Scalability in Multi-Agent Systems: The DIET Project

Paul Marrow
Intelligent Systems Laboratory,
BTexaCT, Adastral Park, Ipswich
IP5 3RE, UK
+44-1473-645166
paul.marrow@bt.com

ABSTRACT
The Decentralised Information Ecosystem Technologies (DIET) project is concerned with the development of a robust, adaptable and scalable software platform for multi-agent systems applied to information processing. Features of the platform, and of research directions in the project that relate to the scalability of multi-agent systems, are discussed.

Keywords
Information, Ecosystem, Decentralised, Software, Processing.

1. INTRODUCTION
The *Decentralised Information Ecosystem Technologies* project [1] (hereafter, DIET) is concerned with the construction of a multi-agent platform to support information manipulation applications. Scalability is central to this project because of the potentially unbounded nature of the information that applications based on DIET may be exposed to (for example, on the World Wide Web).

DIET is *decentralised* because it involves the construction of a Multi-Agent System (MAS) where functionality arises where possible from the local interactions between individual agents, rather than from centralised control of the system.

DIET concerns *information ecosystems*. Here, an analogy is drawn between the web of informational interactions among entities in the global information infrastructure, or among agents in a MAS, and that between living organisms in a natural ecosystem. Accordingly in the DIET project we seek to use nature-inspired computing techniques [2] in the MAS we construct. The term *information ecosystem* comes from the European Commission [3].

DIET refers to *technologies* because the intention is to develop applications based upon the MAS platform. Scalability is going to be an important issue in generating acceptable levels of application performance, as users' demands on the system change.

Previous work on scalability of MAS has identified alternative definitions of scalability [e.g., 4,5]. In DIET, we consider scalability to refer to those properties that ensure that systems can operate effectively without any implicit limits on their size. In thinking of scalability in this way, we are aware that it includes issues of demands on system resources made by the MAS, together with the size of the population of agents included, and the performance of applications developed. These issues are likely to be closely linked in any application built upon the DIET platform.

DIET is a collaboration that began in July 2000. Here I describe the aims of the project, with particular reference to the problems of scalability in MAS, and outline some results together with directions for future research. For further information readers are referred to the project web site [1] and a recent paper [6].

2. THE DIET APPROACH
The goals of the DIET project are:

- To design and implement a novel agent framework via a substantially bottom-up and ecosystem-oriented approach leading to an open, robust, adaptive and scalable software platform.
- To validate and demonstrate the usefulness of the platform via four tasks: information retrieval, information alert, information mining, and information trading.
- To research into the effects of alternative forms of interaction among agents under ecologically inspired software models.

In order to achieve the first goal, it seems appropriate to keep the functionality, memory and processing requirements associated with each individual agent to a minimum. However agents should be designed so that they can interact in a variety of ways so as to carry out functions collectively that they are not capable of individually. This relates to the inspiration from natural ecosystems which underpins the information ecosystem concept. It should also help in making the emergent multi-agent framework scalable as higher numbers of agents are needed, or larger amounts of information are processed. The software platform that is produced should assist the addressing of the second and third goals of the project, through extension to a variety of different information manipulation applications. These applications will provide a basis for the study of some of the outstanding research issues concerning interactions in multi-agent systems.

3. THE DIET PLATFORM

The DIET software platform is based upon a three-layer architecture, incorporating modularity that allows for the flexible extension of the framework. The bottom layer, the *core layer*, contains the DIET kernel that provides basic services to all implementations in the DIET architecture, as well as enforcing the constraints under which all DIET agents must operate. Above this is the *application reusable component layer* that contains software components that are shared between several applications. At the top is the *application layer*, where application-specific code is situated. The architecture also incorporates modules for visualisation and debugging, some of which will be shared between applications, and others that will be application-specific.

The DIET kernel defines DIET agents (also called *infohabitants*, following [3]), and environments, which are locations that DIET agents can reside in. The definition of location is deliberately kept general because an agent is embedded in various different spaces, which will differ in their relevance to different applications. The kernel has been designed to be as lightweight as possible, and thus to facilitate a scalable MAS. Robustness is encouraged by directly exposing agents to potential failure, and thus encouraging them to adapt their behaviour if required.

The DIET kernel provides for: the creation of agents; the formation of connections between agents in the same environment; communication between agents that are connected; movement of agents between environments; destruction of agents. There can be multiple environments on one machine, or spread across several machines. Environments are stand-alone, restricting the system resources that agents within an environment can access, and encouraging scalability through decentralisation. The DIET kernel can impose limits on system resources such as the number of threads that are used and the size of message buffer used by each agent. These limits allow larger populations of agents to be handled without overloading system resources, and thus should encourage scalability. In addition, the computational cost of environmental services does not depend upon the number of agents. In particular, a communication connection between agents can be set up quickly and in constant time. Once a connection has been established it is still under control of the kernel but agents communicate directly, ensuring that there is no communication bottleneck when the number of agents grows. Another feature that should assist the development of highly scalable applications is that DIET agents can give up their threads temporarily when they are not required, only to retake them when they need to restart some behaviour.

An application that illustrates the DIET platform's potential for scalability has been constructed: it sorts agents according to their numerical identity (see [6] for more information). In it, Trigger agents connect Linker agents in sorted sequences using a localised algorithm that relies on messages being passed along sorted sequences of Linkers. As sequences grow longer a single message sent by a Trigger initiates more activity in the system. Triggers therefore adapt their behaviour to the perceived system load and thus ensure that performance remains close to optimal at all times. Because the DIET kernel and agents are lightweight and can cope with high system activity, a single computer can run the application with more than 100,000 agents. These results are a very encouraging support to the aim of producing a scalable MAS platform, although of course future applications can be expected to include very different types of agents and agent behaviour, necessitating more research into the properties of systems based on the DIET platform.

4. FUTURE RESEARCH IN DIET

Following the development of the DIET software platform, research in the DIET project will focus on the development of several information processing applications, involving information retrieval, information alert, information mining, and information trading. Research will also be carried out into the properties of intermediary agents, or information brokers, that will be needed to facilitate interaction between diverse other types of agents. In each case the understanding of scalability will be important in producing applications that perform effectively and robustly under realistic conditions.

Particular areas that will be of interest in DIET are systems with multiple or distributed brokers, or "revolutionary" brokering systems where there are no specialised brokers, and all interactions are carried out locally between agents. In both cases the scaling of interactions, with the size of the agent population, or the throughput of information, will be important to study so as to ensure scalable applications.

5. CONCLUSIONS

Here I have introduced the DIET project and outlined how it addresses scalability in multi-agent systems. The emphasis in the project on decentralised approaches to information processing means that scalability is always going to be significant in developing systems within it. I have presented here an application that demonstrates the scalability of the DIET system: future work should clarify the scalability of other applications using the DIET platform.

6. ACKNOWLEDGMENTS

This work was supported in part by the IST Programme of the European Union, under the FET Proactive Initiative – Universal Information Ecosystems [3], through project DIET (IST-1999-10088). I would also like to thank all those involved in the project.

7. REFERENCES

[1] DIET project web site. http://www.dfki.uni-kl.de:8080/DIET/public/index.html

[2] Marrow, P. Nature-inspired computing technology and applications. BT Technology Journal 18(4), 13-23, 2000.

[3] FET Universal Information Ecosystems web site. http://www.cordis.lu/ise/fetuie.htm

[4] Rana, O.F. and Stout, K. What We Mean When We Talk About Multi-Agent Scalability. Journal of Autonomous Agents and Multi-Agent Systems, 2001.

[5] Turner, P.J. and Jennings, N.R. Improving the Scalability of Multi-Agent Systems. Proceedings of the 1st International Workshop on Infrastructure for Scalable Multi-Agent Systems, Barcelona, 2000.

[6] Marrow, P., Koubarakis, M., van Lengen, R.H., Valverde-Albacete, F., Bonsma, E., et al. Agents in Decentralised Information Ecosystems: The DIET Approach. In: Proceedings of AISB01 Symposium on Intelligent Agents for E-commerce, York, 200