Using Mobile Agent Strategies for Reducing the Response Time in Network Management

Marcelo G. Rubinstein[1,2,*] Otto Carlos M. B. Duarte[1] Guy Pujolle[2]

[1] Grupo de Teleinformática e Automação
Universidade Federal do Rio de Janeiro
COPPE/EE - CP 68504 - 21945-970
Rio de Janeiro RJ - Brazil
{rubi, otto}@gta.ufrj.br

[2] Laboratoire PRiSM
Université de Versailles
45, Avenue des Etats-Unis - 78035
Versailles Cedex - France
Guy.Pujolle@prism.uvsq.fr

Abstract

This paper aims to investigate mobile agent tradeoffs in network management by performing management task simulations. The analysis focuses on the performance of two unload strategies in order to reduce the response time for the mobile agent: in the first approach, the mobile agent returns to the management station; in the second one, the mobile agent just sends data to the management station, both after visiting a fixed number of nodes. Results show that the response time decreases by 62, and 64%, for the first and second strategies, when compared with gathering all data without returning or sending these data to management station before finishing the task.

1 Introduction

Currently, most network management systems use the SNMP (Simple Network Management Protocol) [1] protocol, which is based on a centralized paradigm. This protocol 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).

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 intense traffic and computational overload on the management station, resulting in scalability problems. 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.

Network management can be distributed and scaled by the use of mobile agents, which are programs that help users to perform tasks on the network, acting on behalf of these users. These agents can 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 agent-manager communication, distribute processing load, and increase the flexibility of the management agents' behavior.

Some preliminary results related to the gain due to the unload procedure of the mobile agent after visiting a fixed number of nodes are presented in [2]. This paper overcomes this unload-procedure analysis and evaluates mobile agent performance on 272-node-transit-stub topologies, which are similar in shape to the Internet ones. In this paper, a new strategy, sending data to the management station after visiting a fixed number of nodes, is evaluated.

*Grant holder from CNPq - Brazil.

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 agent's state that include information obtained from previous task executions. As the number of visited nodes grows, the mobile agent size also increases, turning migration harder. One possible solution to this problem is to visit a fixed number of nodes, return or send all data to the agent home (reducing the mobile agent size), and continue the task on the remaining nodes (Section 3.2). The initial size of a mobile agent also affects agent performance because the larger the size, the more difficult the migration [2].

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 may solve this problem and improve the management efficiency.

2.1 Simulation Model

The applicability of mobile agents in carrying out network management tasks is assessed by comparing the mobile agent performance with the SNMP one. The Network Simulator (NS) [3] is used. This discrete-event simulator provides several implemented protocols and mechanisms to simulate computer networks with node and link abstractions.

All simulations use transit-stub topologies, which are similar in shape to the Internet ones. The GT-ITM [4] generates random topologies, on which each node represents switches or routers and edges represent forwarding paths between switches (Figure 1). 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.

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). Three different 272-node topologies are used. The management station controls groups of 16 network elements, which is the number of nodes of a stub domain.

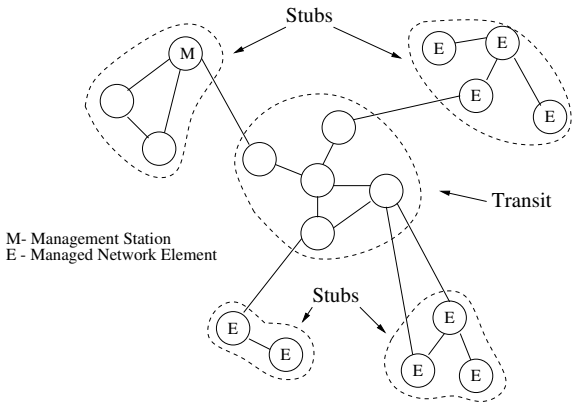

Figure 1: Network management on a transit-stub topology.

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. Every request or response of a variable is sent on a different message and the considered performance parameter is response time in retrieving the 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

Various experiments are carried out in order to evaluate the effect on the management performance of several parameters.

3.1 Effect of the Number of Managed Network Elements

This section evaluates the effect of the number of managed network elements on response time. Each request or response of a variable is 50 bytes long. All links have a 2 Mbps data rate and latency links are about a few milliseconds. The initial size of the mobile agent is 5 kbytes.

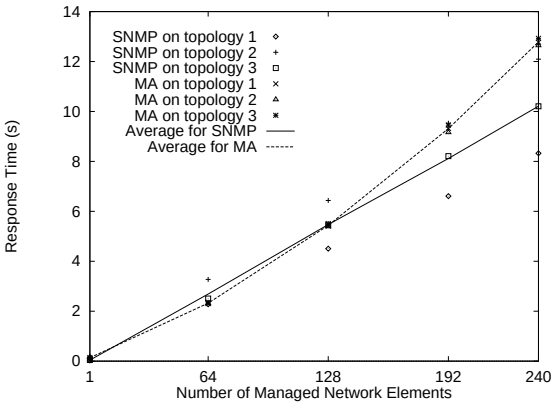

Figure 2: Response time per number of managed elements.

Figure 2 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 that depend on the topology. Figure 2 also presents the average response time per number of managed network elements. 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.

3.2 Effect of Unload Strategies

This section evaluates the performance improvement of two unload strategies: in the first one, s_1, the mobile agent returns to the management station, and for s_2, the mobile agent just sends the results to the management station. The simulations consider that the mobile agent visits a fixed number of nodes "per trip", i.e., the number of visited network elements before returning or just sending the results to the management station. Three tasks are used: task t_1 is the same task used in previous section, task t_2 has a 100-byte-long request or response, and task t_3 is associated with a 150-byte-long variable. Their purpose is to retrieve the variables from all the 240-managed-network elements. The topologies and the mobile agent size are the same of the previous experiment. The number of elements per trip varies from 1 to 240.

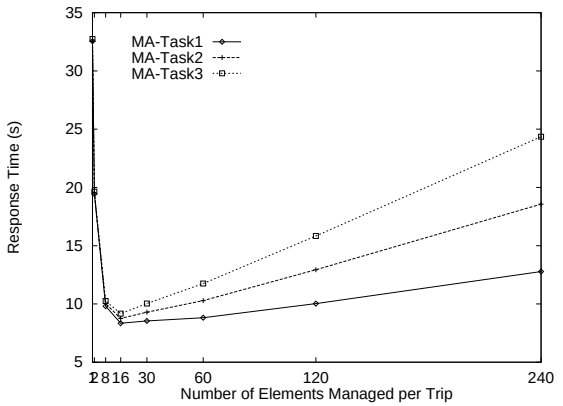

Figure 3: Response time for strategy s_1.

Figure 3 shows that the response time for strategy s_1 decreases sharply when the number of elements managed per trip is small since the agent visits few nodes and returns to the management station. As the number of nodes visited per trip goes on raising, response time decreases up to a point when it starts to increase due to agent migration difficulty related to the agent size. An "optimum point" for this strategy is associated with visiting 16 nodes and then returning to the management station to unload. The response time decreases by 35, 53, and 62%,

for tasks t_1, t_2, and t_3, when comparing each optimum point with gathering all 240 variables in one trip. This optimized strategy performs better than the SNMP for all tasks. For a small number of elements per trip, the task size does not have a great influence since the mobile agent size is small, but when the number of elements per trip is large, e.g., 240, the difference among the tasks is significant. Thus, the greater the task size the greater the improvement of this strategy.

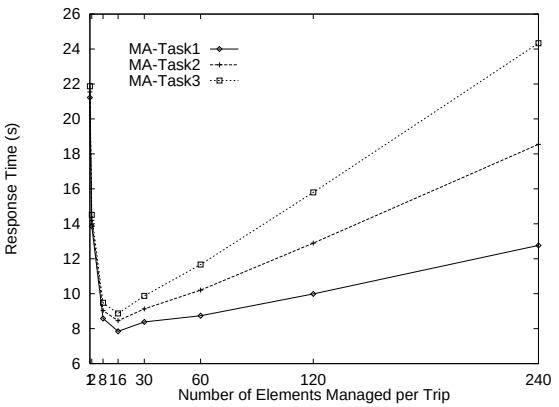

Figure 4: Response time for strategy s_2.

Figure 4 shows that the response time for s_2 decreases by 39, 54, and 64%, for tasks t_1, t_2, and t_3, when comparing each optimum point with gathering all 240 variables in one trip.

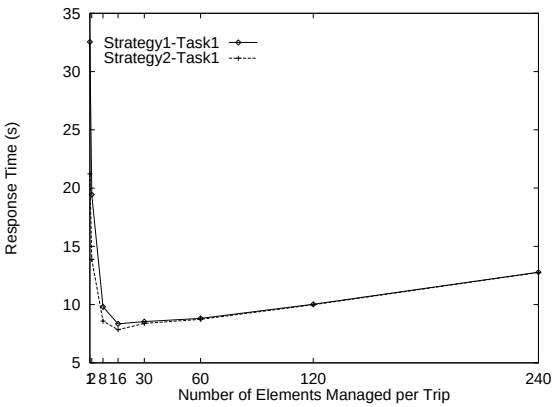

Figure 5: Response time for different tasks.

Comparing both the strategies for task t_1, strategy s_2 performs better when the number of elements managed per trip is small (Figure 5), since the mobile agent's code is not sent to the management station. As this number of elements per trip increases, both the approaches behave almost the same way, due to the fact that the size of data gathered by the agent is much greater than the size of the mobile agent's code.

4 Conclusion

Response time results show that the mobile agent performance increases when the agent returns/sends data to the management station (respectively strategies s_1 and s_2), after visiting a fixed number of nodes. Response time decreases by 35, 53, and 62%, for tasks t_1, t_2, and t_3, when compared with gathering all variables without returning to the management station before finishing the task. Using strategy s_2, the response time is reduced by 39, 54, and 64%, for tasks t_1, t_2, and t_3. For both the strategies, the response time of the mobile agent is smaller than the SNMP one.

Network management using mobile agents seems to be a good approach and the use of the presented strategies reduces the response time.

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] 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.

[3] K. Fall and K. Varadhan, "NS Notes and Documentation", tech. rep., The VINT Project, Jan. 1999.

[4] 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.