Detailed Simulation of Large-Scale Wireless Networks

Gabriele D'Angelo

joint work with

Michele Bracuto

University of Bologna
Department of Computer Science

DS-RT 2007 – Crete Island (Greece)

Presentation outline

- Basic assumptions and goals
- Parallel and Distributed Simulation (PADS)
- Synchronization and load balancing
- Tools and mechanisms: ARTÌS, GAIA+, WiFra
- Experimental evaluation
- Conclusions and future work

Basic assumptions

- Many of the systems of interest are composed by a very large number of entities, with a really dynamic nature and evolution
- Wireless networks, under the simulation viewpoint, have very strict requirements in terms of level of detail
- Networks composed of hundreds of thousands up to millions of nodes will be widely diffused in the next years
- Monolithic simulators are unable to fulfill such scalability requirements due to memory constrains and excessive amount of time required to obtain the results

Basic assumptions

- Parallel And Distributed Simulation can be used to aggregate memory and computational resources
- A set of interconnected Physical Execution Units (PEUs) is responsible to manage and evolve the simulation state
- This approach is not free from drawbacks, the distributed execution architecture needs appropriate mechanisms for:
  - Synchronization (Communication)
  - Load-balancing

Synchronization

- Under the synchronization viewpoint, the detailed (and efficient) simulation of wireless networks is a complex problem
- The Medium Access Control (MAC) protocols commonly used in wireless network requires very fine-grained models to represent the state of the shared medium and the behavior of wireless devices
- In example, in the case of 802.11 DCF the Short Interframe Space (SIF) is set to 10 μs. This potentially means a huge number of synchronizations in the distributed execution architecture

Goals

Main goals of this work, demonstrate that:

- Large scale wireless networks (200.000 node) can be efficiently simulated following an appropriate distributed simulation approach
- Specifically tailored techniques can improve the communication efficiency and therefore increase the simulation speed
- Very large-scale wireless networks (1.000.0000 nodes) can be simulated using clusters composed of Commercial-Off-The-Shelf (COTS) hardware

Distributed simulation: definitions

- The distributed simulation is run by a set of Physical Execution Units (PEUs)
- Each PEU is responsible to manage the execution of a Logical Process (LP)
- Each LP, simply works as a container: it manages the evolution of a set of Simulated Model Entities (SMEs)

Synchronization + Load Balancing

- As said before, synchronization (more generally the communication overhead) and load-balancing are two of the main problems to address in PADS
- The reduction of the **communication overhead** and the **computation load-balancing** can be seen as different aspects of the same problem and therefore should be addressed together
- **A single mechanism should manage both of them**
- The presence of heterogeneous hardware (e.g. different CPU models) can be seen as a special load balancing requirement

Migration-based approach

- We propose a migration-based approach: every entity in the simulation can be dynamically relocated (migrated), moving from a source LP to a new destination LP

- Clustering in the same LP the highly interacting Simulated Mobile Entities (SMEs) it is possible to reduce the costly **inter**-LP communication and conversely increasing the rate of low cost **intra**-LP communication

Load-balancing

- This approach can be also used to improve the computation load-balancing of the execution architecture
- The synchronization points in the distributed architecture can be used to tag each LP as "**fast**" or "**slow**"
- A LP is "slow" if:
  - Its PEU (i.e. CPU) is overloaded
  - Its communication network has a higher delay with respect to other parts of the execution architecture
- In both cases, to speed up the simulation, a solution is to migrate some entities from the slow LPs to faster LPs, therefore reducing the imbalance

Load-balacing

- This description of the mechanism is very high level, the "real world" implementation has to take in account many subtle details, some of them:
  - Each reallocation has a cost that depends on many factors
  - The algorithm has to be "fully distributed" (without any centralization point) and therefore without any "global vision" of the distributed system
  - The mechanism has to quickly react to internal (i.e. the creation of new entities) and external events (i.e. a burst of CPU or network load), without introducing oscillatory behaviors

ARTÌS: parallel and distributed simulation middleware

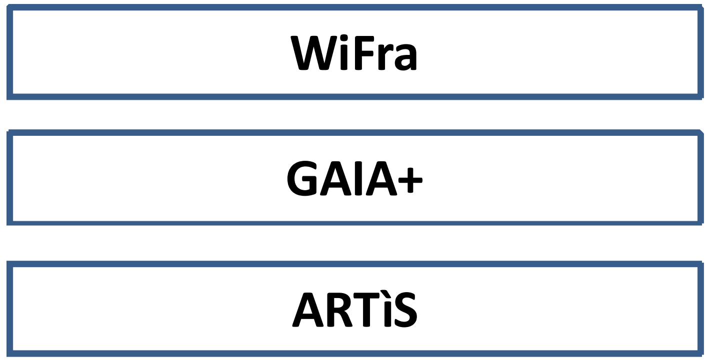

- The (Advanced RTI System) ARTÌS is a middleware for Parallel and Distributed Simulation
- Adaptively adjusts the communication behavior with respect to network technology (that is, adaptively selects Shared-Memory, R-UDP, SCTP, TCP etc. to improve the communication performances)
- Homepage: http://pads.cs.unibo.it

GAIA+: communication and computation load-balancing

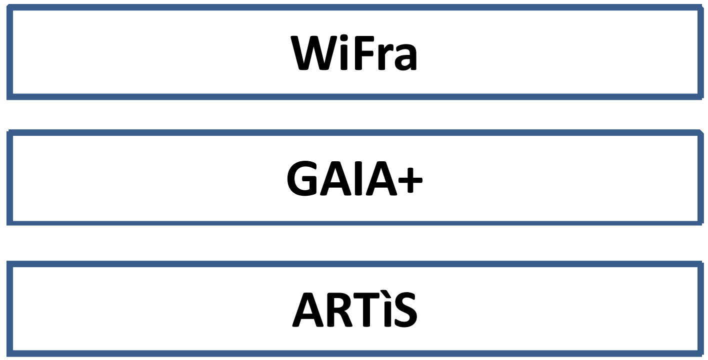

- GAIA+ framework is responsible to adaptively manage the migration of the simulated entities
- It is composed by a set of tightly coupled heuristics, that take care of the computation and communication load-balancing
- It provides an easy to use abstraction level for the definition of new simulation models

WiFra: Wireless Framework

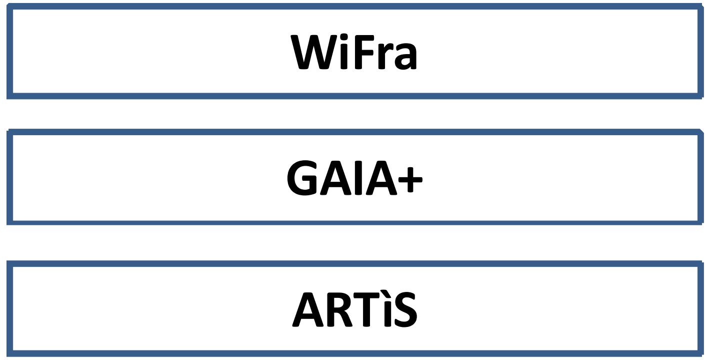

- The Wireless Framework (WiFra) is a collection of wireless models
- In this case we have implemented a IEEE 802.11 DCF MAC Layer protocol and on top of it a very simple info-mobility application
- Each entity defined in WiFra can be automatically migrated by GAIA+, depending on the defined heuristics and the tuning parameters

Experimental evaluation

Wireless model parameters	
Simulation time-step	10 μs (802.11 SIFS)
MAC layer	IEEE 802.11 DCF
Packet size	1024 bytes
Packet rate	4 pkt/s
Propagation model Transmission range	Free space propagation 250 meters
Simulated area	Variable size Fixed density of nodes
Nominal channel bit rate	2 Mbps
Mobility model	Random WayPoint (RWP)
Simulated devices (SMHs)	200.000, 1.000.000

Experimental evaluation: distributed simulation environment

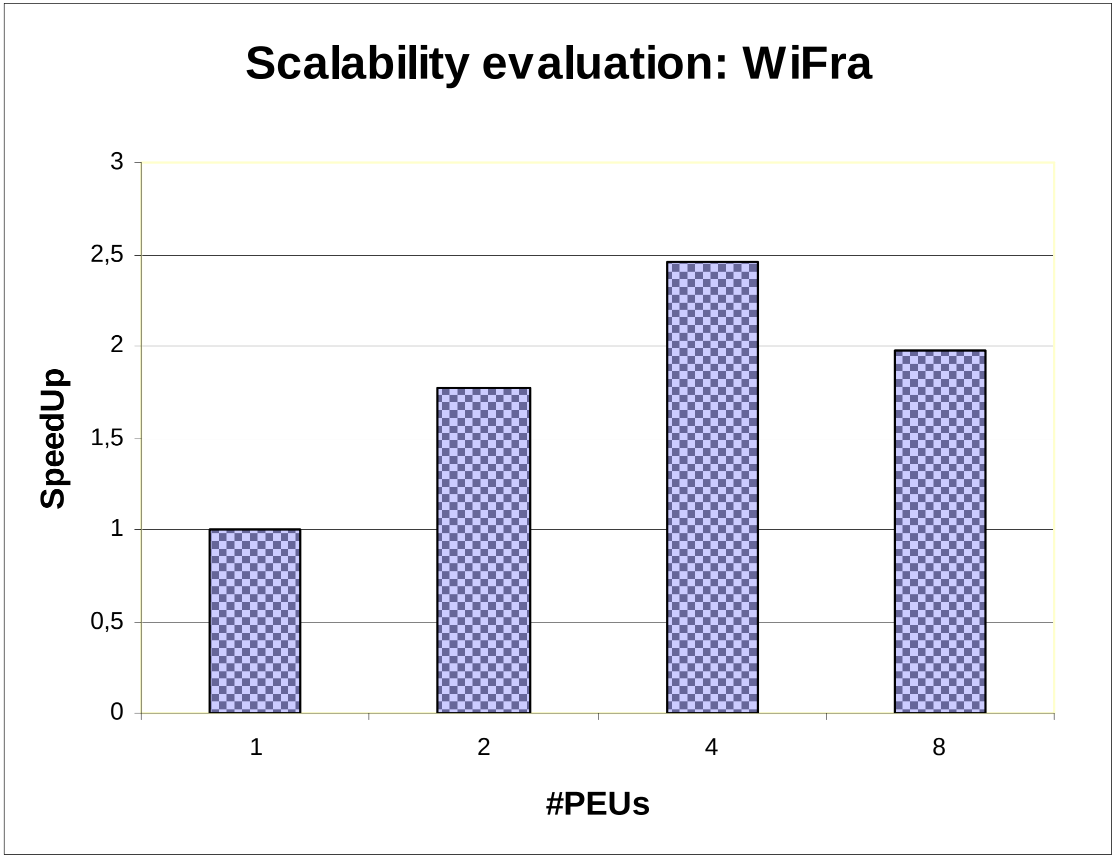

200.000 SMHs. 1 up to 8 PEUs: Dual Core Intel Pentium IV CPU 3.0 GHz with 1GB of RAM, interconnected by Fast Ethernet

Experimental evaluation: GAIA+ ON/OFF

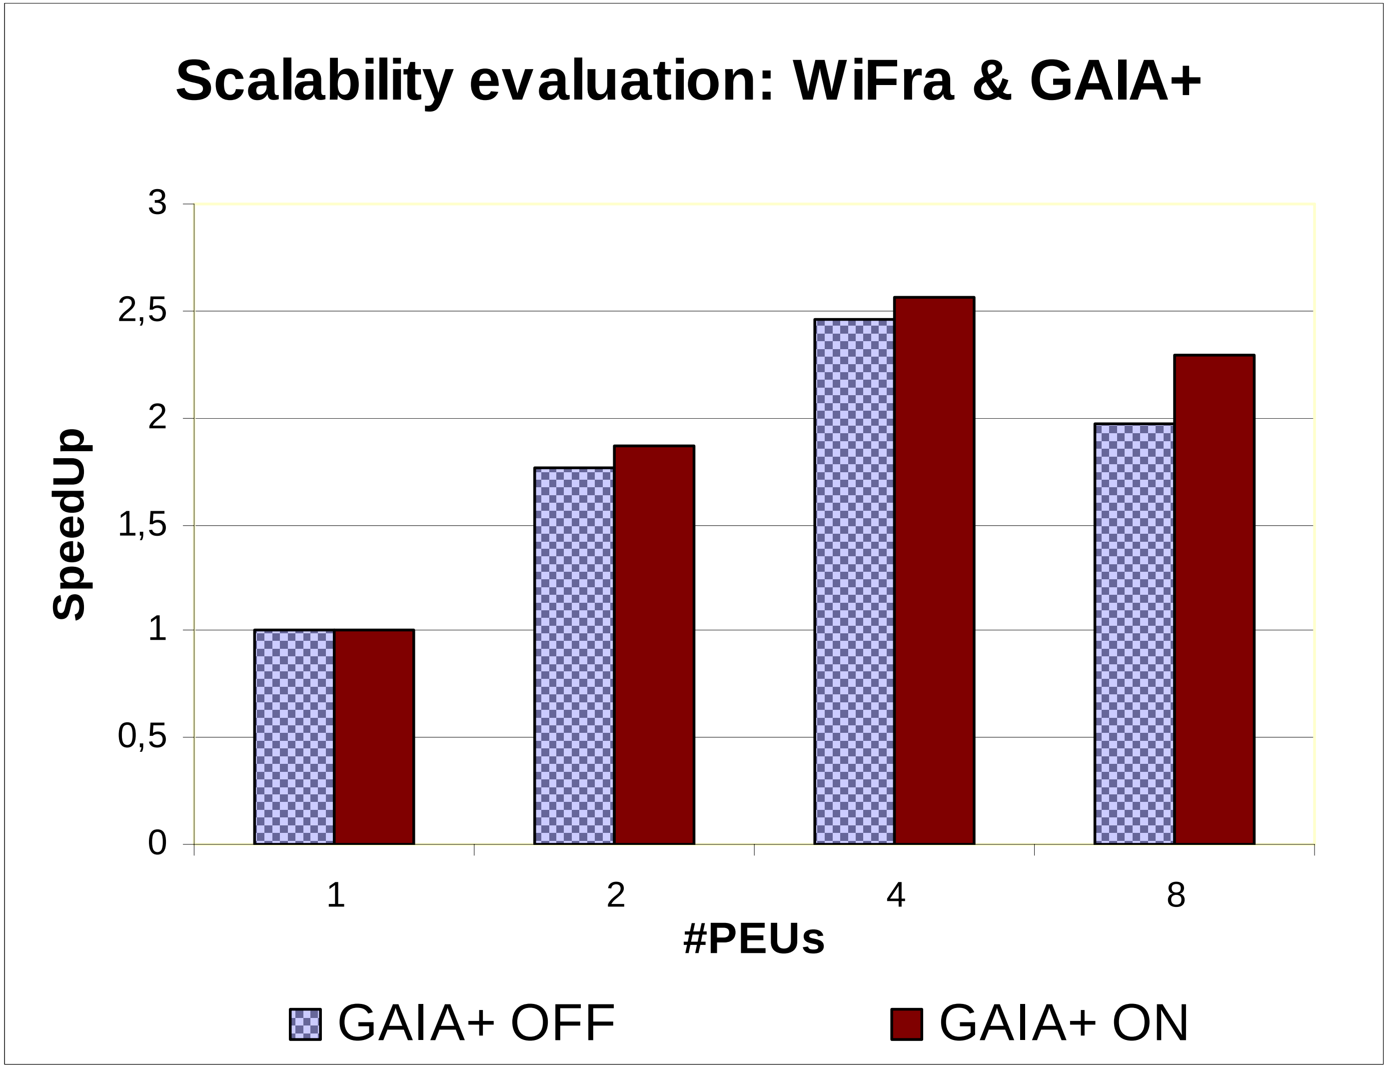

Experimental evaluation: massively populated scenario

- Very large-scale wireless network composed of **1.000.000** nodes

- Execution architecture:
  - **32** PEUs: Dual Core Intel Pentium IV CPU 3 GHz 1GB RAM, interconnected by Fast Ethernet
  - Each PEU initially allocates 32.250 SMHs

- During the simulation the GAIA+ mechanism triggers re-allocations to adapt the load of each PEU to the performance of the hardware and background load (e.g. other running tasks)

Experimental evaluation: massively populated scenario

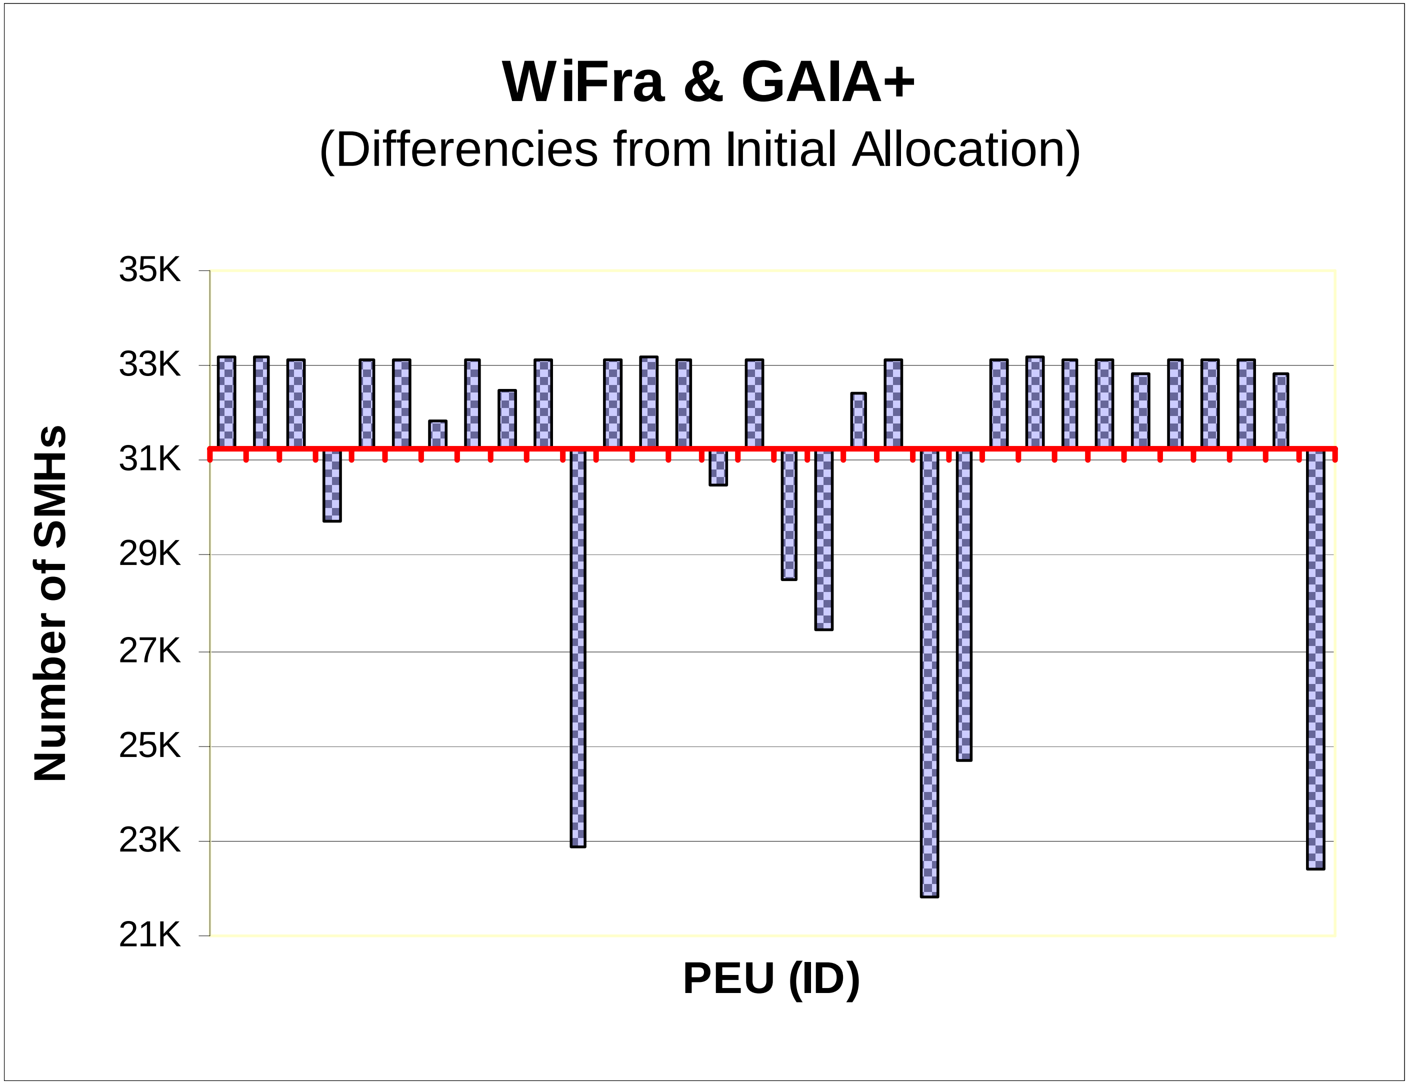

Experimental evaluation: massively populated scenario

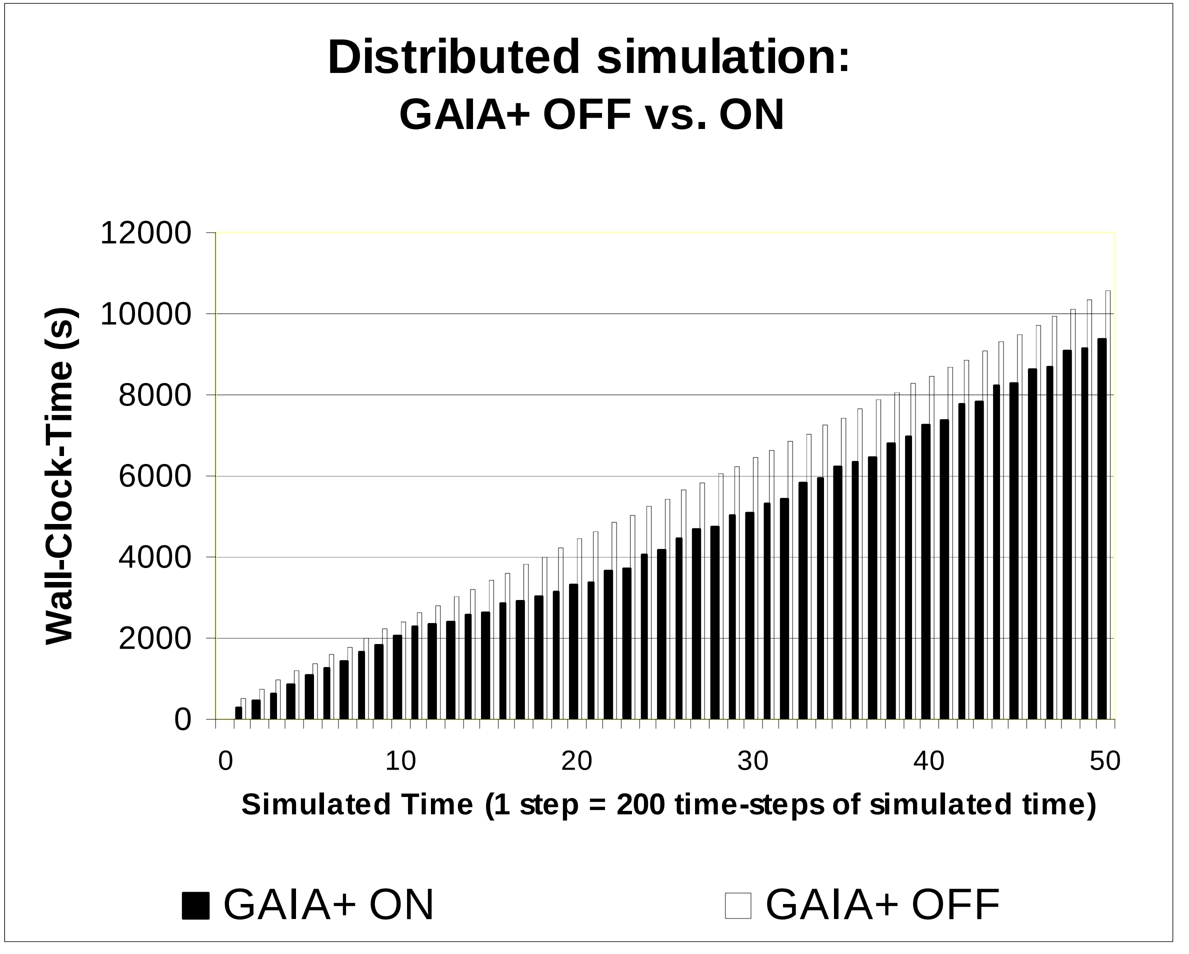

The gain in terms of performance is up to 21%

More complex user level protocols would increase this gain

Conclusions and future work

- Distributed simulation is a feasible approach for the detailed (fine-grained) simulation of large-scale wireless networks
- An approach based on adaptive migration of the simulated entities can both enhance the communication and computation load balancing, therefore increasing the simulation speed-up

- We plan to add more protocol models to WiFra
- Further develop the GAIA+ framework to adaptively activate more CPUs (on-demand) and to reduce the number of used CPUs when are not necessary, therefore reducing the synchronization overhead

Detailed Simulation of Large-Scale Wireless Networks

Gabriele D'Angelo

joint work with

Michele Bracuto

University of Bologna
Department of Computer Science

DS-RT 2007 – Crete Island (Greece)