Advanced RTI System (ARTÌS)

A new middleware for parallel and distributed simulation

Gabriele D'Angelo
Michele Bracuto
Luciano Bononi

{gda, bracuto, bononi}@cs.unibo.it

http://www.cs.unibo.it/~gdangelo
http://www.cs.unibo.it/~bracuto
http://www.cs.unibo.it/~bononi

Dipartimento di Scienze dell'Informazione
Università degli Studi di Bologna

Presentation outline

- Simulation of large scale and complex models
- The ARTìS software architecture
- Simulation of massive, complex and dynamic systems
  - High Performance Computing (HPC)
- Gaia framework: entities migration
- Examples: Ad hoc and Sensors network models
- Concurrent Replication of PADS (CR-PADS)
- Further optimizations for PADS
- Conclusions and future work

Simulation of large scale and complex models

- Simulation models currently of interest may include a potentially huge number of simulated objects
- Large scale and complex simulation models may be unpractical to simulate on a single-processor execution unit: **huge memory requirements**, large amount of **time** required to complete the simulation runs
- The **memory** bottleneck reduction, **scalability** and **speed-up** can be achieved by using *parallel/distributed models and execution architectures*

- Goal: to **increase the simulation speed**, reduce the **Wall Clock Time (WCT)** required to complete the simulation runs

The ARTìS software architecture

Advanced RTI System (ARTìS), parallel and distributed simulation middleware:

- Model components' heterogeneity, distribution and reuse
- Adaptive evaluation of the communication bottlenecks and dynamic adaptation of the inter-process communication layer
- Generic Adaptive Interaction Architecture (GAIA): model components' migration mechanism to support load balancing and data distribution management (DDM) overhead reduction
- Support for High Performance Computing clusters (HPC): scalability evaluation of the middleware
- Concurrent Replication of Parallel and Distributed Simulation (CR-PADS) and cloning

ARTìS: logical architecture

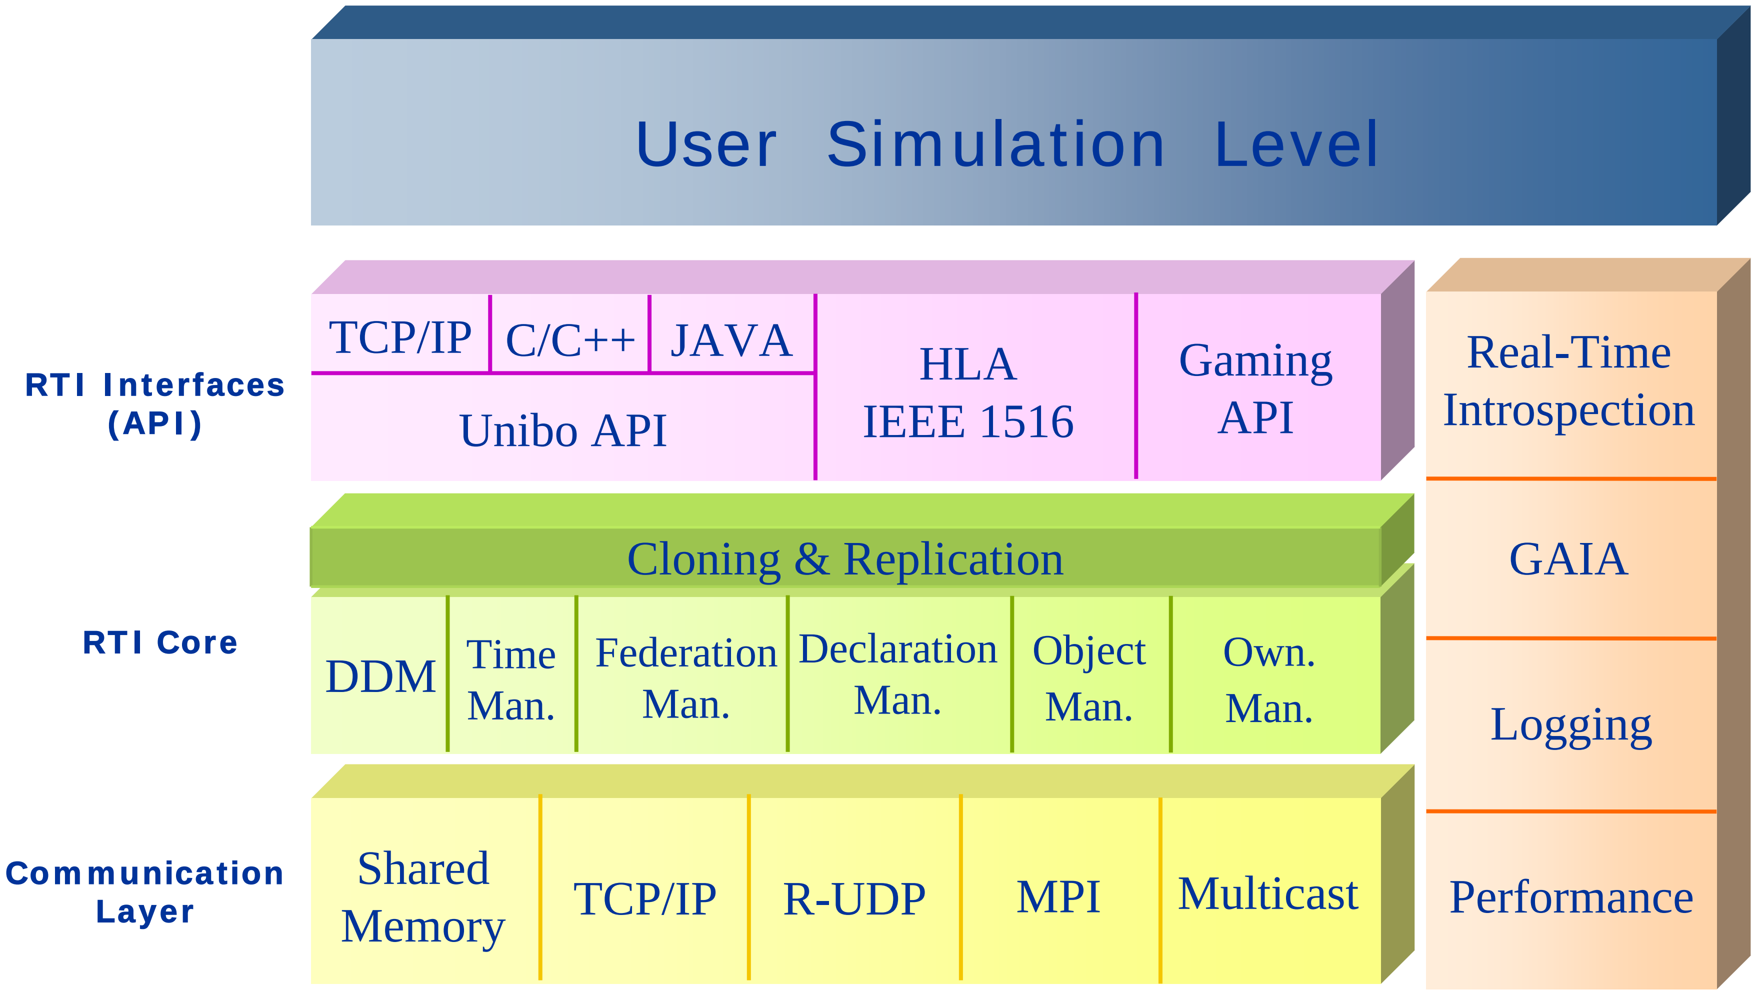

Communication architecture

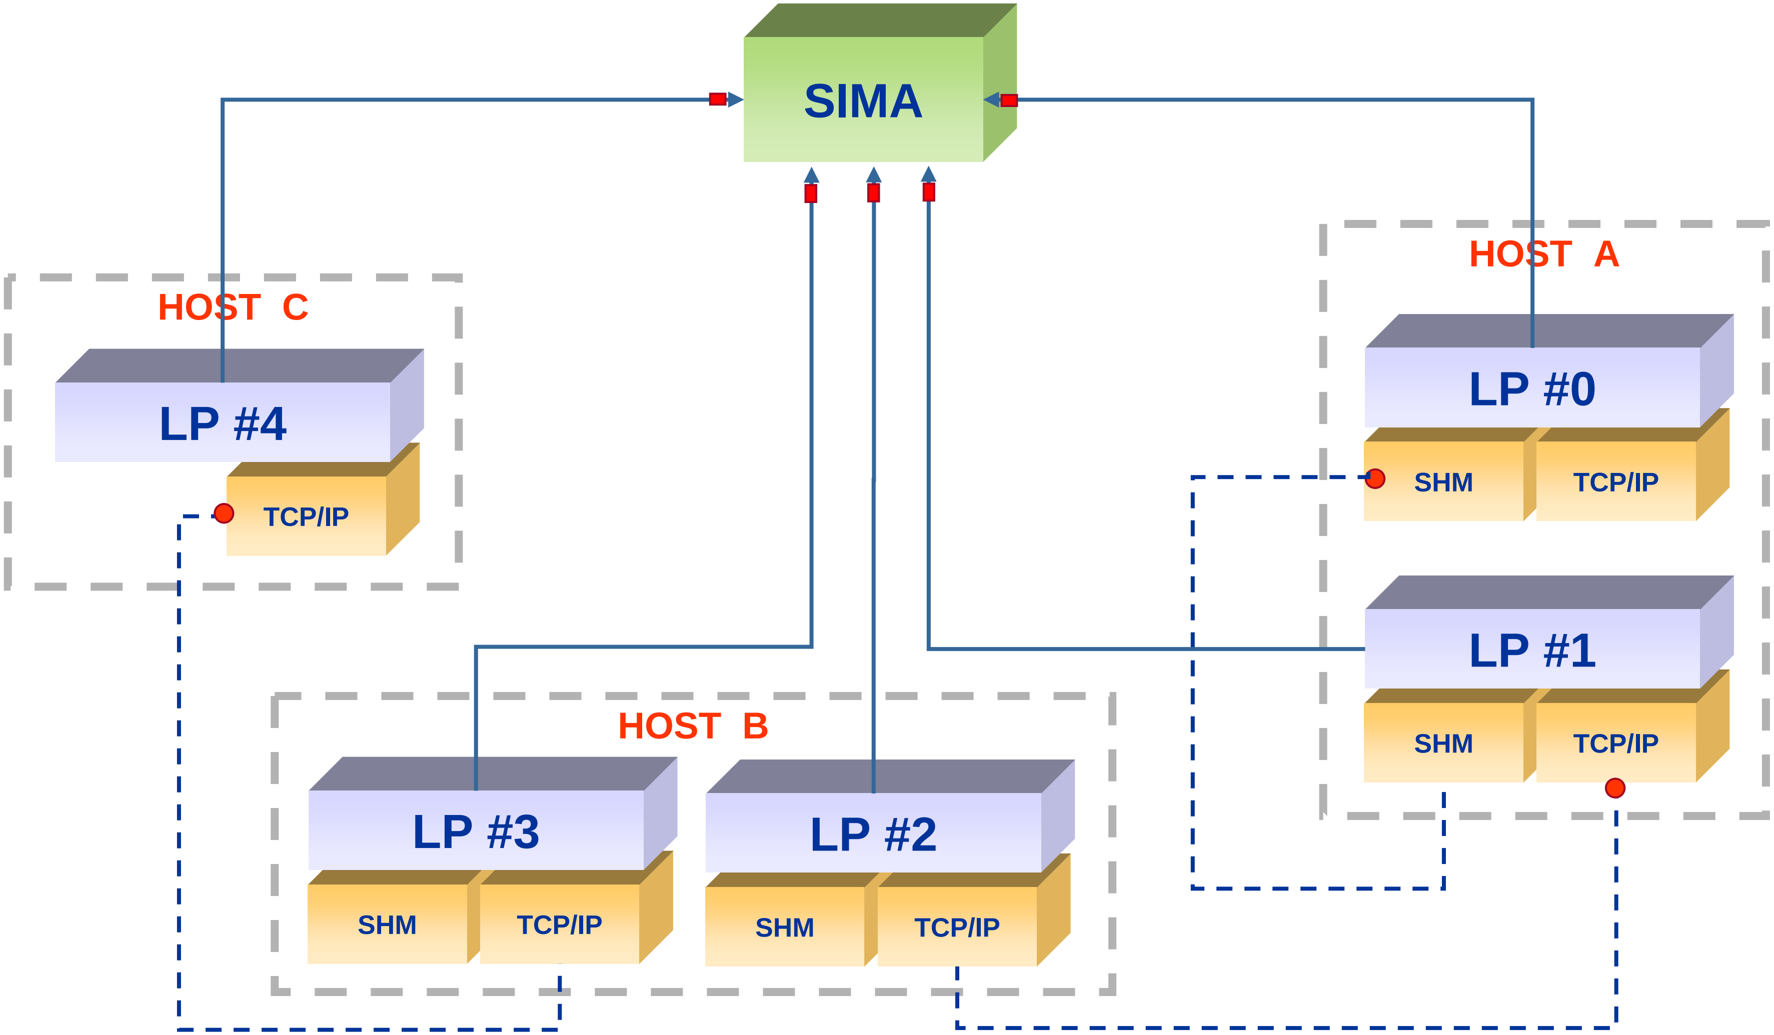

Performance evaluation: Ad-Hoc wireless network model

Simulated model:

- A set of **Simulated Mobile Hosts** (SMHs)
- Mobility model:
  - Random Mobility Motion model (RMM)
  - uncorrelated SMHs' mobility
- Traffic model:
  - ping messages (CBR) by every SMH to all neighbors within the wireless communication range (250 m)
- Propagation model
  - open space (neighbor-SMHs within detection range)

Ad-Hoc network model characterization

Computation and communication issues:

- The **computation** required for each SMH per time-step is in the order of **O(#SMH^2)**: to determine the neighbor set
- The **communication** required among SMHs is in the order of **O(K*#SMH)** per time-step, with K defined as a constant value based on SMHs density (assumed as constant)

Scalability evaluation: High Performance Computing

IBM CLX/1024 – IBM Linux cluster 1350 - CINECA

- 512 2-way nodes (IBM X335)
  - 768 Xeon Pentium IV, 3.06 GHz
  - 256 Xeon Pentium IV EM64T (Nocona), 3.00 GHz
  - 2 GB of RAM for each node
  - Global peak performance: 6.1 TFlops

- All the nodes are interconnected by a low latency Myrinet network, maximum bandwidth between each pair of nodes: 256 MB/s

ARTìS on High Performance Computing (HPC)

1 million of Simulated Mobile Hosts (*1 Timestep of simulated time*)

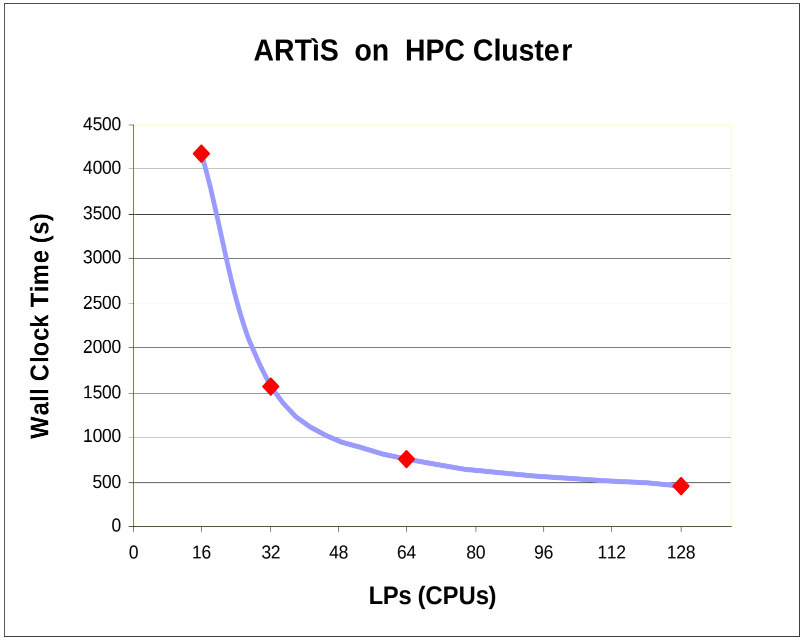

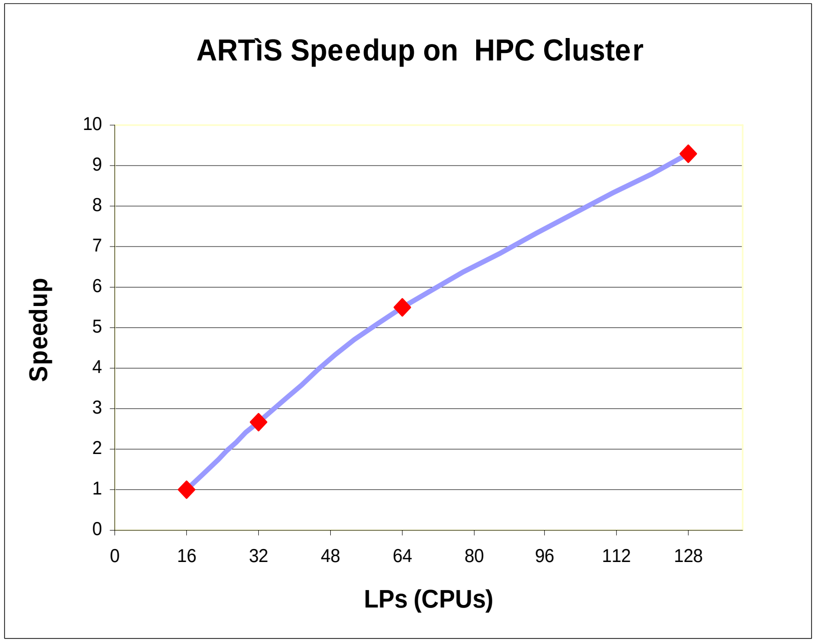

Mobile and Wireless Networks' model characteristics

- **"Open broadcast" nature of the wireless transmissions**
  - "space-locality" of causality between neighbor-hosts
  - neighbor-hosts should be notified about transmission events anyway, e.g. to model interference, detection, MAC, etc.
- **Wireless devices can be mobile**
  - the set of neighbor-hosts change as simulated time elapses
- **Communication between hosts is "session-based"**
  - determines a "time-locality" effect
  - the set of neighbor-hosts is interested by transmission events, for a significant time-window

The group of model entities in the shared causal-domain can be highly dynamic:

high degree of communication is required to maintain full synchronization

GAIA framework: entities migration

Wireless ad hoc network scenario: (evaluating migration of SMH x)

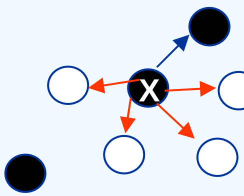

model entities of SMHs executed over PEUs A and B

Host X executed over PEU A "transmission-event"

Physical Execution Units for the simulation

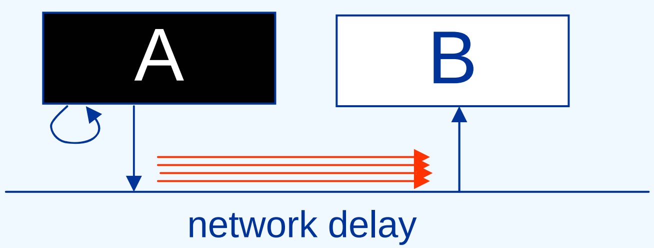

X's "transmission-event" must be notified to the 4 model entities executed over B

SMH = Simulated Mobile Host

PEU = Physical Execution Unit

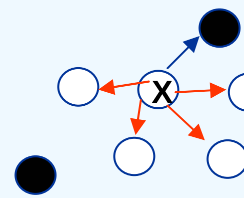

model entities of SMHs executed over PEUs A and B

Same scenario after migration of X from A to B

Physical Execution Units for the simulation

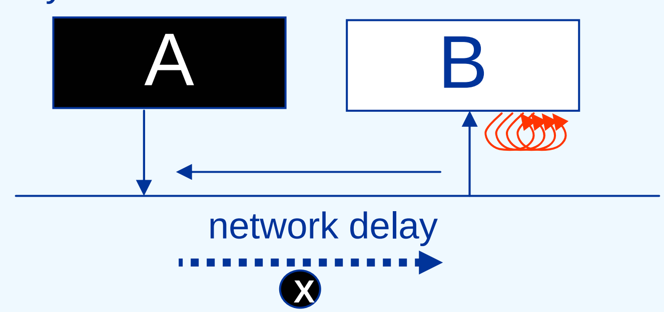

After X's migration, X's "transmission-event" must be notified to one model entity executed over A

Example: Ad-Hoc network model implementation

Modeling issues:

- A set of Simulated Mobile Hosts (SMHs)
- Mobility model:
  - Random Mobility Motion model (RMM)
  - Fast- and Slow-RMM (100, 25, 10 m/s)
  - uncorrelated SMHs' mobility (worst case)
- Traffic model:
  - ping messages (CBR) by every SMH to all neighbors within the wireless communication range (250 m)
  - low local computation model (worst case)
- Propagation model
  - open space (neighbor-SMHs within detection range)

Ad hoc network: migration mechanism “off” and “on”

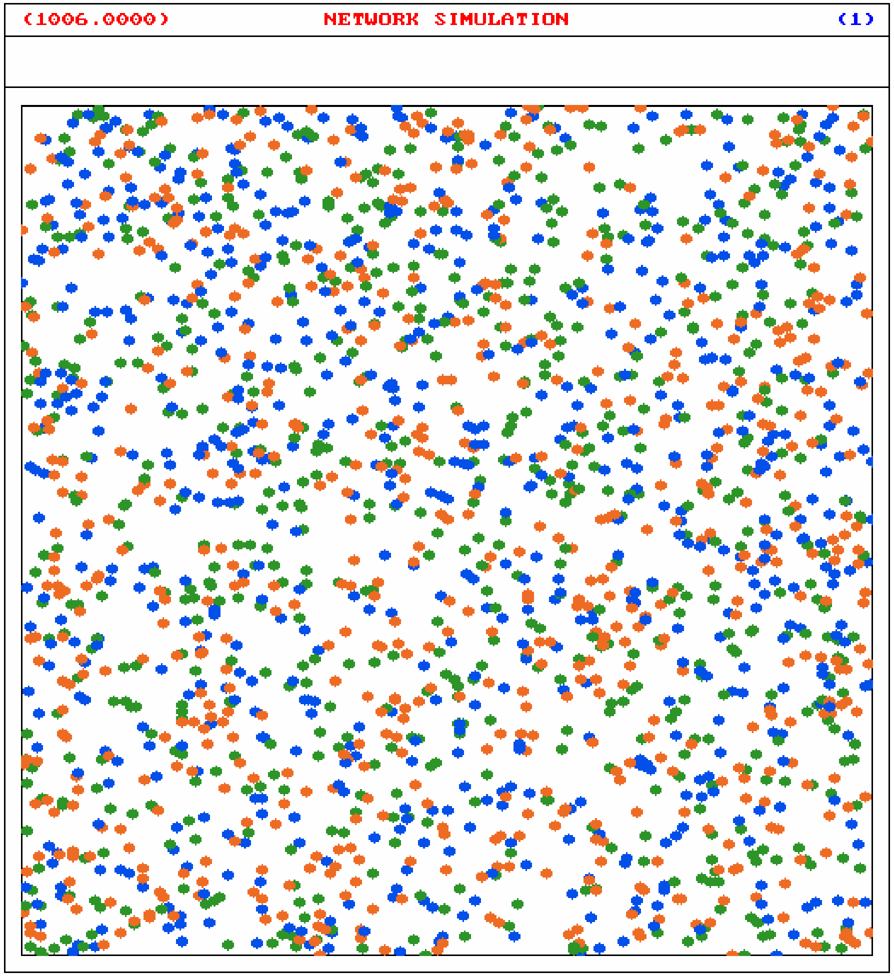

Migration “OFF”

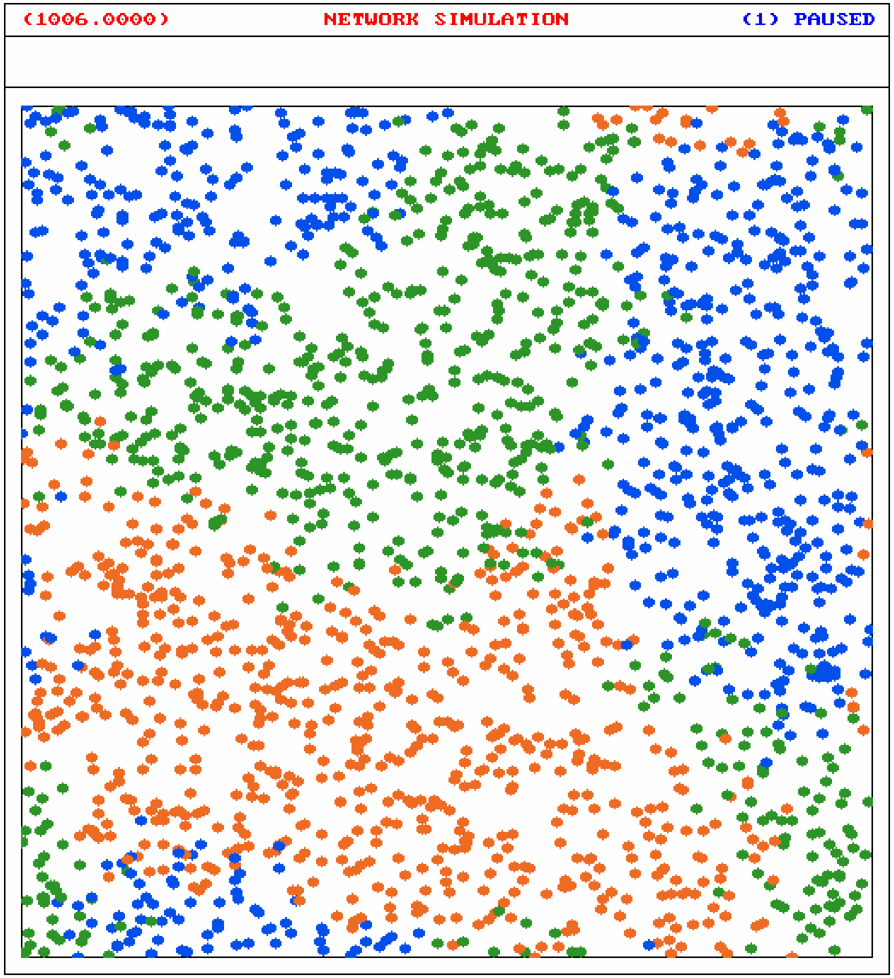

Migration “ON”

Performance analysis: number of migrations / timestep

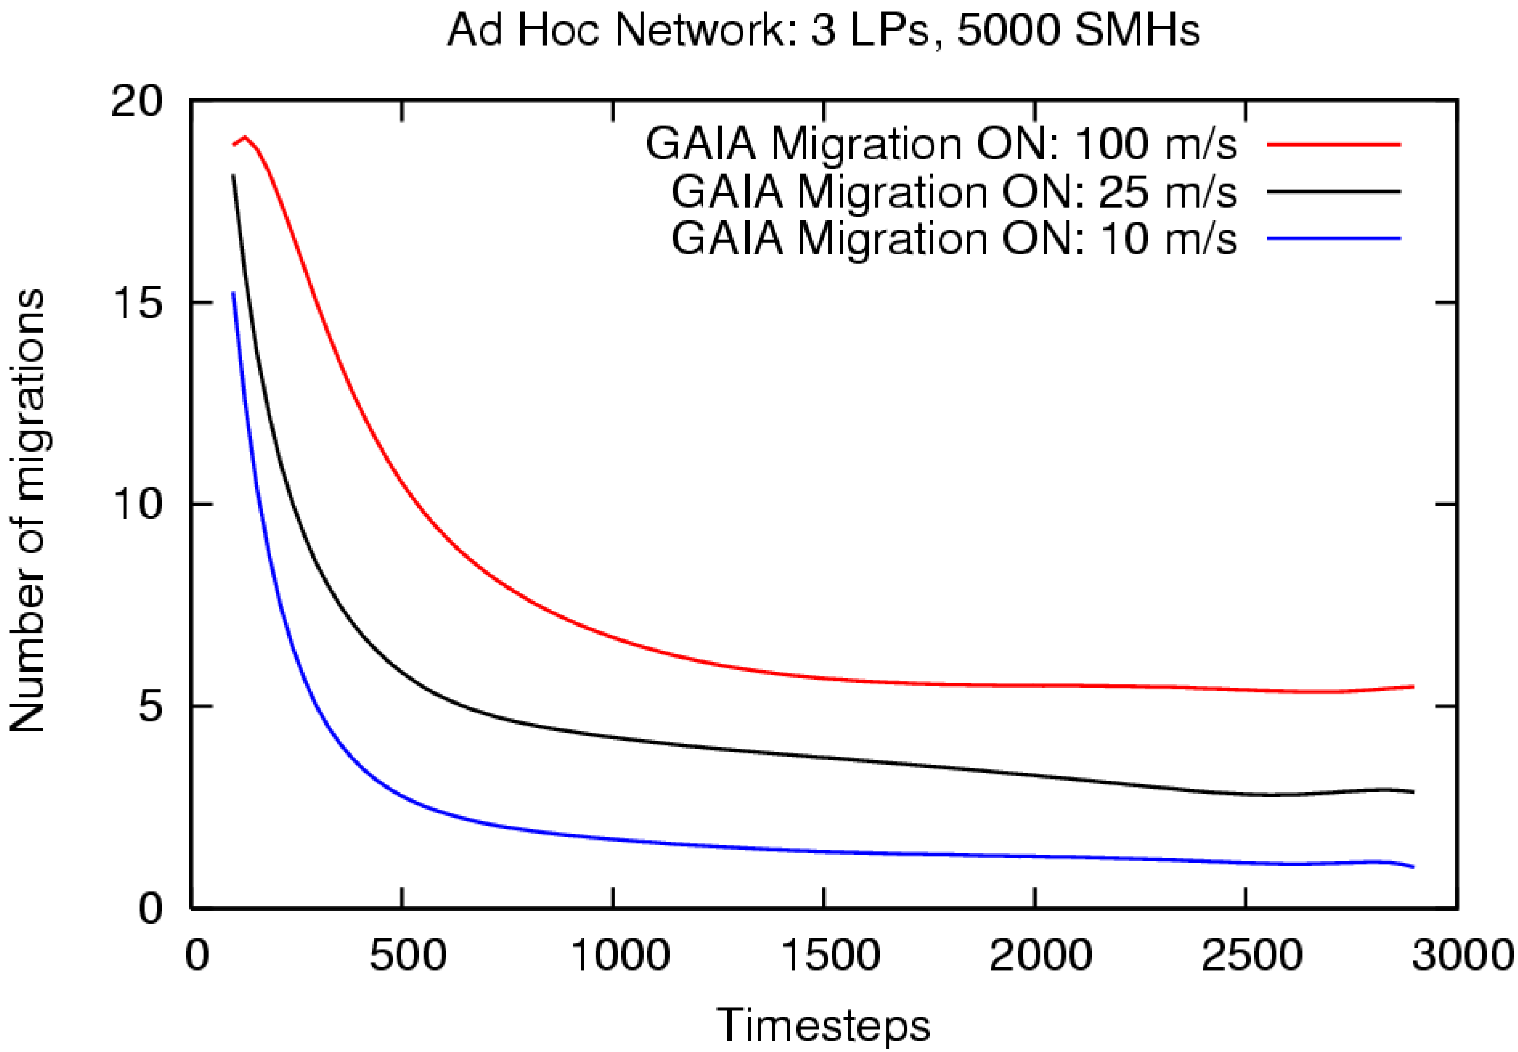

Performance analysis: Local Communication Ratio

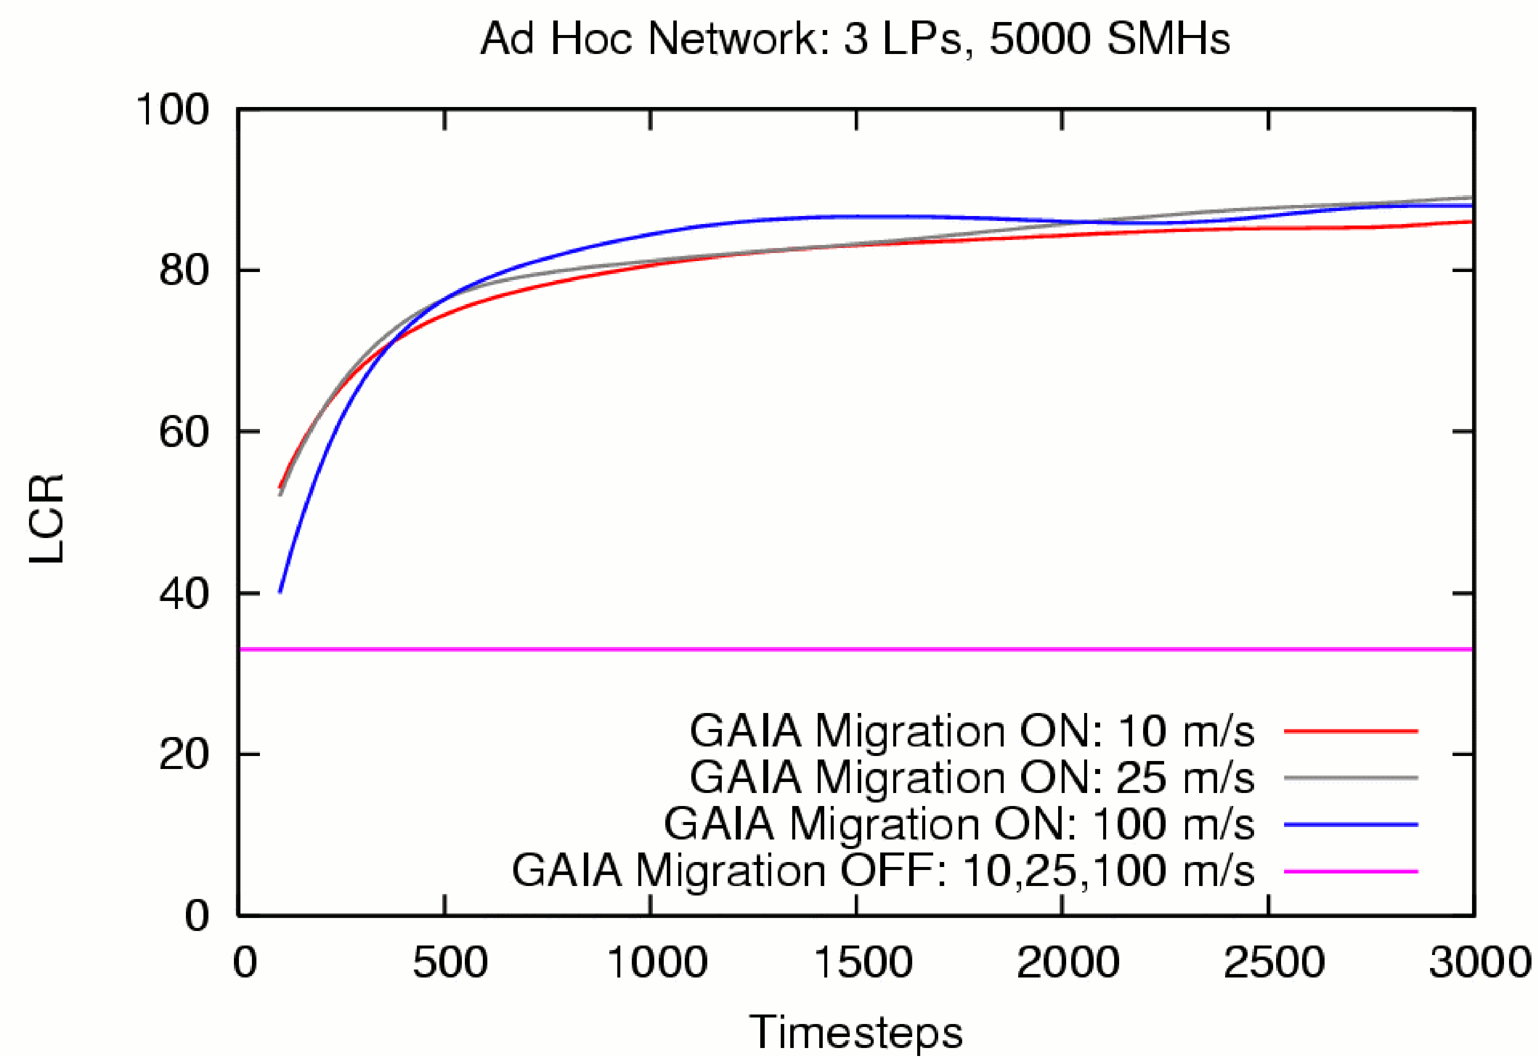

Performance analysis: speed-up

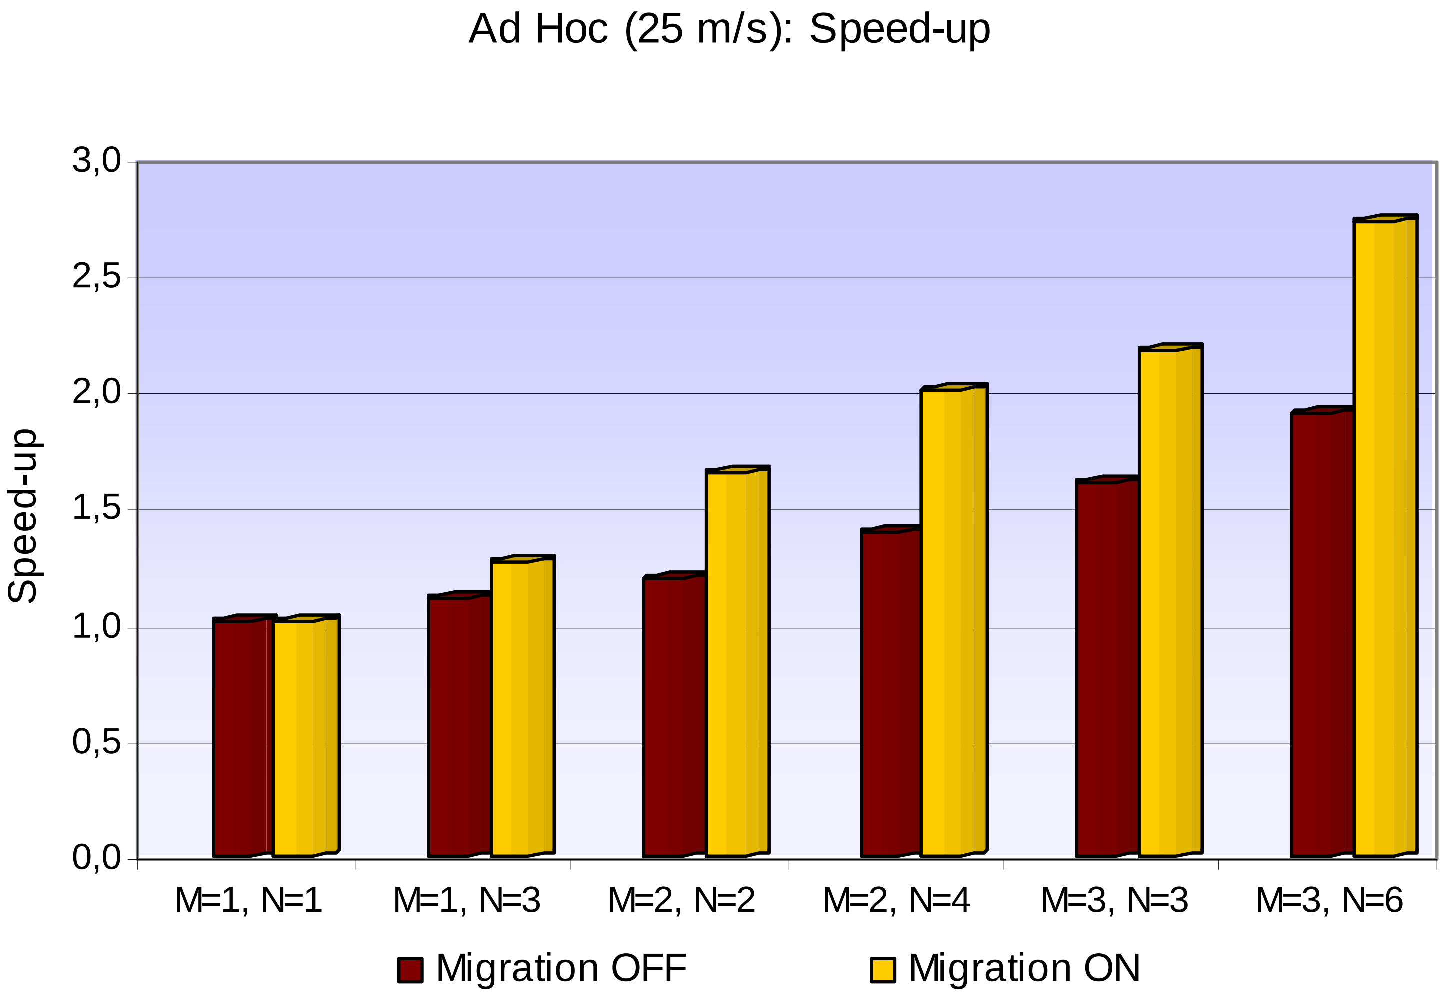

M = # PEU, N = # LP

Example: sensor network model implementation

Design of a new energy aware Medium Access Control protocol:

- Up to 40.000 sensors, limited battery resources
- Sensors are static (no mobility model)
- Every sensor implements the MAC protocol, no centralization
- 4 different sensor states (active, power saving, listening, died)
- Every sensor implements a “pressure variation” detector and sends broadcast alerts flooding toward a set of detection points

- The energy aware MAC protocol increases the network lifetime managing the sensors’ state (dinamically switching to power-save state the sensors within areas covered by other active sensors)

Sensor network simulation example

1000 sensors (for this example),

Limited battery resources

4 different states:

- **red = awake**
- **green = power saving**
- **white = listening**
- **black = died**

Pink circle =

active transmitting range

Blue circle =

alert-message transmission

GOAL: extend the network lifetime maintaining network connectivity

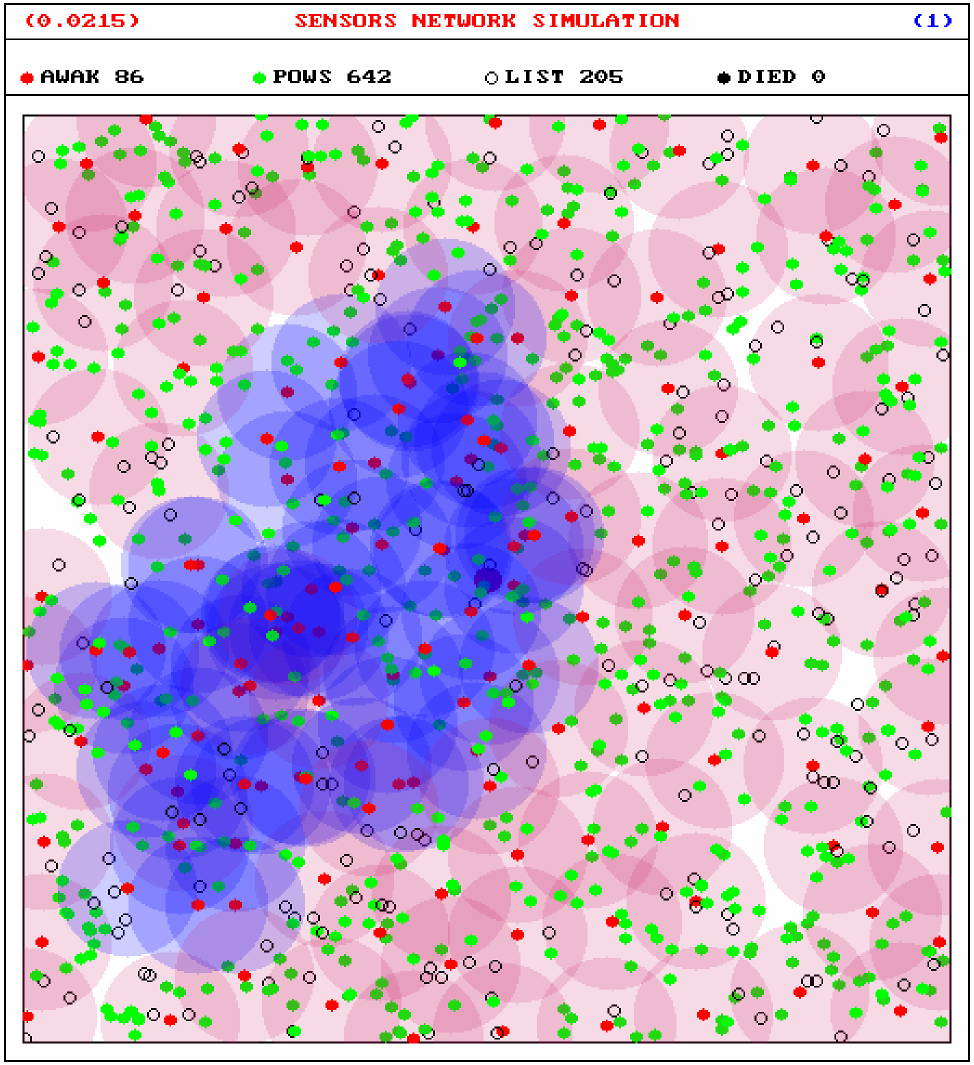

Performance analysis: speed-up

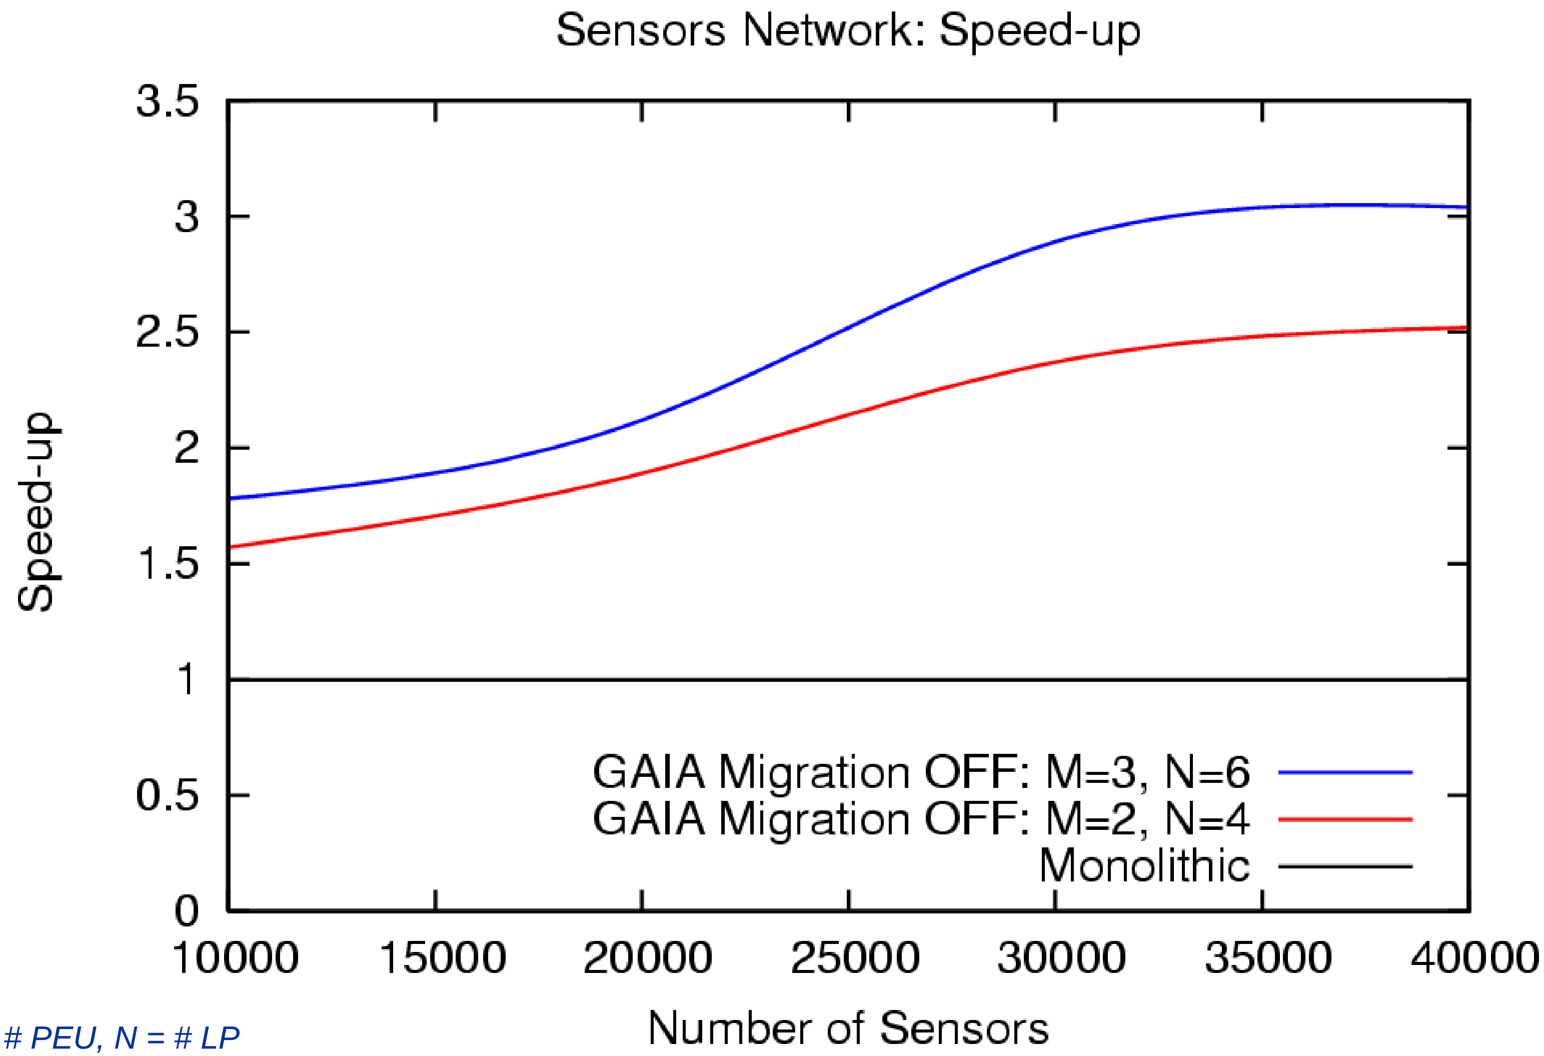

M = # PEU, N = # LP

Conclusion and Future work

- ARTìS is a scalable, optimized parallel and distributed simulation middleware, used to simulate dynamic complex systems

- Careful and enhanced evaluation of the proposed optimizations

- Further improve the ARTìS middleware
  - Data structures optimization (event-management)
  - GAIA: detailed load balancing and communication heuristics
  - IEEE 1516 full compatibility
  - New models (es. Detailed MAC protocols)
  - Migration based middleware for Internet games

Advanced RTI System (ARTìS)

A new middleware for parallel and distributed simulation

Gabriele D'Angelo
Michele Bracuto

{gda, bracuto}@cs.unibo.it

http://www.cs.unibo.it/~gdangelo
http://www.cs.unibo.it//~bracuto

Dipartimento di Scienze dell'Informazione
Università degli Studi di Bologna