

Simulating Agents' Mobility and Inaccessibility with \mathcal{A} -GLOBE Multi-agent System

David Šišlák, Michal Pěchouček, Milan Rollo, Martin Reháč, and Jan Tožička

Department of Cybernetics
Czech Technical University in Prague
Technická 2, Prague 6, 166 27 Czech Republic
sislakt@feld.cvut.cz
pechoucek@labe.felk.cvut.cz

Abstract. People use multi-agent systems for three kinds of activities – (i) integration (ii) simulation (iii) complex problem solving. In this contribution we describe architecture of a newly developed multi-agent platform that supports real world simulation, with particular emphasis to all kinds of mobility and communication inaccessibility. \mathcal{A} -GLOBE performance benchmarks compared against other agent platforms are also stated in this paper.

1 Introduction

The reason why people use multi-agent systems for modelling and simulation (in contrary to classical mathematical and symbolic modelling techniques) is easy and natural transition from the modelling process into the real distributed operation. For distributed operation (e.g. control, diagnostics, planning, ...) can be deployed the same set of agents that are used in the modelling phase. Industrial partners often require substantial testing and experimental verification of the advanced AI systems before they are ready for industrial deployment. Agent technology is an ideal concept for this.

At present several Java-based multi-agent platforms from different developers are available, but none of them fully supports agent mobility and communication inaccessibility. They are thus no suitable for experiments with real-world simulation. In this paper we describe architecture of newly developed agent platform \mathcal{A} -GLOBE. It is fast and lightweight platform with agent mobility support. Beside the functions common to most of agent platforms it provides Geographical Information System service to user, so it can be used for experiments with environment simulation and communication inaccessibility. \mathcal{A} -GLOBE performance benchmarks compared against other agent platforms are also stated in this paper.

\mathcal{A} -GLOBE is an agent platform designed for testing experimental scenarios featuring agents' position and communication inaccessibility, but it can be also used without these extended functionalities [1]. The platform provides functions for the permanent and mobile agents, such as communication infrastructure, store, directory services, migration function, service deployment, etc.

Comparing to the classical agent platforms and development environment, \mathcal{A} -GLOBE is lightweight and is optimized in order to consume just a limited amount of resources. Communication in \mathcal{A} -GLOBE is very fast. Comparison with the others agent platforms can be found in Section 3. The price we have to pay for this is the fact that \mathcal{A} -GLOBE platform is not fully compliant with the FIPA [2] specifications. This is why it does not support communication between different agent platforms (e.g. with JADE, JACK, FIPA-OS, etc). Technically, it is not a very difficult problem and custom tailored interoperability can be implemented upon request. Once we are ready to sacrifice interoperability we can significantly improve ease of communication, stability of the system and overall lightweightness. For large scale scenarios the problems with system performance (memory requirements, communication speed) that interoperability brings are not negligible. Interoperability is also unnecessary when developing closed systems, where no communication outside these systems is required (e.g. agent-based simulations).

\mathcal{A} -GLOBE is suitable for real-world simulations including both static (e.g. towns, ports, etc.) and mobile units (e.g. vehicles). In such case the platform can be started in extended version with Geographical Information System (GIS) services and Environment Simulator (ES) agent. The ES agent simulates dynamics (physical location, movement in time and others parameters) of each unit. In Section 4 we present two different simulation scenarios where the \mathcal{A} -GLOBE system has been deployed.

2 System Architecture

The system integrates one or more agent platforms. The \mathcal{A} -GLOBE design is shown in Figure 1. Its operation is based on several components:

- **agent platform**¹ – provides basic components for running one or more agent containers, i.e. container manager and library manager (section 2.1);
- **agent container** – skeleton entity of \mathcal{A} -GLOBE, ensures basic functions, communication infrastructure and storage for agents (section 2.2);
- **services** – provide some common functions for all agents in one container;
- **environment simulator (ES) agent** – simulates real-world environment and controls visibility among other agent containers (section 2.4);
- **agents** – represent basic functional entities in specific simulation scenario.

Simulation scenario is defined by a set of actors represented by agents residing in the agent containers. All agent containers are connected together to one system by the GIS services. Beside the simulation of dynamics the ES agent can also control communication accessibility among all agent containers. The GIS service applies accessibility restrictions in the message transport layer of the agent container.

¹ We shall distinguish between the agent development (or integration) platform such as JADE and the agent platform used in this context. The agent platform here is a highest level component of the system architecture and resides on one Java Virtual Machine (JVM). Communication between these platform is obviously possible.




Fig. 1. System Architecture Structure

2.1 Agent Platform

The main design goals were to develop the platform as lightweight as possible and to make it easily portable to different operating systems and devices (like PDA). The platform is implemented as an application running on Java Virtual Machine (JVM version 1.4 or higher is required). Several platforms can run simultaneously (maximum 1000) on one computer, each in its own JVM instance. When new agent container is started, it can be specified in which platform it will be created and running.

The platform ensures the functionality of the rest of the system using two main components:

- **Container Manager.** In one agent platform can run one or more agent containers. Container Manager takes care of starting, execution and finishing these containers. Containers are mutually independent except for the shared library manager. Usage of one agent platform for all containers running on one computer machine is beneficial because it rapidly decreases system resources requirements (use of single JVM), e.g. memory, processor time, etc.
- **Library Manager.** The Library Manager takes care of the libraries installed in the platform and monitors which agents/services use which library. Descriptor of each agent and service specifies which libraries the agent/service requires. The Library Manager is also responsible for moving libraries of any agent migrating to other platform when required libraries are not available there. Whenever an agent migrates or agent/service is deployed, Library Manager checks which libraries are missing on the platform and obtains them from the source platform. The inter-platform functionality of the Library Manager is realized though the service library/loader (this service

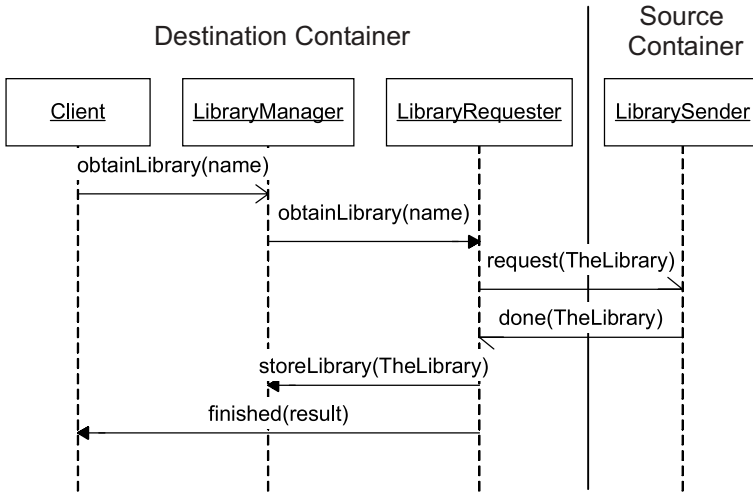


Fig. 2. Library Deployment Sequence Diagram

is present on every agent container). Library deployment sequence diagram is shown in Figure 2. The user can add, remove and inspect libraries using the container GUI.

2.2 Agent Container

The agent container hosts two types of entities that are able to send and receive messages: agents and services. Agents do not run as a stand-alone applications, but are executed inside the agent containers, each in its own separate thread. Container provides the agents and services with several low level functions (message transport, agent management, service management). Most of the higher level container functionality (agent deployment, migration, directory facilitator, etc.) is provided as standard container services.

The agent container components are:

- **Container Core.** The Container Core starts and shuts down all container components.
- **Store.** The purpose of *Store* is to provide permanent storage through interface which shields its users from the operating system filesystem. It is used by all container components, agents and services. Each entity in the agent container (agent, service, container components) is assigned its own virtual storage, which is unaffected by the others. Whenever an agent migrates, its store content is compressed and sent to the new location.
- **Message Transport.** The Message Transport is responsible for sending and receiving messages from and to the container.
- **Agent Manager.** The Agent Manager takes care of creation, execution and removal of agents on the container. It creates agents, re-creates them

after platform restart, routes the incoming messages to the agents, packs the agents for migration and removes agent’s traces when it migrates out of the platform or dies.

- **Service Manager.** The Service Manager takes care of starting and stopping the services present in the agent container and their interfacing to other container components. The user can start, stop and inspect the services using GUI. There are two types of services – user services and system services. The system services are automatically started by the container and form a part of the container infrastructure (agent mover, library deployer, directory services etc.). The system services cannot be removed. The user services can be started by user or any agent/service. The user services can be either permanent (started during every container startup) or temporary (started and stopped by some agent).

The *container name* must be unique inside one system build from several containers. This name is also used for determination specific store subdirectory for the agent container and registered to the *Environment Simulator Agent*.

Message Transport: The message transport is responsible for sending and receiving messages. Shared TCP/IP connection for message sending is created between every two platforms when the first message is exchanged between them. The message flow inside the platform is shown on figure 3.

The message structure respects FIPA-ACL [2]. Messages are encoded either in XML or Byte64 format. Message content can be in XML format or String. The structure of each message content in XML format is described by Document




Fig. 3. Message Flow

Type Definition (DTD). For coding and decoding XML messages the Java APIs for XML Binding (JAXB) [3] package is used. For transport, all binary data (e.g. libraries, serialized agents, etc.) are encoded using the open source Base64 coding and decoding routines. Agent may receive messages without using conversation discrimination (all messages incoming to this agent are processed in one method), otherwise it must use **conversation manager** with tasks.

Conversation Manager and Tasks: Usually, an agents deals with multiple jobs simultaneously. To simplify a development of such agents, the *A-GLOBE* offers *tasks*. A task is able to send and receive messages and to interact with other tasks. The Conversation Manager takes care of every message received by the agent to be routed to the proper task. The decision, to which **Task** a message should be routed, depends on the message **ConversationID**. The **ConversationID** should be viewed as a 'reference number'.

Agents: The agents are autonomous entities with unique name and ability to migrate. There is a separate thread created for each agent. A wrapper running in the thread executes the agent body. Whenever an agent enters an error state or agent body exits, the control is passed back to the wrapper, which handles the situation. The return value of the agent state is used to determine agent's termination type (die, migrate, suspend). Therefore potential agent failures are not propagated to the rest of the agent container. Agents could be deployed on remote containers.

Services: The services are bound to particular container by their identifier. There could be the same service on several containers. Services do not have their own dedicated thread and are expected to behave reactively on response to incoming messages and function calls. Services can be deployed to remote container. The agents (and services or container components) have two ways how to communicate with a service. Either via normal messages or by using the *service shell*. The service shell is a special proxy object that interfaces service functions to a client.

The advantage of service shell is an easy agent migration (for migration description see section 2.3): while the service itself is not serializable, the service shell is. When an agent migrates, the shell serializes itself with information what service name it was connected to. When the agent moves to the new location, the shell reconnects to its service at the new location. When a service is shut down, it notifies it's shells so that they refuse subsequent service calls. There are several common services such as remote command execution, searching agents/services directories, service deployment, etc. These services are automatically started by the agent container and provide common functions for all agents.

Agent Name: The agent name is globally unique and is normally generated by platform during agent creation. The service name is unique only within one agent container (services cannot migrate) and is specified by the service creator. An address has the following syntax:

```
aglobe://platform_ip:port/[agent|service]/name.
```

2.3 Migration Procedure

In order to successfully migrate, the agent has to support serialization. The migration sequence is shown in Figure 4.

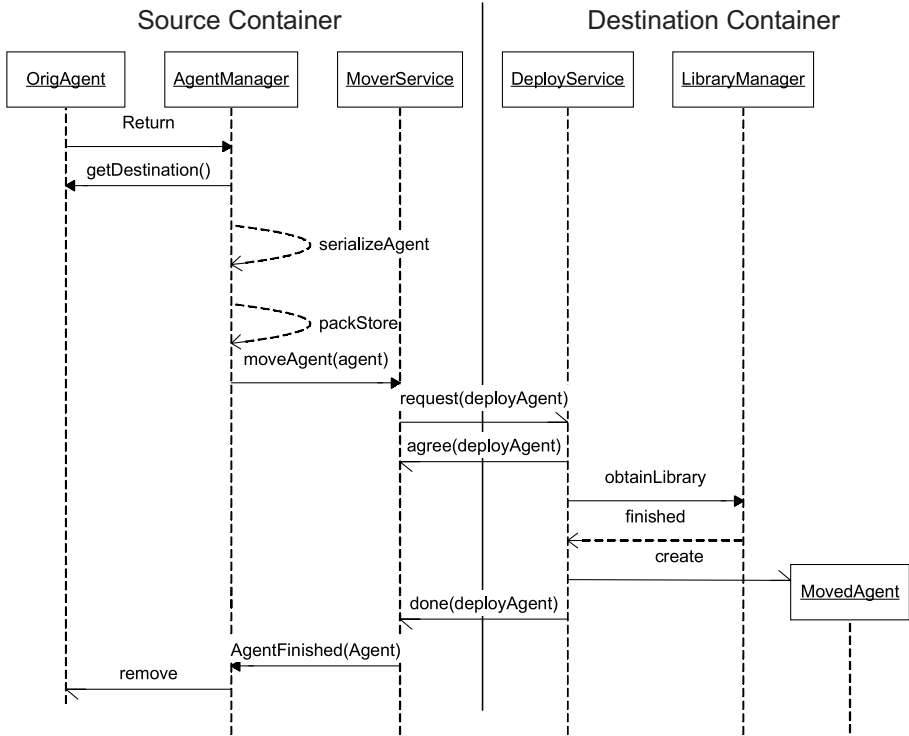


Fig. 4. Agent Migration

All exceptions that might occur during the process are properly handled and the communication is secured by timeouts. If the migration cannot be finished for any reason, the agent is re-created in its original container.

If the **done** message is successfully sent by the agent destination container but never received by the source container, two copies of the agent emerge. If the **done** message is received by the source container, but the agent creation fails at the destination container, the agent is lost. These events can never be fully eliminated due to the possible communication inaccessibility, but maximum caution was given to minimize their probability.

2.4 Environment Simulator

The purpose of ES is to simulate the real world environment. More specifically the ES models the platform mutual accessibility and informs each container

about other platforms inside its communication range. Besides visibility, the ES can inform the containers about any other parameters (eg. position, temperature, ...). The ES consists of two parts:

- **ES Agent**, which generates the information and
- **GIS services** that are present at every platform connected in the scenario.

Presence of the ES agent allows the container freely communicate with all other containers. The ES agent architecture allows simulation of complicated container motion and environment dynamics. There are two ES agents implemented:

- Manual (Matrix) ES Agent This agent provides simple user-checkable visibility matrix. The user simply checks which containers can communicate together and which can not.
- Automatic ES Agent This agent is fully automatic environment simulator. It moves mobile agent containers representing mobile units in virtual world and automatically controls accessibility between them. The visibility is controlled by means of simulation of the short range wireless link. Therefore each container can communicate only with containers located inside pre-defined radius limit. As the containers move, connections are dynamically established and lost.

GIS services distribute visibility information to all container message transport layers. There are two types of GIS services - one for server side and one for client side, as presented in Figure 1.

Server side is container where ES Agent and `gis/master` service are running. On the client side `gis/client` service is running. After the container startup, the service subscribes with the `gis/master` to receive the environmental updates (visibility, etc) and this information is than accessible to any container component (agent or service) interested in it. This client services are connected to the message transport layer and control message sending (when agent tries to send message to agent whose container is not visible, this message is returned as undeliverable). If no ES Agent is started, all platforms are connected without any restrictions.

2.5 Sniffer Agent

The Sniffer Agent is an on-line tool for monitoring all messages and their transmission status (delivered or not-reachable target). This tool helps to find and resolve communication problems in system during development phase.

The sniffer can be started only on an agent container where `GIS/master` service is running. After the sniffer starts, all messages between agents and services inside any container or among two agent containers are monitored. Messages can be filtered by the sender or receiver of the message. All messages matching the user-defined criteria are shown in the sniffer gui. The message transmission status is emphasized by type of line. The color of the message corresponds to the message performative.

Table 1. Message delivery time results for selected agent platforms

JAVA-based Agent Development Toolkits/Platforms - Benchmark Results						
April 2004, Rockwell Automation in Prague						
PIII, 600MHz, 256MB Agent Platform	Message sending - average roundtrip time (RTT)					
	agents: 1 pair messages: 1.000 x ↔		agents: 10 pairs messages: 100 x ↔		agents: 100 pairs messages: 10 x ↔	
	serial [ms]	parallel [s]	serial [ms]	parallel [s]	serial [ms]	parallel [s]
JADE v3.1	0,8	0,36	7,5	0,19	76,3	0,49
JADE v3.1 1 host, 2 JVM, RMI	10,3	4,92	111,9	6,35	1 190,5	7,14
JADE v3.1 2 hosts, RMI	5,79	3,30	68,8	3,71	770,3	2,48
FIPA-OS v2.1.0	28,6	14,30	607,1	30,52	2 533,9	19,50
FIPA-OS v2.1.0 1 host, 2 JVM, RMI	20,3	39,51	205,2	12,50	x	x
FIPA-OS v2.1.0 2 hosts, RMI	12,2	5,14	96,2	5,36	x	x
ZEUS v1.04	101,0	50,67	224,8	13,28	x	x
ZEUS v1.04 1 host, 2 JVM, ?	101,7	51,80	227,9	x	x	x
ZEUS v1.04 2 hosts, TCP/IP	101,1	50,35	107,6	8,75	x	x
JACK v3.51	2,1	1,33	21,7	1,60	221,9	1,60
JACK v3.51 1 host, 2 JVM, UDP	3,7	2,64	31,4	3,65	185,2	2,24
JACK v3.51 2 hosts, UDP	2,5	1,46	17,6	1,28	165,0	1,28
A Globe v1.0	0,3	0,10	2,8	0,04	28,4	0,09
A Globe v1.0 1 host, 2 JVM, TCP/IP	2,4	0,33	24,6	0,88	242,7	0,98
A Globe v1.0 2 hosts, TCP/IP	2,2	0,33	13,9	0,31	96,5	0,44

3 Platform Comparison

This section presents the results of comparison of available JAVA-based agent development frameworks evaluated by an independent expert Pavel Vrba from Rockwell Automation Research Center in Prague [4], which were carried out in a cooperation with the Gerstner Laboratory.

These benchmarks were focused especially on the platform performance which is a crucial property in many applications. Detailed description of the particular features is beyond the scope of this paper. Firstly, the particular benchmark criteria, which the agent platform should provide are identified (e.g. small memory footprint and message sending speed). These benchmarks were carried out for following agent platforms² - JADE [5] [6], FIPA-OS [7], ZEUS [8], JACK [9] and A-GLOBE [10].

3.1 Message Speed Benchmarks

The agent platform runtime, carrying out interactions, should be fast enough to ensure reasonable message delivery times. The selected platforms have been put

² GRASSHOPPER's licence does not allow to use it for benchmarking and other comparison activities.




Fig. 5. Message delivery time - serial test results

through a series of tests where the message delivery times have been observed under different conditions.

In each test, so called *average roundtrip time* (avgRTT) is measured. This is the time period needed for a pair of agents (let say A and B) to send a message from A to B and get reply from B to A. The roundtrip time is computed by the agent A when a reply from B is received as a difference between the receive time and the send time. This message exchange was repeated several times (depending on the type of experiment) and the results were computed as an average from all the trials.

The overall benchmark results are presented in the table 1. More transparent representation of these results in the form of bar charts is depicted in Figure 5. Three different numbers of agent pairs have been considered: 1 agent pair (A-B) with 1000 messages exchanged, 10 agent pairs with 100 messages exchanged within each pair and 100 agent pairs with 10 messages per pair. Moreover, for each of these configurations two different ways of executing the tests are applied.

In the *serial* test, the A agent from each pair sends one message to its B counterpart and when a reply is received, the roundtrip time for this trial is computed. It is repeated in the same manner N-times (N is 1000/100/10 according to number of agents). The *parallel* test differs in such a way that the A agent from each pair sends all N messages to B at once and then waits until all N replies from B are received.

Different protocols used by agent platforms for the inter-platform communication are mentioned: Java RMI (Remote Method Invocation) for JADE and FIPA-OS, TCP/IP for ZEUS and A-GLOBE and UDP for JACK. Some of the tests, especially in the case of 100 agents, were not successfully completed mainly because of communication errors or errors connected with the creation of agents. These cases are marked by a special symbol.

3.2 Memory Requirements Benchmark

This issue is mainly interesting for running thousands of agents on the one computer at the same time. Approximate memory requirements per agent can be seen in Figure 6.




Fig. 6. Approximate memory requirements per agent

4 Simulation

The functionality and usefulness of the \mathcal{A} -GLOBE approach to agent modelling and simulation has been tested on two different scenarios, that we will describe briefly below:

4.1 NAIMT – Underwater Mines Detection

Features of the \mathcal{A} -GLOBE platform were verified on the simulation of identification/removal of mines situated in given area using a group of autonomous robots. This simulation was developed within the Naval Automation and Information Management Technology (NAIMT) project. This software simulation of real-life hardware robots was required to enable scalability experiments and efficient development and verification of embedded decision making algorithms.

The goal of the group of robots is to search the whole area, detect and remove all mines located there. To allow mine removal a video transmission path must be established between the base (operated by human crew that gives the robot a permission to remove the mine) and the robot who has found the mine. Typically, relying via the other robots is necessary, because the video transmission range is limited (e.g. wi-fi connection or acoustic modems in underwater environment). Figure 7 shows an example of robots transmitting a video to base. In this scenario two types of communication accessibility are included:

- **High bandwidth** accessibility, necessary for video transmissions, is limited.
- **Signaling** accessibility, used for coordination messages and position information, is currently assumed to be perfect.

All robots in the simulation are autonomous and cooperative. Their dedicated components (coordinators) negotiate in peer-to-peer manner when preparing the




Fig. 7. Relayed communication (link between the robot and base through 3 relays)

transmission path. \mathcal{A} -GLOBE ES agent and GIS services are utilized during this phase to inform the robot about others within its video transmission range.

Each robot consists of several components, implemented as \mathcal{A} -GLOBE agents running within one agent container:

- **Robot Pod** simulator, computing robot moves and updating its position with GIS server via GIS service.
- **Mine Detector** simulator, providing the decision making components with information about found mines.
- **Video** data acquisition and transmission element. This subsystem creates the data feed from the source provided by the simulation and prepares transmission path by remotely spawning one-use transmission agents along the path. Video is then transmitted as a stream of binary encoded messages.
- **Robot Coordinator** implementing search algorithm, transmission coalition establishment and negotiation.

Next step within the NAIMT project would be integration of this \mathcal{A} -GLOBE simulator with the KAoS policy and domain services [11]. This system (developed at Institute for Human and Machine Cognition [12]) provides an organizational structure to an agent community which facilitates policy management of agent actions. It structures the groups of agents into domains (e.g. dynamic task-oriented teams like in this case). KAoS policy services then allow specification, management, conflict resolution and disclosure of policies within these domains.

The goal is especially the automatic conflict resolution (e.g. the potential deadlock problem solving when several robots find mine simultaneously). In current version the human crew in base station decides about resolution of all conflict situations, which could be heavily time consuming in complex scenarios.

4.2 Ad-Hoc Networks – Distribution of Humanitarian Aid

The second scenario has been motivated by another real-life situations and was developed within InX (Inaccessibility in Multi-Agent System) project supported by EOARD – European Office for Airspace Research and Development. Let us assume that several delivery vehicles are operating in an area without available communication infrastructure (e.g. land lines, cellular phone infrastructure, etc.). In the scenario we have three kinds of actors:

- **ports** – areas where the humanitarian aid is located and distributed from
- **villages** – locations in where the humanitarian aid is required
- **trucks** – autonomous vehicles that dispatch the humanitarian aid upon request from the ports to the villages

These are represented in the \mathcal{A} -GLOBE system by *containers*, each located on one of physically distributed platforms. Communication between the platforms is possible, while within a limited range. Communication is reliable (with no lags or dropouts) and complete (at all times) among the agents residing on one agent container, while communication across the agent containers is implemented by means of the short range radio links (for example IEEE802.11 [6]). Therefore each agent container can communicate only with containers located within its radio range.

Specific experiment scenario is shown in Figure 8. There are five villages, six ports and six trucks. Positions of villages and ports in the agents' world are chosen in such a way that no village can directly communicate with any port. They can communicate only by sending information over trucks. A truck can communicate with villages, ports and other trucks only if they are within the reach of the simulated radio link. There is at least one truck going from each village where some resource is needed to some port where this resource is distributed from. In other words there is no need for some middle storehouse to transfer the commodities through, but under special occasions the resource can be delivered into target village through others villages or ports.

Three specific sets of experiments have been carried out:

- agents that are planning the port-to-village humanitarian aid provision reside each on one truck and they collect requests when at the village and then travel to port for supplies
- on each of the vehicles we have a **relay agent** that broadcasts the request for aid to all 'visible' agents in hope to be forwarded to the right recipient (once the right couple of trucks will meet)



Fig. 8. Humanitarian Relief Scenario

- introducing the community **stand-in** agents that represent and act on behalf of the village; the stand-in agents swarm in the environment so that soon a stand-in agent from each of the village is located at each vehicle and each port. Non-trivial techniques of knowledge synchronization have been put in place.

5 Conclusion

\mathcal{A} -GLOBE agent platform supports communication inaccessibility, agent migration and deployment on remote containers. These features make \mathcal{A} -GLOBE well suited platform for simulation and implementation of physically distributed agent systems with applications ranging from mobile robotics to environmental surveillance by sensor networks. \mathcal{A} -GLOBE was designed as stream line lightweight platform which will operate on classical (PC) as well as mobile devices (PDA).

As can be seen, \mathcal{A} -GLOBE has the best results in all message sending speed benchmarks (table 1) from all selected agent platforms. In comparison with its main competitors, JADE and FIPA-OS, the \mathcal{A} -GLOBE is at least two times faster than JADE and six times faster than FIPA-OS. \mathcal{A} -GLOBE has not any communication errors. Also in memory benchmark (Figure 6) \mathcal{A} -GLOBE has one of the smallest memory requirement per agent.

Acknowledgement

Authors wish to express acknowledgement to Rockwell Automation Research Center in Prague for mutually beneficial cooperation in the platform evaluation process. \mathcal{A} -GLOBE agent platform was developed within the project "Inaccessibility in Multi-Agent Systems" (contract no.: FA8655-02-M-4057). The NAIMT deployment has been supported in part by ONR project no.: N00014-03-1-0292.

References

1. Pěchouček, M., Mařík, V., Šišlák, D., Reháč, M., Lažanský, J., Tožička, J.: Inaccessibility in multi-agent systems. final report to Air Force Research Laboratory AFRL/EORD research contract (FA8655-02-M-4057) (2004)
2. FIPA: Foundation for intelligent physical agents. <http://www.fipa.org> (2004)
3. JAXB: JAVA API for XML Binding. <http://java.sun.com/xml/jaxb> (2004)
4. Vrba, P.: Java-based agent platform evaluation. In Mařík, McFarlane, Valckenaers, eds.: *Holonic and Multi-Agent Systems for Manufacturing*. Number 2744 in LNAI, Springer-Verlag, Heidelberg (2003) 47–58
5. JADE: Java Agent Development Framework. <http://jade.tilab.com> (2004)
6. Bellifemine, F., Rimassa, G., Poggi, A.: Jade - a fipa-compliant agent framework. In: *Proceedings of 4th International Conference on the Practical Applications of Intelligent Agents and Multi-Agent Technology*, London (1999)
7. Poslad, S., Buckle, P., Hadingham, R.: The fipa-os agent platform: Open source for open standards. In: *Proceedings of 5th International Conference on the Practical Applications of Intelligent Agents and Multi-Agent Technology*, Manchester (2000) 355–368
8. Nwana, H., Ndumu, D., Lee, L., Collis, J.: Zeus: A tool-kit for building distributed multi-agent systems. *Applied Artificial Intelligence Journal* **13** (1999) 129–186
9. Fletcher, M.: Designing an integrated holonic scheduler with jack. In: *Multi-Agent Systems and Applications II.*, Manchester, Springer-Verlag, Berlin (2002)
10. A-Globe: A-Globe Agent Platform. <http://agents.felk.cvut.cz/aglobe> (2004)
11. Bradshaw, J.M., Uszok, A., Jeffers, R., Suri, N.: Representation and reasoning for daml-based policy and domain services in kaos and nomads. In: *Autonomous Agents and Multi-Agent Systems (AAMAS 2003)*, Melbourne, Australia, New York, NY: ACM Press (2003)
12. IHMC: Institute for Human and Machine Cognition. <http://www.ihmc.us> (2004)