	NODA Intelligent Systems AB, SE
99	SUPREMA	E-government	Knowledge Genesis, RU
102	MAS-Dispo	Logistics and manufacturing	Saarstahl AG/DFKI, DE
112	KOWLAN MACROLAN	Telecommunications	Telefónica España (MACROLAN), ES
115	CORMAS	Simulation	Francois Bousquet, FR
116	INNSIST	E-commerce	Grupo TCA. Monterrey, MX
123	ASE (Autonomous Sciencecraft Experiment)	Aerospace	NASA, US
126	OCA Management System (OCAMS)	Aerospace	The Work Systems Design & Evaluation group at NASA Ames Research Center, US
131	CogniTAO (Think As One)	Robotics	Cogniteam, Ltd., IL
133	EV2G (Electric Vehicles to Grid)	Energy	Willett Kempton, University of Delaware, US

(continued)

Table 3.2 (continued)

App id	Application name	Sector	Owner/developer
139	Living Systems Process Suite (LSPS)	Business process/IT Management	Whitestein AG, CH
142	Tacsim	Defense	AOS, Australian Defence Department, AU
150	KOWLAN CZ IP Connect	Telecommunications	Telefónica España (MACROLAN), ES
151	KOWLAN Digital probes	Telecommunications	Telefónica España, ES
152	KOWLAN Iberbanda	Telecommunications	Telefónica España, ES
153	KOWLAN BA	Telecommunications	Telefónica España, ES

Acknowledgements Numerous people, agent researchers and practitioners alike, have contributed to this work by participating in the online survey, thus pointing us to a large number of interesting applications, by providing information about applications they have been involved in, by email and in some cases via phone interviews. We are indebted to Amal El Fallah Seghrouchni, Ana Bazzan, Ladislav Bölöni, Francois Bousquet, Jeff Bradshaw, Jean-Pierre Briot, Birgit Burmeister, Stefan Bussmann, Giovanni Caire, Jiri Chabera, David Clarke, Juan Manuel Corchado, Paul Davidsson, Keith Decker, Frank Dignum, Virginia Dignum, Klaus Dorer, Emad Eldeen Elakehal, Yehuda Elmaliyah, Javier Garcia Algarra, Mike Georgeff, Maria Gini, Andrey Glaschenko, Jorge J. Gomez-Sanz, Amineh Gorbani, Danilo Gotta, Dominic Greenwood, Hiromitsu Hattori, Axel Hessler, Koen Hindriks, Benjamin Hirsch, Tom Holvoet, Michael Huhns, Nick Jennings, Catholijn Jonker, Olaf Junker, Gal Kaminka, Jan Keiser, Victor Lesser, Renato Levy, Andrew Lucas, Mike Luck, Philippe Mathieu, Peter McBurney, Felipe Meneguzzi, Tim Miller, Tsunenori Mine, Anabel Montero, Prabhu Natarajan, Bernhard Nebel, Pablo Noriega, Ingrid Nunes, James Odell, Steve Osborn, Julian Padget, Lin Padgham, David Parkes, Michal Pěchouček, Raymont Perrault, Valentin Robu, Alex Rogers, Tuomas Sandholm, Paul Scerri, Glenn Semmel, Alexei Sharpanskykh, Onn Shehory, Maarten Sierhuis, David Sislak, Petr Skobelev, Nikolaos Spanoudakis, Patrick Storms, Milind Tambe, Simon Thompson, Pavel Tichy, Luca Trione, Wim van Betsbrugge, M. Birna van Riemsdijk, Ondrej Vanek, Giuseppe Vizzari, Pavel Vrba, Michael Wellman, Danny Weyns, Michael Winikoff, and Peter Wurman.

References

1. AAI conference on artificial intelligence (2013) Association for the advancement of artificial intelligence. <http://www.aaai.org/Conferences/AAAI/>. Last accessed 20 March 2013
2. AAIJ: Applied Artificial Intelligence Journal (2013) Taylor and Francis. <http://www.tandfonline.com/toc/uaai20/current>. Last accessed 20 March 2013
3. AAMAS: international joint conference on autonomous agents and multiagent systems. International foundation for autonomous agents and multiagent systems. <http://ifaamas.org>. Last accessed 20 March 2013
4. Agentlink case studies. AgentLink website. <http://www.agentlink.org/resources/casestudies.html>. Last accessed 11 April 2013
5. Balke T, Hirsch B, Lützenberger M (2013) Assessing agent applications: r&D vs. R&d. In: Ganzha M, Jain LC (eds) Multiagent systems and applications, vol 1, practice and experience. Springer, Berlin, pp 1–20

6. Belecheanu R, Munroe S, Luck M, Payne T, Miller T, Pechoucek M, McBurney P (2006) Commercial applications of agents: lessons, experiences and challenges. In: Proceedings of industrial track, fifth international joint conference on autonomous agents and multi-agent systems (AAMAS 2006), ACM, Hakodate
7. Berry P, Peintner B, Conley K, Gervasio M, Uribe T, Yorke-Smith N (2006) Deploying a personalized time management agent. In: Proceedings of the fifth international joint conference on autonomous agents and multiagent systems (AAMAS'2006), ACM, New York, pp 1564–1571
8. Bubela T, Strotmann S (2008) Toward a new era of intellectual property: from confrontation to negotiation. In: Designing metrics to assess the impacts and social benefits of publicly funded research in health and agricultural biotechnology. The international expert group on innovation and intellectual property, Toronto
9. Dignum F, Dignum V (2010) Designing agent systems: state of the practice. *Agent Oriented Software Eng* 4(3):224–243
10. Edelman B, Ostrovsky M, Schwarz M (2007) Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords. *Am Econ Rev* 97(1):242–259
11. Fischer K, Kuhn N, Müller JP (1994) Distributed, knowledge-based, reactive scheduling in the transportation domain. In: Proceedings of the tenth IEEE conference on artificial intelligence for applications, San Antonio, March 1994, pp 47–53
12. Franklin S, Graesser A (1996) Is it an agent, or just a program?: a taxonomy for autonomous agents. In: Müller JP, Jennings NR, Wooldridge MJ (eds) *Intelligent agents III: agent theories, architectures, and languages*. Lecture notes in artificial intelligence, vol 1193. Springer, Berlin, pp 21–35
13. Gage D (2012) The venture capital secret: 3 out of 4 start-ups fail. *The Wall Street Journal*, (September 2012). <http://online.wsj.com/article/SB10000872396390443720204578004980476429190.html>. Last accessed 20 March 2012
14. Godin B, Doré C (2012) Measuring the impacts of science: beyond the economic dimension, inrs urbanisation, culture et société. In: Proceedings of international conference “science impact - rethinking the impact of basic research on society and the economy”, Vienna, May 2007. The Austrian Science Fund (FWF) and the European Science Foundation (ESF). http://www.csiic.ca/PDF/Godin_Dore_Impacts.pdf. Last accessed 20 March 2012
15. Groff E, McEwen T (2012) Identifying and measuring the effects of information technologies on law enforcement agencies. US Department of Justice and Institute for Law and Justice, Alexandria. <http://www.cops.usdoj.gov/Publications/e08084156-IT.pdf>. Last accessed 20 March 2012
16. Huhns MN (1987) *Distributed artificial intelligence*. Pitman/Morgan Kaufmann, San Mateo
17. IAT (2013) International conference on intelligent agent technology. <http://cs.gsu.edu/wic2013/iat> (URL for the 2013 conference, no series website available). Last accessed 20 March 2013
18. ICAART (2013) International conference on agents and artificial intelligence. <http://icaart.org>. Last accessed 20 March 2013
19. JAAMAS (2013) *Autonomous agents and multiagent systems*. Springer, New York. <http://link.springer.com/journal/10458>. Last accessed 20 March 2013
20. KER (2013) The knowledge engineering review. *Cambridge Journals*. <http://journals.cambridge.org/action/displayJournal?jid=ker>. Last accessed 20 March 2013
21. Leitao P, Vrba P (2011) Recent developments and future trends of industrial agents. In: Marik V, Vrba P, Leitao P (eds) *Holonic and multi-agent systems for manufacturing (HoloMAS 2011)*. Lecture notes in computer science, vol 6867. Springer, Berlin
22. Mankins JC (1995) Technology readiness levels: a white paper. NASA Office of Space Access and Technology. http://decadal.gsfc.nasa.gov/PACE-IDL/Final_Report/Additional_Docs/Technology_Readiness_Levels-expanded.pdf. Last accessed 20 March 2012
23. MATES (2013) German conference on multiagent technologies. <http://www-ags.dfki.uni-sb.de/~klusch/mates-series/index.html>. Last accessed 20 March 2013

24. Measuring the impact of information and communication technology for development. United Nations conference on trade and development, 2011. http://unctad.org/en/docs/dtlstict2011d1_en.pdf. Last accessed 20 March 2012
25. Merriam-Webster Online Dictionary (2013) Merriam-Webster. <http://www.merriam-webster.com>. Last accessed 2 April 2013
26. Müller JP, Bauer B, Friese T, Roser S, Zimmermann R (2006) Software agents for electronic business: opportunities and challenges (2005 re-mix). In: Chaib-Draa B, Müller JP (eds) Multi-agent-based supply chain management. Studies in computational intelligence. Springer, Berlin, pp 63–102
27. OECD guide to measuring the information society 2011. The OECD's working party on indicators for the information society, 2009. <http://www.oecd.org/sti/measuring-infoeconomy/guide>. Last accessed 20 March 2012
28. Osterweil LJ, Ghezzi C, Kramer J, Wolf A (2008) Determining the impact of software engineering research on practice. *Computer* 41(3):39–49
29. PAAMS (2013) International conference on practical applications of agents and multi-agent systems. <http://www.paams.net>. Last accessed 20 March 2013
30. Parunak HVD (1987) Manufacturing experience with the contract net. In: Huhns MN (ed) Distributed artificial intelligence. Morgan Kaufmann Publishers, San Mateo, pp 285–310
31. Sandholm T (1993) An implementation of the contract net protocol based on marginal cost calculations. In: Proceedings of the eleventh national conference on artificial intelligence (AAAI). AAAI Press, 1993
32. Subirana B, Eckes C, Herman G, Sarma S, Barrett M (2003) Measuring the impact of information technology on value and productivity using a process-based approach: the case for RFID technologies. Technical report 4450-03.CCS 223, MIT Sloan
33. Veloso M, Stone P, Han K (1998) The CMUnited-97 robotic soccer team: perception and multiagent control. In: Proceedings of the second international conference on autonomous agents, ACM, pp 78–85
34. Wagner T, Gasser L, Luck M, Odell J, Carrico T (2005) Impact for agents. In: 4th international joint conference on autonomous agents and multi-agent systems (industry track), ACM, pp 93–99