Reducing the Response Time in Network Management by Using Multiple Mobile Agents

Marcelo G. Rubinstein[†,⌘,1], Otto Carlos M. B. Duarte[†], and Guy Pujolle[⌘]
† Grupo de Teleinformática e Automação, Universidade Federal do Rio de Janeiro, COPPE/EE, CP 68504, 21945-970, Rio de Janeiro, RJ, Brazil; ⌘ Laboratoire PRiSM, Université de Versailles, 45, Avenue des Etats-Unis, 78035, Versailles, Cedex, France

Key words: Mobile Agents, network management, and performance evaluation

Abstract: This paper aims at investigating mobile agent tradeoffs in network management by performing management task simulations. Results show that the mobile agent performs better than the SNMP when the number of managed network elements ranges between two limits: an inferior bound, associated with the number of messages that pass through backbone links, and a superior bound, related to incremental size of mobile agent. Furthermore, if the network is divided into several domains and one mobile agent is assigned to each domain, instead of using a single mobile agent to manage the whole network, the results show that the response time decreases up to 90% when using 8 mobile agents.

1. INTRODUCTION

Most network management systems use SNMP (Simple Network Management Protocol) [1] and CMIP (Common Management Information Protocol) [2] protocols, which are based on a centralized paradigm. These protocols use the client-server model, on which the management station acts as a client that provides a user interface to the network manager and interacts with agents, which are servers that manage remote access to local information stored in a Management Information Base (MIB).

[1] Grant holder from CNPq - Brazil

The operations available to the management station for obtaining access to the MIB are very low level. This fine grained client-server interaction, called micro-management, generates an intense traffic that overloads the management station [3], resulting in scalability problems. But at the same time that the networks grow, the computational power of the managed element increases, making possible to perform management functions in a distributed way [4].

In this sense, mobile agents is an option to distribute and scale the network management. These agents move to the place where data are stored and select information the user wants by using intelligence; saving bandwidth, time, and money. Mobile agents decentralize processing and control, and, as a consequence, reduce the traffic around the management station, turn asynchronous the communication between the agent and the manager (useful when there are unreliable or lossy links), distribute processing load, and increase the flexibility of the management agents' behavior.

Research activities related to mobile code in network management are recent [5]. Magedanz et al. [6] have been one of the first researchers to present issues for the deployment of mobile agents in telecommunication and network management. Labetoulle et al. [7] present a state-of-the-art of intelligent agents in network management, from the management by delegation paradigm [8] to the mobile agent one, while Karmouch and Pham [5] describe the current use of mobile agents in telecommunications and network management. Management by Delegation (MbD) [8] has been the first paradigm to address decentralization and automation of management tasks by dynamically delegating management functions to agents. Pagurek et al. [9] implement a code mobility infrastructure and a suite of simple tools that interact with agents located on network components. Puliafito et al. [10] use mobile agents to collect information about the state of the network and to perform a micro-management of network devices through multiple variables. Morin et al.[11] implement a management architecture that consists of a set of servers and managers communicating with SNMP agents located on the managed network elements.

The performance of mobile agents in network management is also being investigated. Geihs et al. [12] and El-Darieby and Bieszcad [13] provide simple quantitative evaluation of mobile agent and client-server approaches. Baldi et al. [3] evaluate the tradeoffs of mobile code design paradigms in network management applications by developing a quantitative model that provides the bandwidth used by traditional and mobile code design of management functionalities. Gavalas et al. [14] analyze bandwidth utilization for the SNMP and the mobile agent. Experimental implementation results are presented in terms of bandwidth consumption and

response time to obtain an aggregation of multiple variables. Sahai and Morin [11] perform measurements of bandwidth utilization of mobile agent and client-server applications on an Ethernet LAN. They also present a single case comparison of response time for the mobile agent and the client-server. Rubinstein and Duarte [15, 16] simulate management tasks performed by mobile agents and SNMP ones, comparing both the approaches, on a topology that consists of a LAN of managed network elements connected to the management station by a bottleneck link.

In this paper, we assess the performance improvement obtained by using multiple mobile agents. The network is partitioned into several domains and one mobile agent is assigned to each domain. Transit-stub topologies, which are similar in shape to the Internet, are used in the simulations. Response time results show that the management performance increases when using multiple mobile agents.

This paper is organized as follows. Section 2 discusses the use of mobile agents in network management. Section 3 reports simulation results concerning the applicability of mobile agents in management tasks. At last, concluding remarks are presented in Section 4.

2. NETWORK MANAGEMENT USING MOBILE AGENTS

Mobile agents are created to perform their tasks in different network computers. A mobile agent can interrupt its execution and migrate from a machine to another, carrying data about its state that include information obtained from previous task executions. As the number of visited nodes grows, mobile agent size also increases, turning migration harder. One possible solution to this problem is to visit a fixed number of nodes, return to the agent home or send all data to it (reducing mobile agent size), and start the task again on the remaining nodes [15]. Another solution is to use multiple mobile agents to manage a partitioned network (Section 3). The initial size of a mobile agent also affects agent performance since the larger the size, the more difficult the migration [17]. This size depends on the task to be carried out and on the language used to implement it.

In the SNMP, network management does not scale when the size or the complexity of the network increases because of the centralized processing and control. Mobile agents can be used in order to solve this problem and it is important to investigate when they improve management efficiency.

2.1 Simulation Model

The applicability of mobile agents in carrying out network management tasks is assessed by comparing mobile agent performance with the SNMP one.

The Network Simulator (NS) is used in these simulations [18]. This discrete-event simulator provides several implemented protocols and mechanisms to simulate computer networks with node and link abstractions. Some UDP modules of the NS have had to be modified in order to effectively receive a packet, sending it to the next layer.

All simulations use transit-stub topologies, which are similar in shape to the Internet. The GT-ITM [19] generates random topologies, on which each node represents switches or routers and edges represent forwarding paths between switches (Figure 1).

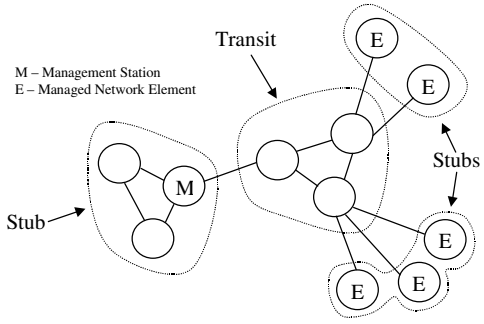

Figure 1. Network management on a transit-stub topology

Each routing domain in the Internet can be classified as either a stub domain or a transit domain. A domain is a stub domain if the path connecting any two nodes *u* and *v* goes through that domain only if either *u* or *v* is in that domain. Transit domains do not have this restriction. The purpose of transit domains is to interconnect stub domains efficiently.

A transit domain comprises a set of backbone nodes, which are typically fairly well connected to each other. In a transit domain, each backbone node also connects to a number of stub domains, via gateway nodes in the stubs.

We consider that the management station belongs to a node of a stub domain and managed network elements are located in other stub domains (Figure 1). We use the same input parameters to generate three random 272-node topologies. The management station controls groups of 16 network elements, which is the number of nodes of a stub domain.

The simulation model assumes that links and nodes have no load, links are error-free, and processing time in application layer is not taken into account. UDP is used in all simulations and UDP and IP headers are considered. The Maximum Transmission Unit (MTU) used for all links is 1500 bytes, therefore, there is no fragmentation of SNMP messages since they are small. Every request or response of a variable is sent on a different message and the considered performance parameters are bandwidth consumption on the links that connect the management station to the backbone nodes and response time in retrieving management variables.

SNMP sends requests to all elements to be managed (one after receiving the response from the other); the mobile agent goes to an element to be managed, gathers the variable, and visits all other elements, from a stub domain to another. After finishing, the mobile agent returns to the management station.

3. SIMULATION RESULTS

First, we evaluate the effect of the number of managed network elements on response time. We assume that each request or response is 50 bytes long, all links have a 2 Mbps data rate and latency links are about a few milliseconds, and the initial size of the mobile agent is 5 kbytes.

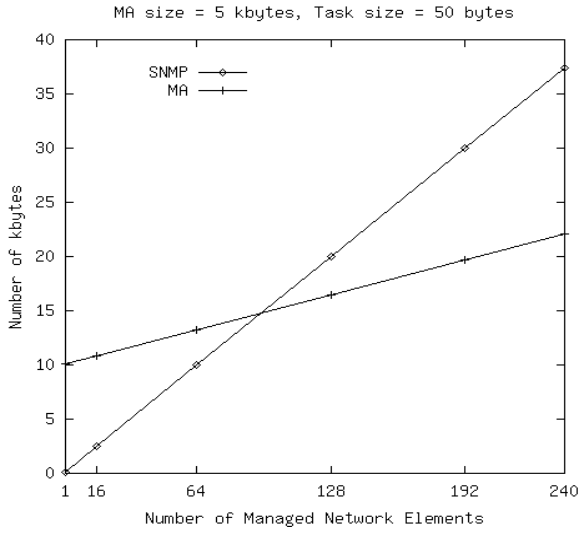

Figure 2. Bandwidth consumption per number of managed elements

For a small number of managed network elements, the SNMP uses fewer bytes than the mobile agent to perform the task (Figure 2). Nevertheless, as the number of managed elements increases, the overhead associated with several retrievals (PDU GetRequest, UDP, and IP's headers) turns mobile agent utilization more suitable.

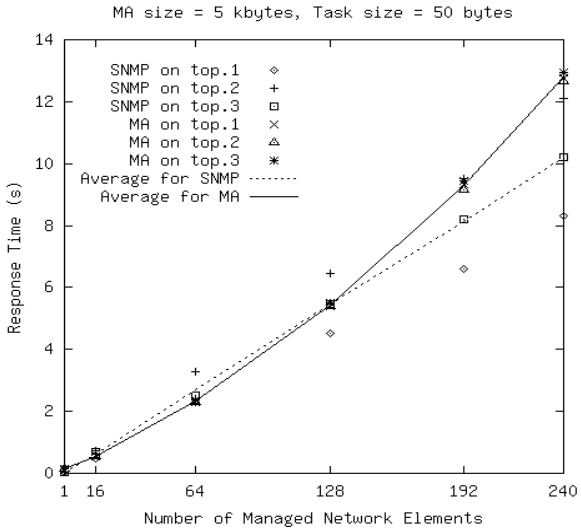

Figure 3. Response time per number of managed elements

Figure 3 shows that the mobile agent's behavior does not change with the topology, but for the SNMP, the difference of performance among response times for the three topologies increases faster when the number of managed elements is large, due to the fact that SNMP packets traverse oftentimes the backbone links, which depend on the topology. Figure 3 also presents the average response time. For a small number of managed elements, the SNMP performs better than the mobile agent due to the fact that an SNMP message is smaller than the initial size of the mobile agent. As the number of managed elements increases, response time for the SNMP grows proportionally, since the time to manage a stub domain is approximately the same for all stub domains. For the mobile agent, response time increases faster when the number of managed elements grows, due to the incremental size of the mobile agent. Therefore, the mobile agent performs better that the SNMP when the number of managed network elements ranges between two limits.

It has been show that the mobile agent size increases with the number of visited nodes, and, as a consequence, the migration becomes difficult. Moreover, when the number of elements to be managed is large, the

response time for the mobile agent will also be high. In order to improve the management performance, we divide the network into several domains and assign one mobile agent to each domain. The number of mobile agents used in the simulations is 1, 2, 4, and 8. Different initial sizes of the mobile agents (S) and task sizes (T) are used, and one of the three topologies is considered.

In order to also analyze the effect of the task on the management performance, two different tasks are used: task t_1 is the same task used in the previous experiment and task t_2 is associated with a 150-byte-long variable.

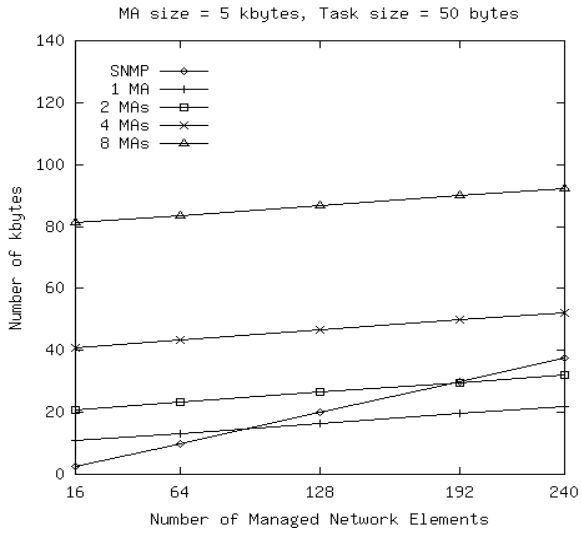

Figure 4. Bandwidth consumption for S = 5 kbytes and T = 50 bytes

Figures 4 and 5 show the bandwidth consumption per number of mobile agents for S = 5 kbytes and T = 50 and 150 bytes. The SNMP and the mobile agents are both influenced by the task to be performed. When the number of managed elements is 240, bandwidth consumption for the SNMP is 37.44 and 85.44 kbytes, respectively for T = 50 and 150 bytes; for one mobile agent, bandwidth consumption is 22.05 and 46.15 kbytes, for T = 50 and 150 bytes and 240 managed elements. When we increase the number of mobile agents, the bandwidth consumption also increases, since all mobile agents return to the management station after finishing their tasks. For example, in Figure 4, when the number of managed network elements is 240 and the number of mobile agents increases from 1 to 8, the bandwidth consumption increases from 22.05 to 92.4 kbytes.

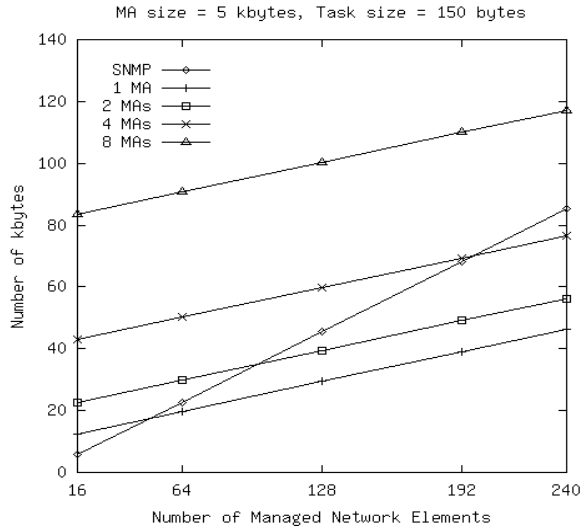

Figure 5. Bandwidth consumption for S = 5 kbytes and T = 150 bytes

For the response time, as depicted in Figures 6 and 7, the mobile agents and the SNMP also have different behaviors. For the SNMP, the task does not have a great influence due to the fact that the number of bytes exchanged between the management station and the managed network elements is small; but for the mobile agents, as the number of bytes exchanged increases, the response time also grows since the mobile agents will move with higher difficulty. For the mobile agents, as the number of managed network elements increases, the performance difference among response times for the different number of mobile agents increases faster (Figures 6 and 7). For example, when T = 50 bytes (Figure 6) and the number of managed network elements doubles from 64 to 128, the response time for 1 and 8 mobile agents goes from 2.3209 to 5.43483 s, and from 0.518067 to 0.778 s, respectively.

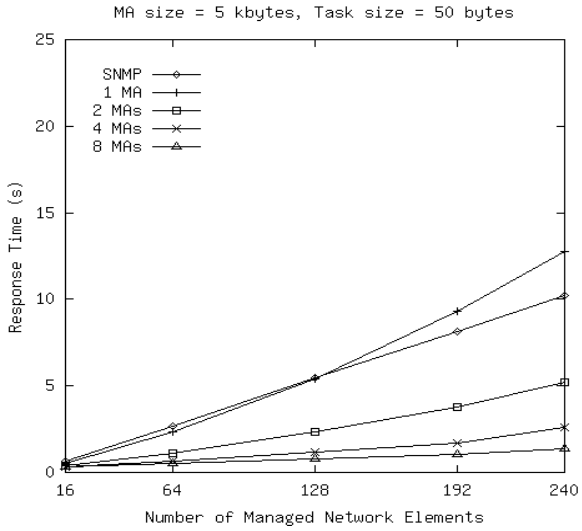

Figure 6. Response time for S = 5 kbytes and T = 50 bytes

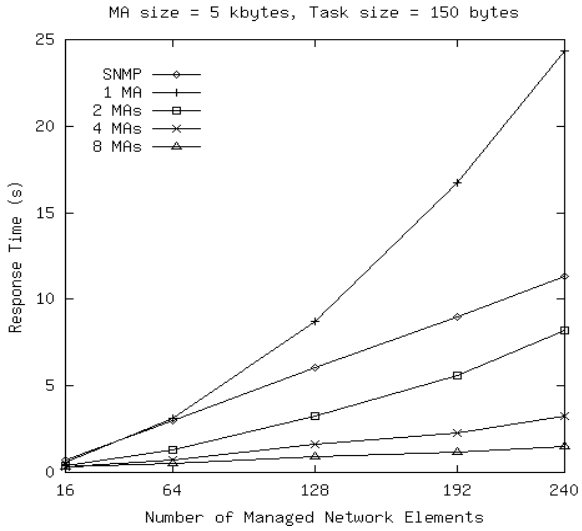

Figure 7. Response time for S = 5 kbytes and T = 150 bytes

Two different initial sizes of mobile agents are used: 1 and 8 kbytes. Figures 8 and 9 show the bandwidth consumption per number of mobile agents for T = 50 bytes and S = 1 and 8 kbytes. Comparing Figures 8 and 9, we conclude that the initial size of the mobile agent has a great influence on the bandwidth consumption.

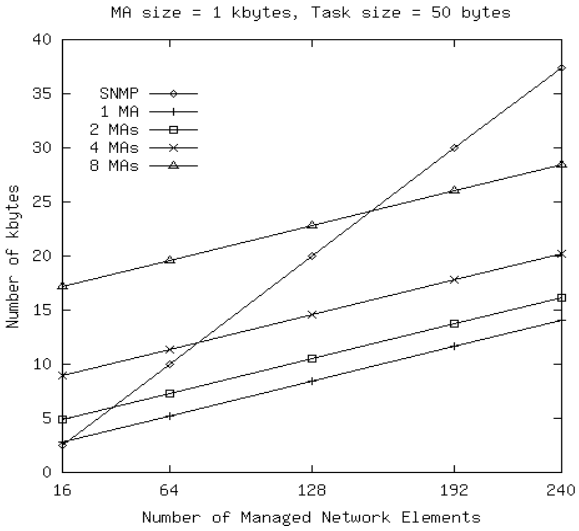

Figure 8. Bandwidth consumption for S = 1 kbytes and T = 50 bytes

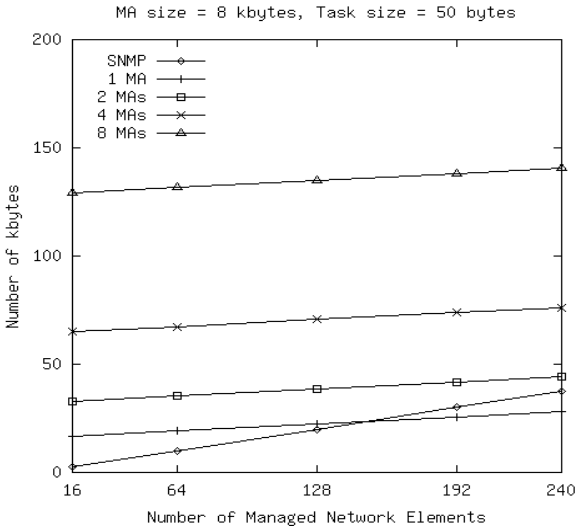

Figure 9. Bandwidth consumption for S = 8 kbytes and T = 50 bytes

Figures 10 and 11 show the response time for different mobile agent sizes. For example, when the number of managed elements is 240 and the initial size of each mobile agent goes from 1 to 8 kbytes, the response time for 1 and 8 mobile agents goes from 8.9154 to 15.6741 s, and from 0.614267 to 1.89 s, respectively.

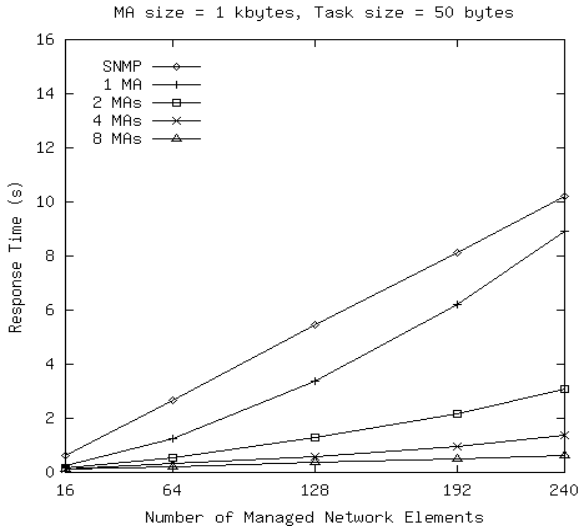

Figure 10. Response time for S = 1 kbytes and T = 50 bytes

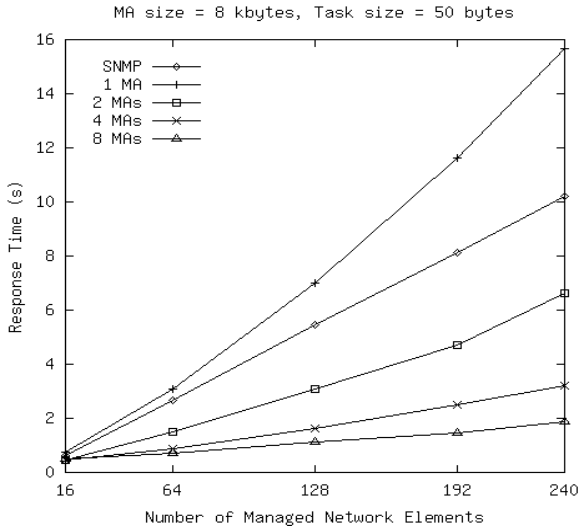

Figure 11. Response time for S = 8 kbytes and T = 50 bytes

4. CONCLUSION

This paper has analyzed tradeoffs of mobile agents in network management tasks. The performance of mobile agents has been compared with the SNMP one in several simulations.

Response time results show that the mobile agent performs better than the SNMP when the number of managed network elements ranges between two limits: an inferior bound, associated with the number of messages that pass through backbone links, and a superior bound, related to incremental size of mobile agent, which turns migration difficult.

Moreover, the management performance increases when we use multiple mobile agents. For an initial size of a mobile agent equal to 5 kbytes and a 50-byte-long task, response times for gathering 240 variables decrease by 59, 80, and 90%, for 2, 4, and 8 mobile agents, when comparing with using one mobile agent.

Network management using mobile agents seems to be a good approach and the use of multiple mobile agents enhances the management performance.

5. REFERENCES

[1] W. Stallings, “SNMP and SNMPv2: The infrastructure for network management”, IEEE Communications Magazine, vol. 36, no. 3, pp. 37-43, Mar. 1998.

[2] Y. Yemini, “The OSI network management model”, IEEE Communications Magazine, vol. 31, no. 5, pp. 20-29, May 1993.

[3] M. Baldi and G. P. Picco, “Evaluating the tradeoffs of mobile code design paradigms in network management applications”, 20th International Conference on Software Engineering (ICSE'98), Kyoto, Japan, pp. 146-155, Apr. 1998.

[4] K. Meyer, M. Erlinger, J. Betser, C. Sunshine, G. Goldszmidt and Y. Yemini, “Decentralizing control and intelligence in network management”, 4th International Symposium on Integrated Network Management, Santa Barbara, USA, May 1995.

[5] V. A. Pham and A. Karmouch, “Mobile software agents: An overview”, IEEE Communications Magazine, vol. 36, no. 7, pp. 26-37, July 1998.

[6] T. Magedanz and T. Eckardt, “Mobile software agents: A new paradigm for telecommunications management”, IEEE/IFIP Network Operations and Management Symposium (NOMS'96), Kyoto, Japan, Apr. 1996.

[7] M. Cheikhrouhou, P. Conti and J. Labetoulle, “Intelligent agents in network management: a state-of-the-art”, Networking and Information Systems Journal, vol. 1, no. 1, pp. 9-38, 1998.

[8] G. Goldszmidt and Y. Yemini, “Delegated agents for network management”, IEEE Communications Magazine, vol. 36, no. 3, pp. 66-70, Mar. 1998.

[9] A. Bieszczad, B. Pagurek and T. White, “Mobile agents for network management”, IEEE Communications Surveys, vol. 1, no. 1, Sept. 1998.

[10] A. Puliafito and O. Tomarchio, “Using mobile agents to implement flexible network management strategies”, Computer Communications, vol. 23, no. 8, pp. 708-719, Apr. 2000.

[11] A. Sahai and C. Morin, “Towards distributed and dynamic network management”, IEEE/IFIP Network Operations and Management Symposium (NOMS'98), New Orleans, USA, Feb. 1998.

[12] M. Zapf, K. Herrmann and K. Geihs, “Decentralized SNMP management with mobile agents”, Sixth IFIP/IEEE International Symposium on Integrated Network Management (IM'99), Boston, USA, pp. 623-635, May 1999.

[13] M. El-Darieby and A. Bieszczad, “Intelligent mobile agents: Towards network fault management automation”, Sixth IFIP/IEEE International Symposium on Integrated Network Management (IM'99), Boston, USA, pp. 610-622, May 1999.

[14] D. Gavalas, D. Greenwood, M. Ghanbari and M. O'Mahony, “Using mobile agents for distributed network performance management”, 3rd International Workshop on Intelligent Agents for Telecommunication Applications (IATA'99), Stockholm, Sweden, Aug. 1999.

[15] M. G. Rubinstein and O. C. M. B. Duarte, “Evaluating tradeoffs of mobile agents in network management”, Networking and Information Systems Journal, vol. 2, no. 2, pp. 237-252, 1999.

[16] M. G. Rubinstein and O. C. M. B. Duarte, “Analyzing mobile agent scalability in network management”, IEEE Latin American Network Operations and Management Symposium (LANOMS’99), Rio de Janeiro, Brazil, pp. 64-74, Dec. 1999.

[17] M. G. Rubinstein and O. C. M. B. Duarte, “Evaluating the performance of mobile agents in network management”, IEEE Global Telecommunications Conference (GLOBECOM'99), Rio de Janeiro, Brazil, pp. 386-390, Dec. 1999.

[18] K. Fall and K. Varadhan, “NS Notes and Documentation”, tech. rep., The VINT Project, Jan. 1999.

[19] E. W. Zegura, K. L. Calvert and M. J. Donahoo, “A quantitative comparison of graph-based models for internet topology”, IEEE/ACM Transactions on Networking, vol. 5, no. 6, pp. 770-783, Dec. 1997.