Adaptive Middleware for Distributed Simulation of massive Ad Hoc Networks

Gabriele D'Angelo*
Department of Computer Science
University of Bologna
gdangelo@cs.unibo.it
http://www.cs.unibo.it/~gdangelo

July 29, 2003

Abstract

The design and implementation of a new middleware for Parallel and Distributed Simulation involves analyzing the limitations of actual simulation tools. Among the relevant scenarios, Ad Hoc networks are gaining an increased importance but both centralized and distributed simulators does not scale. We have localized the main scalability bottleneck in at least two peculiarities of wireless systems: host-mobility and open-broadcast nature of the wireless medium. Our proposal involves a new paradigm for simulation based on the concept of component migration. Preliminary results shown a speed-up up to 23% given by adaptive distributed simulation compared to to a static distributed simulation.

1 Introduction

Our research team is actually involved in design and implementation of a middleware oriented to parallel and distributed simulation. During the first stages of design we analyzed the behavior and limitations of the most used simulation tools. In the work clearly emerged how these tools are inadequate for the simulation of some valuable scenarios [4, 9, 16, 19].

In recent years the wireless systems have gained an increasing relevance and which consequently led to a wide research work in the field of tools and methodologies for efficient simulation of wireless systems. An Ad Hoc network is formed by wireless mobile hosts, each host has a "double nature" of terminal and router in the network and operates without any form of centralized administration [10].

Wireless networks currently considered of interest usually include a massive number of simulated mobile hosts [11, 20]. The simulation of the mobility, protocol stack and applications of every simulated host can require a lot of computation and have a significant memory footprint. Such requirements make impossible or unfeasible the computation on mono-processor architectures [18]. A simple approach involves increasing the model abstraction level, modelling only a conceptual subset of the interactions. This "reduced approach" is often cause of unrealistic results. A more interesting approach involves parallel and distributed simulation models and architectures (e.g. Parallel Discrete Event Simulation [8, 17]).

Simulation models of wireless systems have at least two peculiarities with respect to wired networks' models, the first is the mobility and the second is the open-broadcast nature of the wireless medium. Topology changes due to simulated hosts' mobility reflects on causality effects of every hosts' interaction. In presence of such peculiarities a simple static or dynamic definition of publishing/subscribing lists, groups and causal domains could be far from optimality [12, 21]. We believe that a solution can be achieved exploiting the concept of components' migration.

*PhD thesis submission date: planned for middle 2004

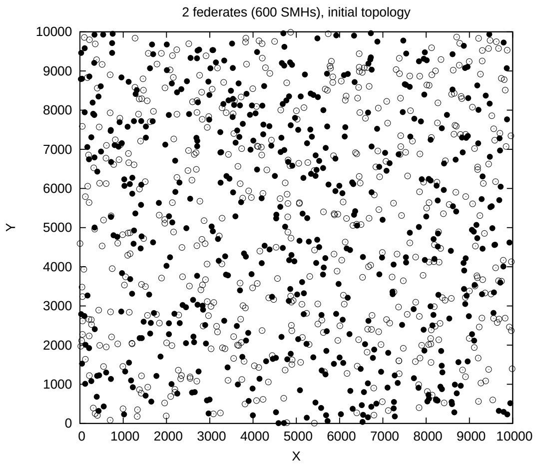

Figure 1: 600 SMHs, initial random distribution over black PEU1 and white PEU2.

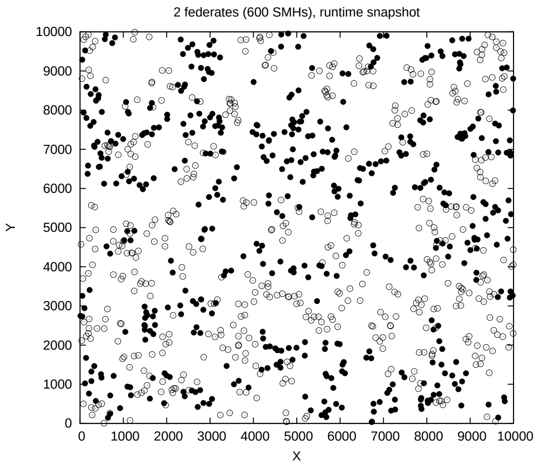

Figure 2: 600 SMHs, steady-state distribution over black PEU1 and white PEU2, migration on.

2 The GAIA framework

High Level Architecture (HLA) is a recently approved standard (IEEE 1516) [5, 6, 7] dealing with component-oriented distributed simulation. It defines rules and interfaces allowing for heterogeneous components' interoperability in distributed simulation [2, 3].

Our research involved the definition and implementation of a mechanism allowing for Simulated Mobile Hosts' (SMHs) migration in HLA-based distributed simulation. This prototype migration framework is a middleware extension called Generic Adaptive Interaction Architecture (GAIA). It adopts heuristic load balancing and migration policies aimed to dynamically reallocate and cluster the interacting SMHs among the Physical Execution Units (PEUs). The reallocation is performed with two targets: reducing the amount of external-PEU communications and performing a load-balancing in the set of PEUs.

The network latency is orders of magnitude bigger than local memory communication so every external-PEU communication causes a big overhead to the whole simulation. Reducing external communications at bare minimum means a lower overhead for the execution of simulation algorithms to achieve causal synchronization. The allocation of every SMH on the same PEU obviously would reduce the total communication overhead but it is an unscalable approach and is not feasible for the amount of computational resources required.

3 Preliminary results

As part of out tests for the evaluation of the GAIA framework we have deployed a high number of simulated mobile hosts (SMHs) moving around a toroidal-shaped 2-D arena. Every SMH follows a mixed variation of the Natural Random and Random Waypoint (RWP) [1]. The uncorrelated and unpredictable mobility pattern of SMHs is an interesting worst case scenario for our migration based framework. The modeled communication between SMHs is a constant flow of ping messages (i.e. constant bit rate), transmitted by every SMH to its nearest device.

The visual effect of the dynamic allocation between EUs is shown in Figure 1 and Figure 2. The first figure is a snapshot of the initial random distribution of 600 (out of 900) SMHs over the testing arena. For readability, two (out of three) EUs are shown: black dots refer to SMHs allocated on EU1 and white dots refer to SMHs allocated on EU2. The Figure 2 shows the allocation after a few timesteps: the dynamic clustering is demonstrated by the presence of clusters, chains and empty spaces (hidden EU3 clusters) of highly interacting SMHs. Figure 3 demonstrate that the heuristic migration

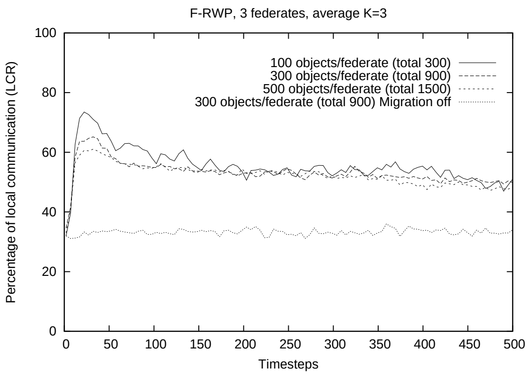

Figure 3: LCR vs. SMH density (only fast mobility F-RWP).

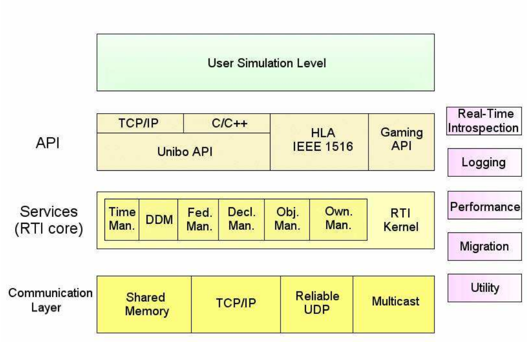

Figure 4: The conceptual structure of ARTIS.

mechanism is able to increase the Local Communication Ratio (LCR) from 33.3% (given by the uniform distribution of SMHs over 3 EUs) up to 50% for an adaptive simulation (migration on) with respect to a static one (migration off). A noteworthy aspect is that the LCR becomes stable in a few timesteps and it is invariant in respect of SMHs' density.

We ran our simulation experiments over a variable set of PEUs equipped by commodity hardware and FastEthernet network. Preliminary results shown a speed-up up to 23% given by adaptive distributed simulation with respect to a static distributed simulation, for 900 SMHs and involving 3 PEUs [13, 14, 15].

This speed-up result is surely encouraging but could be significantly increased performing some code optimizations, tuning the heuristic function used for load-balancing and migration, and also simulating a scenario characterized by less worst case assumptions.

4 Conclusions and Future Work

We propose an adaptive framework (GAIA) for the dynamic allocation of model entities in a HLA-based distributed simulation. GAIA is based on runtime migration and load-balancing policies, to reduce in an adaptive way the amount of external communication between the PEUs. Short term evolutions will increase performance, tuning and optimizing the existent framework. We think our framework can be easily extended to simulate real world protocols and complex mobile hosts' applications.

Likely the migration-based approach described in this work can be applied to a wide set of simulations as multi-agents systems, genetic systems, P2P models and agent-based economic simulations.

Actually we are involved in the first implementation steps of a brand-new middleware for parallel and distributed simulation, having control over design and full access to source code will permit a detailed evaluation of overheads and a better integration of the migration facilities in the middleware core. The new middleware, actually called ARTIS (Adaptable RTI System) will support both optimistic and pessimistic based simulations. We plan to offer simple native Applications Programming Interfaces (APIs), limited IEEE 1516 compatibility, dedicated support for Internet Gaming and Agent Based simulations (Figure 4).

5 Acknowledgments

I wish to thank Professor L. Donatiello and Dr. L. Bononi for their valuable support and contribution to this work.

References

[1] T. Camp, J. Boleng, and V. Davies. A survey of mobility models for ad hoc network re-

search. *Wireless Communications and Mobile Computing (WCMC): Special issue on Mobile Ad Hoc Networking: Research, Trends and Applications*, 2(5):483–502, 2002.

[2] Dahmann J. S., Fujimoto R., and Weatherly R. M. The Department of Defense High Level Architecture. In *Winter Simulation Conference*, pages 142–149, 1997.

[3] Dahmann J. S., Fujimoto R., and Weatherly R. M. The DoD High Level Architecture: an update. In *Winter Simulation Conference*, pages 797–804, 1998.

[4] Wayne J. Davis and Gerald L. Moeller. The high level architecture: is there a better way? In *Winter Simulation Conference*, pages 1595–1601, 1999.

[5] Defense Modeling and Simulation Office. Data Distribution and Management Design Document, V0.2, 1996.

[6] Defense Modeling and Simulation Office. High Level Architecture Interface Specification, V1.3, 1996.

[7] Defense Modeling and Simulation Office. High Level Architecture Object Model Template, 1996.

[8] Fujimoto R. M. *Parallel and Distributed Simulation Systems*. Wiley-Interscience, 2000.

[9] H. Zhao and N. D. Georganas. Collaborative Virtual Environments: Managing the Shared Spaces. *Networking and Information Systems Journal*, 3(2), 2001.

[10] Z. Haas, J. Deng, B. Liang, P. Papadimitatos, and S. Sajama. Wireless ad hoc networks, 2002.

[11] M. Heissenbttel and T. Braun. Ants-based routing in large scale mobile ad-hoc networks.

[12] D. Van Hook and J. Calvin. Data distribution management in rti 1.3. In *Proceedings of Simulation Interoperability Workshop*, 1998.

[13] Bononi L. and D'Angelo G. Dynamic host allocation for HLA-based distributed simulation of mobile ad-hoc networks. In *Proceedings of ISCS Italian Society for Computer Simulation (ISCS '02)*, 2002.

[14] Bononi L. and D'Angelo G. A novel approach for distributed simulation of wireless mobile systems. In *Proceedings of Personal Wireless Communications (PWC '03)*, 2003.

[15] Bononi L., D'Angelo G., and Donatiello L. HLA-based adaptive distributed simulation of wireless mobile systems. In *Proceedings of the 17th Workshop on Parallel and Distributed Simulation (PADS '03)*, 2003.

[16] Johannes Lüthi and Clemens Berchtold. Concepts for dependable distributed discrete event simulation. In R. Van Landeghem, editor, *Proceedings of the Int. European Simulation Multi-Conference, Ghent University, Ghent, Belgium, May 23-26, 2000*, pages 59–66. SCS, 2000.

[17] Misra J. Distributed discrete event simulation. *ACM Computing Surveys*, 18(1):39–65, 1986.

[18] George Riley Mostafa. Simulating large networks - how big is big enough?

[19] V. Paxon and S. Floyd. Why we don't know how to simulate the internet, 1997.

[20] Dhananjai Madhava Rao and Philip A. Wilsey. Simulation of ultra-large communication networks. In *MASCOTS*, pages 112–119, 1999.

[21] Georgios Theodoropoulos and Brian Logan. An approach to interest management and dynamic load balancing in distributed simulation. In *Proceedings of the 2001 European Simulation Interoperability Workshop (ESIW'01)*, pages 565–571, Harrow, London, UK, June 2001. Simulation Interoperability Standards Organisation and Society for Computer Simulation.