MAS platforms as an enabler of enterprise mobilisation: The state of the art *

Patrik Mihailescu
Intelligent Systems Research & Innovation Centre
BT Group
Ipswich, UK
patrik.2.mihailescu@bt.com

John Shepherdson
Intelligent Systems Research & Innovation Centre
BT Group
Ipswich, UK
john.shepherdson@bt.com

Paul Marrow
Pervasive ICT Research Centre
BT Group
Ipswich, UK
paul.marrow@bt.com

Lyndon Lee
Applications Solutions
BT Exact
Ipswich, UK
lyndon.lee@bt.com

Habin Lee[1]
Intelligent Systems Research & Innovation Centre
BT Group
Ipswich, UK
ha.lee@bt.com

Abstract - *One of the main application areas for multi-agent systems technology is enterprise mobilization, wherein the main business process actors are nomadic workers. An agent's autonomy, sociality and intelligence are highly prized features when it comes to supporting those mobile workers who are geographically isolated from the main knowledge source (i.e. the corporate Intranet) and are frequently moving from one location to another. Based on experience gained from two field trials of applications (built using for multi-agent systems technology and running on lightweight handheld devices) that support mobile business processes for telecommunications service provisioning and maintenance, this paper proposes desirable metrics for any multi-agent systems platform intended for enterprise mobilisation use. These metrics are then used to compare a number of existing multi-agent systems platforms, and based on the results, this paper identifies some areas for improvement.*

Keywords: Multi-agent systems, mobile business processes, evaluation metrics.

1 Introduction

As today's market environment is characterised by the dynamic changes of customer needs and shortening product life cycles, agility is considered as one of the most important competencies of modern companies. In this environment, the effective management of a mobile workforce is regarded as one of the key means to increase the agility of a company, as the mobile workforce is one of the main channels in detecting the changing needs of customers and obtaining feedback on new products and/or services.

Despite its importance, the information system (IS) support provided for the management of a mobile workforce is frequently inferior to that provided for in-office workers. This is mainly due to the fact that the IS support requirements for a mobile workforce are different in many ways from those for in-office workers, as the latter work within a reliable computing environment [7]. Furthermore, the behavioural differences between mobile and in-office workforces also make it difficult for mobile workers to directly re-use an IS developed for in-office workers. This proposition has already been proved by some studies in the Task/Technology Fit research area [9][17].

Multi-agent systems (MAS) are considered as one of the main technologies to support mobile workers as the intelligence, autonomy, and sociality of software agents can easily be used to benefit them [1][12]. However, despite this potential usage, most existing MAS platforms have been primarily designed for in-office end users. Only a handful of MAS platforms target the mobile computing environment.

This paper aims to identify features that should be incorporated in a MAS platform in order to support an enterprise's mobile workforces. From this, we derive a set of metrics that are used to evaluate several MAS platforms in order to establish the goodness of fit of each with the needs of the mobile computing environment.

This paper is organized as follows. The next section reviews related work and Section 3 explains how the metrics for the evaluation of MAS platforms have been derived. Section 4 briefly describes the MAS platforms selected with our designed criteria for evaluation and then details the evaluation results. Finally, Section 5 discusses the issues found during the evaluation and concludes this paper.

2 Literature review

The comparison of MAS platforms is considered difficult due to the lack of an agreed set of specific metrics, a consequence of the fact that many MAS platforms have been developed using different design philosophies, and

[1] Equal Authorship

targeted on different domains. As a result, existing evaluations or comparisons of MAS platforms are based on either generic or high level metrics. Mangina [13] reviewed thirty-six MAS platforms mostly focusing on their components and features. At the end of the review, he differentiated the platforms by their licensing policy and origins (academia or commercial). Giang and Tung [8] have also performed a similar study. They however used metrics such as the type of Java virtual machine (JVM) used, message type (KQML or FIPA ACL for example), security features, support for agent mobility, etc, for their evaluation. Ricordel and Demazeau [15] surveyed and compared four MAS platforms from a developer's point of view. Their comparison is focused on evaluating the level of support provided for each stage of a development methodology (analysis, design, development, and deployment). Dikaiakos et al. [6] compared the performance of three mobile agent platforms. For this, they proposed a hierarchical performance evaluation framework, which consisted of four layers of metrics.

Carabelea and Boissier's work [2] has a similar motivation to this paper in that they focused on MAS platforms designed for mobile devices. However, they used basic metrics for the comparison, such as the target device, communication protocols supported, FIPA compliance, target JVM, etc. This paper is distinguished from their study by the metrics used for the evaluation and comparison of platforms. In this paper, we identify metrics specific to mobile computing for our comparison of MAS platforms.

3 Characterizing a MAS platform as an enabler of enterprise mobilization

The unique features that an IS must provide in order to support nomadic workers are derived from the constraints imposed by the mobile computing environment. Based on the literature review in the mobile computing area [2] and the experience gained from two field trials where a MAS platform was used to support mobile teams in the UK and Germany [1], we have identified a set of metrics can be used to evaluate the suitability of a MAS platform to support a mobile workforce. The metrics are classified into four categories as shown in Table 1, and the following sub-sections detail why these metrics were chosen.

3.1 Usability

One of the unique features that characterize the IS support typically provided for the mobile workforce is the inferior computing environment, compared with that available in an office environment. Usability refers to the ability of a MAS platform to overcome the environmental barriers preventing mobile workers from accessing services. These barriers include temporary or long-term network disconnections, foreseen device failure (e.g. low battery), and awkward input mechanisms.

Table 1 Required features of a MAS platform for enterprise mobilisation

Domain	Features
Usability	Platform replication Agency recovery Agency re-connection Ghost agent management Agent mobility Multi-modal human agent interaction Reasoning support
Device adaptability	Supported configurations Modularity support Ease of configuration User interface independence
Communication efficiency	Data compression support Message buffering support Firewall penetration by GPRS
Lightness	Static RAM footprint Dynamic RAM footprint Boot-up time

Usability can be broken down into the following metrics:

Platform replication ensures a robust application that can continue the provision of services even in the event of failure at the platform level. Replication can be achieved either locally or remotely.

Agency recovery ensures an application agent can recover from a previous known stable state in case of unexpected failure.

Agency re-connection support ensures an application agent can re-join an agent society after a temporal disconnection due to network outage.

Ghost agent management support is a service of a MAS platform that resolves the issue of ghost agents, that is, agents that are registered in an agent society but are no longer available due to either device failure or long term network disconnection.

Agent mobility is an important feature in the context of a mobile workforce, as it ensures an application agent can move from one mobile device to another mobile device in case of foreseen problems such as low power.

Multi-modal human agent interaction increases the usability of a mobile application by increasing the reach-ability of the mobile workforce when they are executing jobs in such circumstances that they cannot monitor the screen of their mobile devices. A multi-modal interface (for instance the inclusion of speech based input/output) would be appropriate in such circumstance as a means for communication between an agent and a human user.

Reasoning support is a service of a MAS platform allowing application developers to implement a reasoning capability for an agent on a mobile device.

3.2 Device adaptability

The devices used within a mobile computing environment tend to be heterogeneous, ranging from lightweight mobile phones to heavyweight laptop computers. Therefore, the heterogeneity of computing devices imposes another barrier that should be overcome. That is, an application developed for a specific type of computing device cannot be directly used without modifications on a different computing device. As a result, the capability of a MAS platform to allow developers to configure their systems on different computing devices is considered as one of the main requirements of device adaptability. Device adaptability can be broken down into the following metrics:

Supported configuration measures what types of configurations are supported by a platform for different mobile devices.

Modularity support is a service of a MAS platform, which allows developers to selectively install only those platform components required for a specific application in order to minimize the amount of computing resources required, thereby optimising performance.

Ease of configuration measures the level of difficulty of modularity support of a MAS platform. For example, whether the MAS platform provides in-built support via configuration scripts or if application developers are forced to perform this process manually.

User interface configuration indicates whether a MAS platform provides support for implementing a user interface that is not only independent of the properties of the underlying mobile device, but also decoupled from the agent(s) that may use this user interface. This is an important requirement if agents are to be used on a variety of different mobile devices.

3.3 Communication efficiency

Mobile devices communicate over wireless networks, which in the majority of cases when compared to wireline networks are more expensive to access, more prone to network outages, and provide lower bandwidth. As a result, the capability of a MAS platform to operate efficiently under these network constraints is an important requirement. Communication efficiency can be broken down into the following metrics:

Data compression is the capability of a MAS platform to compress/decompress outgoing/incoming messages or mobile agents in a transparent way.

Message buffering is the capability of a MAS platform to buffer messages if they cannot be delivered to the intended recipient due to temporal network disconnection.

Firewall penetration by GPRS (General Packet Radio Service) is the capability of a MAS platform to enable the communication between two agents (one running within public Internet or private Intranet space, the other running in GPRS network space). This is an important issue, as some GPRS service providers have their own firewalls that do not allow direct connections from the Internet/Intranet to mobile devices within their network space.

3.4 Lightness

The limited computing resources of mobile devices impose a constraint on not just the size of a mobile application but also the level of functionality provided. As a result, the capability of a MAS platform to operate efficiently within a limited computing environment is essential. Lightness can be broken down into the following metrics:

Static RAM footprint represents the total storage size required to install a MAS platform on a mobile device.

Dynamic RAM footprint represents how much memory a platform requires when it is launched without any application agents.

Boot-up time represents the total time taken to launch a MAS platform. This is important in the context of a mobile computing environment, as users of mobile devices only interact for very short periods of time with applications, therefore any start-up delays have a negative impact on users.

4 Evaluation of the platforms

This section details the results from the evaluation of currently available MAS platforms, which target mobile devices. Section 4.1 briefly describes how the MAS platforms were selected and the methodology used for their evaluation is described in Section 4.2. Finally, Section 4.3 summarizes the evaluation results.

4.1 Selection of MAS platforms for evaluation

Anecdotal evidence suggests that there are a large number of MAS platforms in circulation. Therefore to reduce the number to be evaluated against the metrics identified in Section 3, we used three screening criteria: i) publicity, ii) explicitness, and iii) availability. Publicity was measured by checking if the MAS could be found on the Internet using a search engine or was described within published papers. Explicitness was fulfilled if the MAS platform explicitly stated that it was targeting mobile devices. Availability was satisfied if the MAS platform is available for use, either in an unrestricted or restricted (evaluation) mode. Based on these three criteria, five platforms were selected for evaluation, as shown in Table 2.

A detailed description of each MAS platform can be found via the references provided in this paper.

4.2 Evaluation methodology

The evaluation of the selected MAS platforms was completed by checking both the user/developer guides and source code provided within the platform distribution.

Furthermore, where an analytical approach was not feasible, live experiments were conducted using a PDA device and a GPRS network.

Table 2 Selected MAS platforms

Platform	Developer	License	Last Update
JADE-LEAP [10]	LEAP Consortium	LGPL	May 2004
DIET Agents[5]	BT Group	GPL	May 2004
Micro FIPA OS[14]	University of Helsinki	Emorphia Public License	Sept 2001
Cougaar[3]	BBN Technologies	Cougaar Open Source License	January 2003
kSACI[11]	Uni. of São Paulo / Uni. of Federal de Pernambuco	Open Source	February 2002

Figure 1 shows the environment in which the live experiments were conducted. The environment consisted of two devices: i) a wireless PDA (D1, RAM 64M, Processor ARM SA1110), and a laptop (D2). Both devices ran the MAS platform under investigation at the time and a single agent. Device (D1) communicated over a GPRS network, while device (D2) was connected to a wireline network. Communication between the devices was achieved via a virtual private network (VPN).

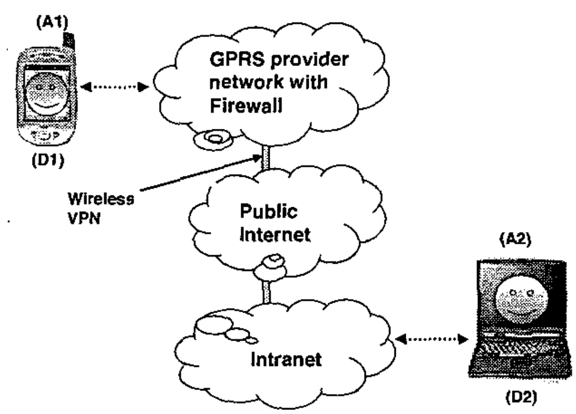

Figure 1 Platform evaluation environment

The metrics evaluated through live experiments were: i) agency re-connection support, ii) ghost agent management support, iii) message buffering support, iv) firewall penetration by GPRS, v) run time RAM footprint, and vi) boot up time. We now discuss how the first four of these were evaluated.

Agency re-connection was evaluated by terminating the network connection on device (D1), and after two minutes re-establishing the network connection, and checking to see if agent (A1) could reconnect to its previous community by sending a message. If the message was received by agent (A2) the platform passed. *Ghost agent management* was evaluated by terminating agent (A1) abnormally by stopping the JVM. Then after two minutes agent (A2) sent a message to agent (A1). If the platform indicated a failure has occurred with the message delivery, then it passed. *Message buffering* was evaluated by terminating the network connection for two minutes on device (D1), during this time agent (A2) sent a message to agent (A1). After two minutes the network connection was re-established for device (D1), and if the message was received, then the platform passed. *Firewall penetration by GPRS* was evaluated by sending a message from agent (A2) to agent (A1). If agent (A2) received the message, then the platform passed.

4.3 Evaluation results

Table 3 summarizes the evaluation result of the platforms. Note that each platform was evaluated 'as is' with no optimisations (such as those described in [15]) applied.

Firstly, Cougaar was the only platform that failed to run within the environment described in section 4.2. This is because it uses a Virtual Machine (KVM) which is no longer supported and is not available.

The majority of the platforms failed many of the usability metrics. In particular, no platform passed metrics such as agency recovery, multi-modal interface, and provision of an inference engine. On the other hand, agent mobility is the most supported metric (by four of the five platforms).

With regard to device adaptability, two platforms provided two different configurations for two types of mobile devices. As a result only these two platforms passed the ease of configuration metric. Support for the development of a user interface for different types of mobile devices is only partially supported by one platform, the rest do not provide any support. This seems to indicate that the platforms consider the development of a user interface as an application-specific task.

Only one platform passed all the metrics for communication efficiency while the other platforms didn't consider this functionality in their design. Particularly, it is notable that most of the platforms except JADE-LEAP failed to penetrate the GPRS network provider's firewall, which prevents the use of these platforms in a GPRS environment.

The size of the static RAM footprint seems proportional to the functionality provided by the platform. The platform that passed most of the metrics had the biggest static RAM footprint while the platform that failed most of the metrics had the smallest one.

Summarizing, most of the platforms evaluated failed most of the metrics. In particular, we consider agency recovery, multi-modal interface, inference engine and UI configuration support as important areas which need to be addressed by developers of MAS platforms. JADE-LEAP is an exception, which failed very few of the metrics.

Table 3 Evaluation results for selected MAS platforms: (1) kSaci, (2) DIET Agents, (3) JADE-LEAP, (4) FIPA-OS, and (5) Cougaar.

Metric	(1)	(2)	(3)	(4)	(5)
Usability					
Platform replication	No	No	Yes	No	No
Agency recovery	No	No	No	No	No
Agency reconnection	Yes	Yes	Yes	Yes	N/A
Managing ghost agent	No	No	Yes	No	N/A
Agent mobility	Yes	Yes	Yes	No	Yes
Multi-modal interface	No	No	No	No	No
Reasoning support	No	No	No	No	No
Comms Efficiency					
Data compression	No	No	Yes	No	No
Message buffering	No	Yes	Yes	No	N/A
Firewall penetration by GPRS	No	No	Yes	No	N/A
Device Adaptability					
Supported configuration	J2ME CLDC	J2ME CDC	J2ME CDC/ CLDC	J2ME CDC	KVM TINI
Modularity support	Low	Low	Low	Low	Low
Ease of configuration	N/A	N/A	High	N/A	Med
UI configuration support	Low	Low	Med	Low	Low
Lightness					
Static RAM footprint	64.8 kb	315 kb	1,034 kb	1,298 kb	N/A
Dynamic RAM footprint	99 kb	95 kb	111 kb	340 kb	N/A
Boot-up time	2606 ms	164 ms	10107 ms	4180 ms	N/A

5 Discussion and Conclusions

From the evaluation results presented in section 4.3, the functionality that is least supported by MAS platforms is usability. This can be attributed to the fact that many of these MAS platforms have been initially designed for a PC desktop environment, and as such only satisfy some minimal criteria such as compliance with Java 1.1, but do not adequately address the constraints of a mobile computing environment.

One of the main characteristics of an agent is the ability to solve unstructured problems using a knowledge base, and the fact that no platform provided an inference engine is worth noting. This seems to be due to limitations in computing resources. However, the lack of computing resources did not stop two platforms from implementing an agent internal architecture based on FIPA specifications. Therefore, the provision of a lightweight inference engine embedded in an agent could be one goal for MAS platforms intended for enterprise mobilization use.

Some readers may disagree with the inclusion of user interface metrics for the evaluation of a MAS platform. However, it should be noted that this paper is addressing a narrow application domain within a mobile computing context, wherein heterogeneous computing devices are used, and user interface issues are highly relevant due to the nature of the user's environment (see more detail regarding this issue in [4]). It should also be noted that the usability metrics have a trade-off relationship with the lightness metrics. For example, the inclusion of a reasoning capability for an agent would increase both the static and dynamic RAM footprint on a mobile device.

There are a range of optimisations that could have been applied in an attempt to improve the figures for the 'lightness' metrics across the board. However, it was never the intention of this paper to evaluate optimised versions, instead the 'out of the box' version was used in each case.

Due to both time constraints and page limitations, certain metrics have been excluded from this paper. For example, message latency, dynamic RAM usage, and platform stability (the length of time a MAS platform can provide its services) have been excluded although they are seen as important metrics within a mobile computing environment. And, it should be noted that the metrics presented within this paper originated from the analysis of requirements for mobile business processes. Therefore the metrics may differ from those for say a MAS for mobile commerce or an entertainment environment. Finally, it must be mentioned that every effort has been made to ensure an extensive search was conducted to locate suitable MAS platforms that could satisfied the initial screening criteria. However, we acknowledge that there may be a number of MAS platforms that satisfy all our metrics, but were not located, despite our best effort.

In summary, the contributions of this paper are twofold. Firstly we identify functionality that is considered useful if supported by a MAS platform for enterprise mobilization, and secondly we evaluate a handful of publicly available MAS platforms to give an overview of the state of the art. Despite the limitations mentioned above, this paper proposes some areas for improvement for MAS platforms to enable them to better support enterprise mobilization.

References

[1] M. Berger, M. Buckland, M. Bouzid, H. Lee, N. Lhuillier, D. Olpp, J. Picault and J. Shepherdson, "An Approach to Agent-based Service Composition and it's Application to Mobile Business Processes", *IEEE Transactions on Mobile Computing*, Vol 2, No. 3, pp. 197-206, 2003.

[2] C. Carabelea and O. Boisseare, "Multi-Agent Platforms on Smart Devices: Dream or Reality?", In *Proceeding of Smart Object Conference*, 2003.

[3] Cougaar web site, http://micro.cougaar.org/.

[4] D. R. McGee and P. Cohen, "Exploring handheld, agent-based, multimodal collaboration", Proceeding of Workshop on Handheld CSCW, Seattle, USA, November 1998.

[5] DIET web site, http://diet-agents.sourceforge.net/.

[6] M. Dikaiakos, M. Kyriakou and G. Samaras, "Performance Evaluation of Mobile-Agent Middleware: A Hierarchical Approach, *Lecture Notes in Artificial Intelligence*, Vol. 2240, pp 244 – 259, 2001.

[7] G. Forman and J. Zahorjan, "The Challenges of Mobile Computing," *IEEE Computer*, Vol. 27, No. 4, pp. 38-47, April 1994.

[8] N. T. Giang and D. T. Tung, "Agent Platform Evaluation and Comparison", *Technical Report, Institute of Informatics*, Slovak Academy of Science, 2002.

[9] D. L. Goodhue and R. L. Thompson, "Task-Technology Fit and Individual Performance", *MIS Quarterly*, Vol. 19, No. 2, pp. 213-236, 1995.

[10] JADE-LEAP web site, http://leap.crm-paris.com/.

[11] kSaci web site, http://www.cesar.org.br/~rla2/ksaci/.

[12] H. Lee, J. Shepherdson and P. Mihailescu, "A Multi-agent System for Team-based Job Management in Telecommunications Service Environment", *Special Issue on Jade of Telecom Italia Journal Exp*, Vol. 3, No. 3, pp. 96-105, Sep 2003.

[13] E. Mangina, Review of Software Products for Multi-Agent Systems, *White Paper of AgentLink*, Applied Intelligence Ltd., June 2002.

[14] Micro FIPA-OS web site, http://www.cs.helsinki.fi/group/crumpet/mfos/.

[15] P. Mihailescu, H. Lee and J. Shepherdson, "Hold the sources: A Gander at J2ME Optimisation techniques", *Proceedings of the 1st International Workshop on Ubiquitous Computing*, pp 73 – 82, April 2004.

[16] P. Ricordel, Y. Demazeau, "From Analysis to Deployment: a Multi-Agent Platform Survey", *Lecture Notes on Artificial Intelligence*, Vol. 1972/2000, Proceeding of Workshop on Engineering Societies in the Agent World, 2000.

[17] I. Zigurs and B.K. Buckland, "A Theory of Task-Technology Fit and Group Support System Effectiveness", *MIS Quarterly*, Vol. 22, No. 3, pp.313-334, 1998.