Adapting a Multi-Agent Soil Simulation on GPU

Journée RGE - Belfort

Guillaume Laville

FEMTO-ST/DISC - Université de Franche-Comté

Feb 2013

Overview

Multi-Agent Systems

Used to simulate complex models

Natural processes, social behaviours...

... based on entities

- Agents (passive ou active actors)
- Environment (store the global properties of the model)
- Behaviours and interactions (often message-based)

Multi-Agent Systems

Generic platforms for MAS simulation

Repast, Madkit, JADE, Netlogo...

... most of these use traditional CPUs

- Sequential execution
- Shared memory parallelism, using IPC or threads
- Distributed unto multiple host, using native sockets or libraries such as MPI

Multi-Agent Systems

Problematic

Could this type of simulation be executed on one or multiple GPU, using generalised rules of mapping, to exploit the inherent parallelism of this new architecture ?

Logic

Would allow faster execution and/or better scaling on readily available GPU, available at a low cost for any researcher, as well as usage of a new type of hardware platform emerging on supercomputers ?

About the Sworm model

Overview

- Simulation of soil evolution
- Written in Java using the Madkit platform
- Developped for the IRD

Soil considered as a multi-scale agent-based simulation

- Macroscopic ones (earth worms)
- Microscopic ones (aerobic bacterias) $\Rightarrow$ **Mior**

Sworm multi-scale model

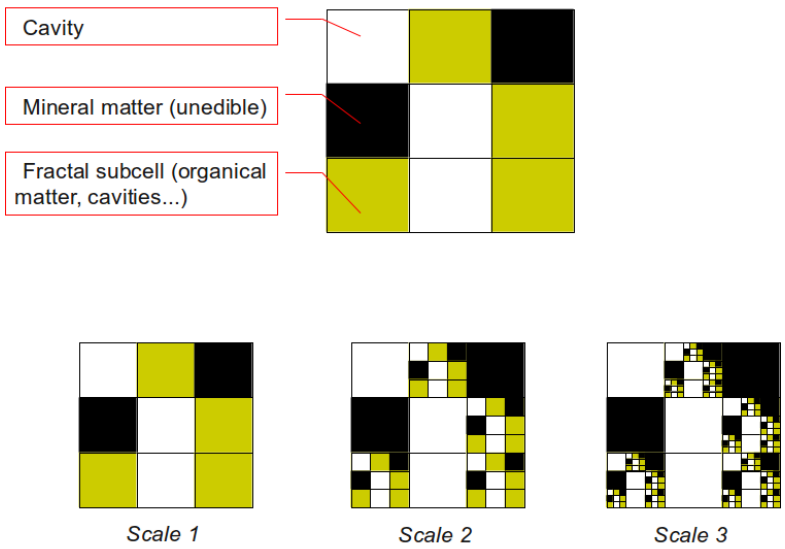

FIGURE: Sworm multi-scale representation

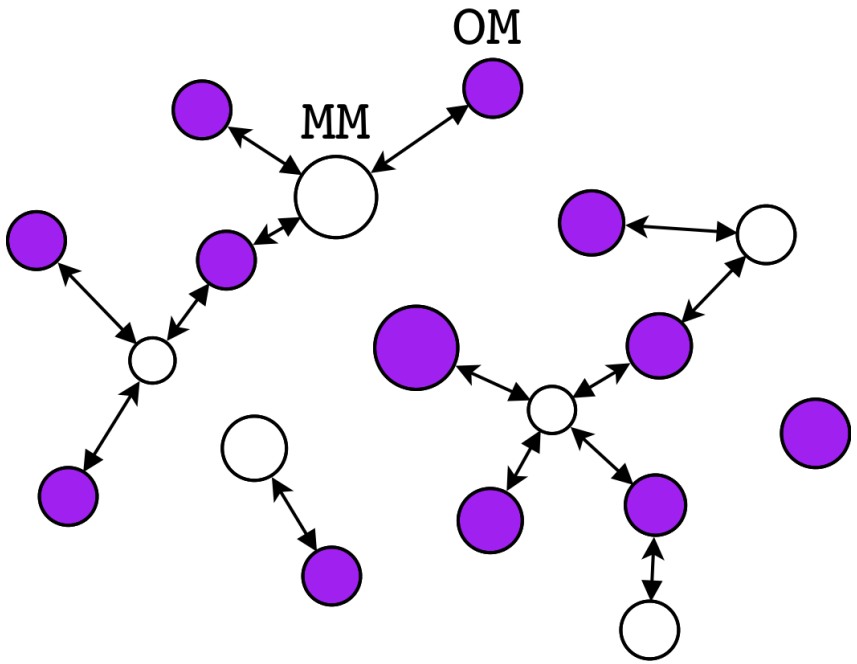

FIGURE: MIOR representation

Microscopic scale : MIOR simulation

Based on two kind of agents

- Microbial colonies (called `MM`)
- Carbon deposits (called `OM`)

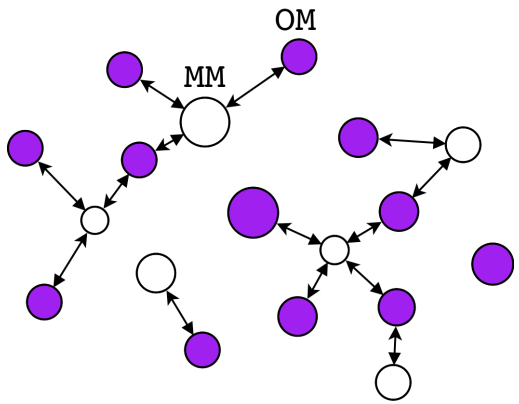

Two main evolution processes linked to the `MM`

- Breathing
- Growth

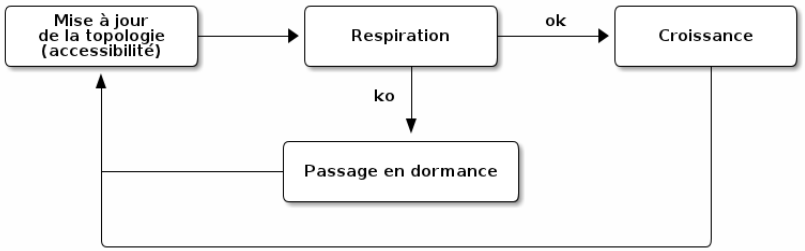

Overview of a graphic card

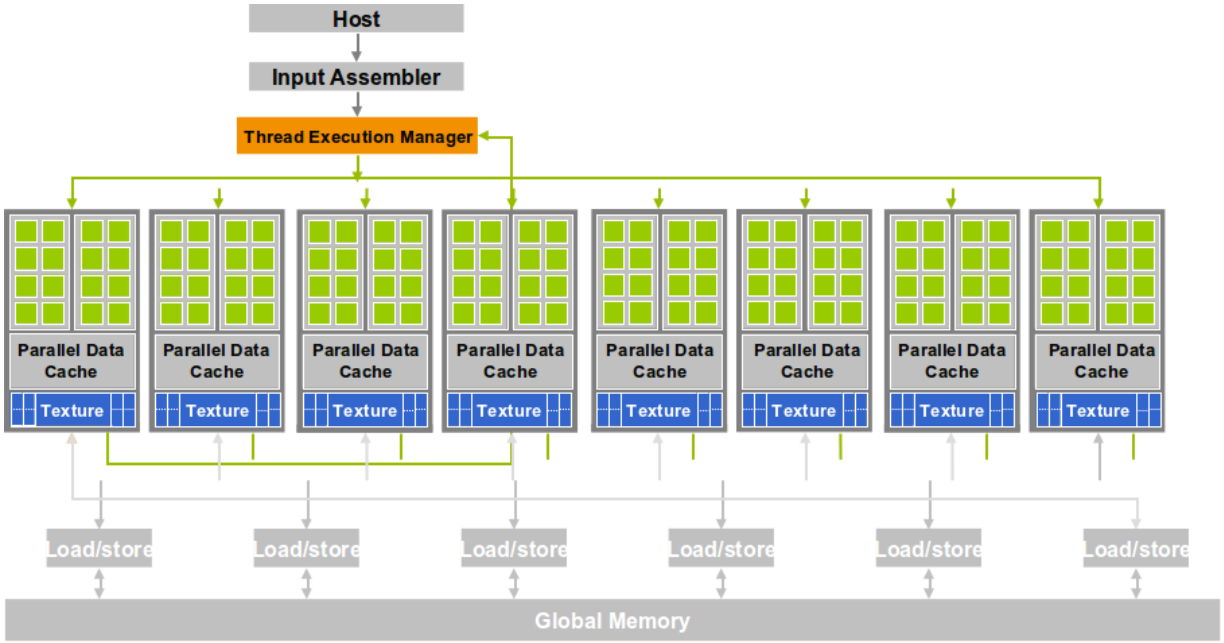

What is GPGPU ?

Principe

Use modern, programmable graphic hardware as a generic execution platform for parallelism.

Two main languages available

- CUDA (NVIDIA)
- OpenCL (AMD, NVIDIA, Intel, Khronos Group)

Multipe Java bridges

- JCUDA (www.jcuda.org)
- **JOCL (www.jocl.org)** $\Rightarrow$ Retained solution
- JavaCL, LWJGL, JOCL (jogamp.org)...

Model mapping on OpenCL concepts

Model	OpenCL
Agent properties	C-like Structures
Agent actions	Device-local functions
Simulation step	OpenCL kernel

Four main data structures

- `MM` array
- `OM` array
- Topology matrix
- Environment (MiorWorld)

```
 5 typedef struct MM {
 6     float   x;
 7     float   y;
 8     int     carbone;
 9     int     dormance;
10 } MM;
11
12 typedef struct OM {
13     float   x;
14     float   y;
15     int     carbone;
16     int     lock;
17 } OM;
18
19 typedef struct MiorWorld {
20     int     nbMM;
21     int     nbOM;
22     int     RA;
23     float   RR;
24     float   GR;
25     float   K;
26     int     width;
27     int     minSize;
28     int     CO2;
29     int     lock;
30 } MiorWorld;
31
32
33 typedef struct RandomState {
34     ulong a;
35     ulong b;
36     ulong c;
37 } RandomState;
```

Execution of multiples agents at the same time

Topology

Each GPU thread is associated to an (OM, MM) couple.

Carbon distribution

Each GPU thread is associated to an OM agent.

Breathing/growth

Each GPU thread is associated to an MM agent.

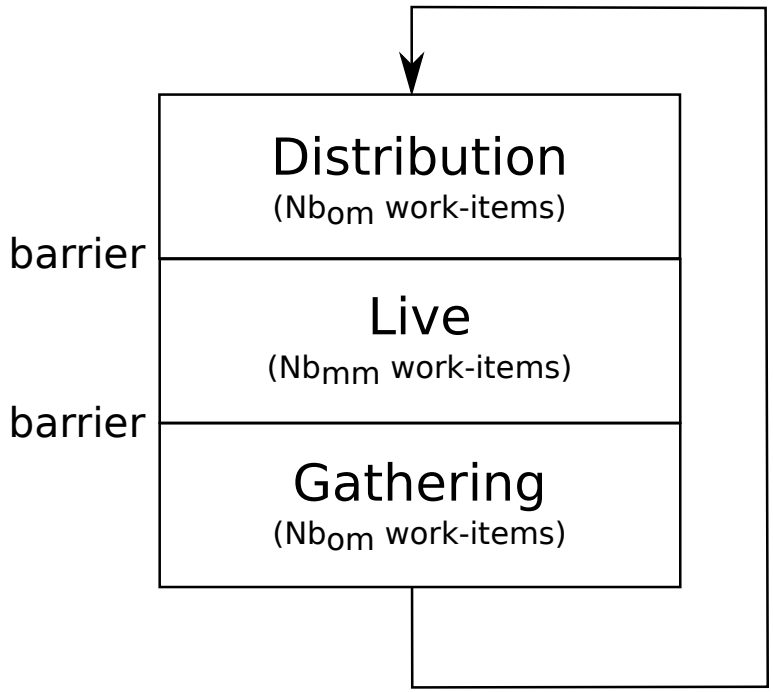

FIGURE: Execution decomposition

Data structures optimizations

Topology

- two-dimension low density matrix
- walkthought-optimized matrix

Carbon parts

- parts in global device memory
- parts cached in local memory

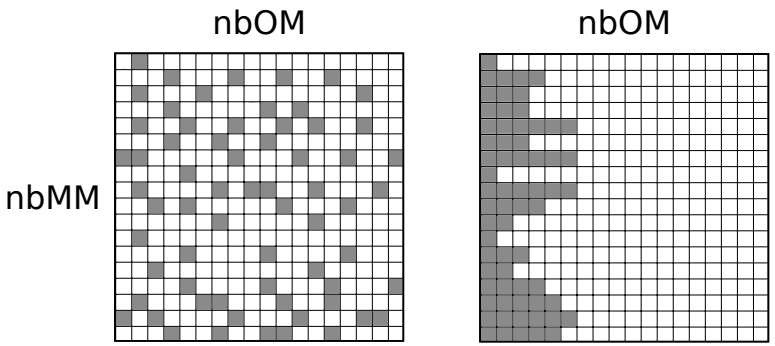

FIGURE: Topology representations

Results for one simulation

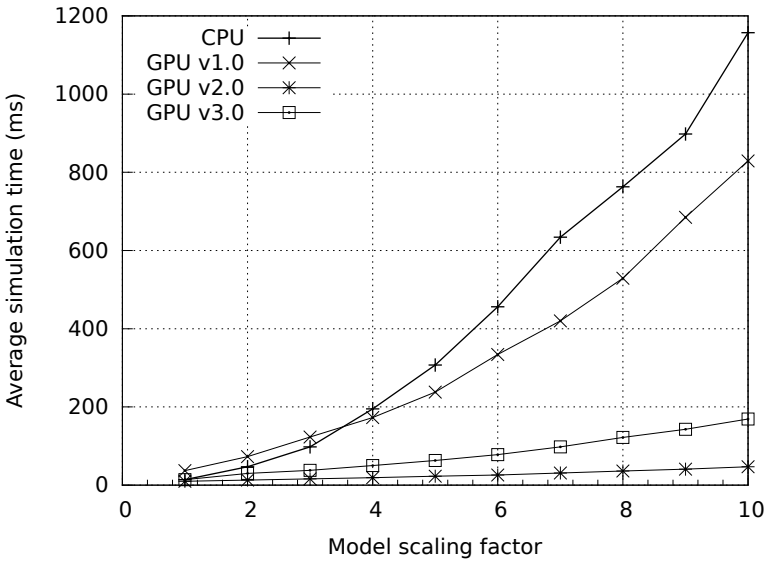

FIGURE: NVIDIA GeForce 8800GT

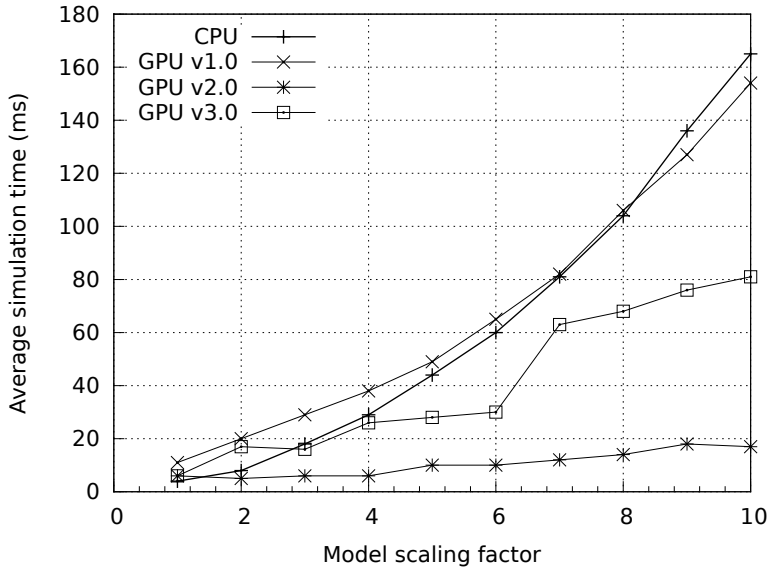

FIGURE: NVIDIA Tesla 1070

Results for a fixed number of simulations

- fixed number of simulations to execute
- varying the number of execution launched in one GPU execution

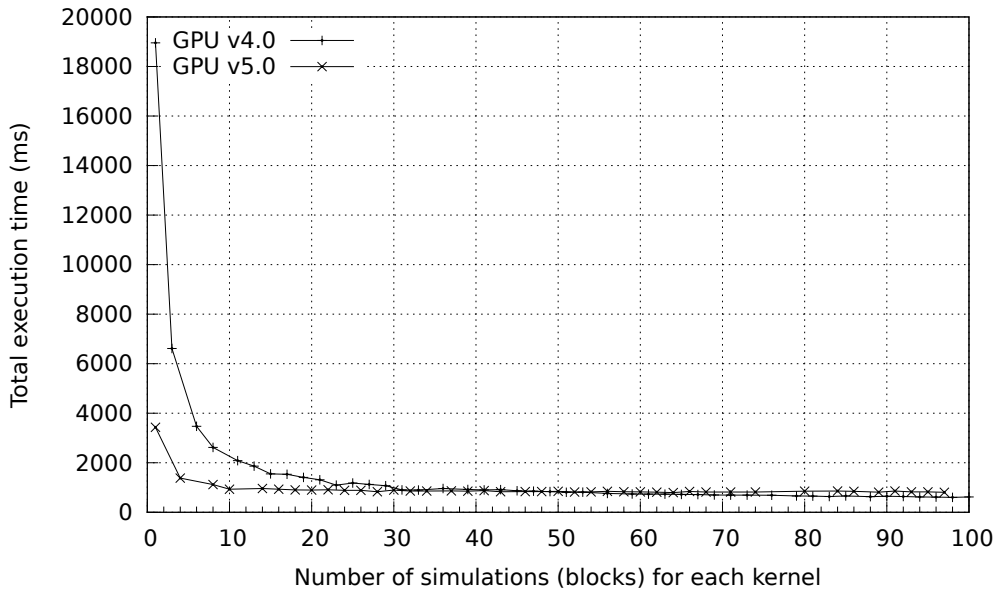

FIGURE: Execution time of 1000 simulations

Future works

- Synchronize MIOR execution on GPU with the existing Sworm CPU computations.
- Allow the distribution of GPU computations on multiple GPUs across multiple nodes.
- Propose a generic execution model for computing-intensive, environment-wide model evolutions in ABMs.

Future works

- Synchronize MIOR execution on GPU with the existing Sworm CPU computations.
- Allow the distribution of GPU computations on multiple GPUs across multiple nodes.
- Propose a generic execution model for computing-intensive, environment-wide model evolutions in ABMs.

Future works

Generic GPU execution platform

- Provide generic implementations of common MAS algorithm (diffusion, computation of distances in 2d and 3d, matrices reduction...)
- Modular : Each set of functions regrouped in modules
- Extensible : Possibility to add modules to treat new problems.

Thanks you

- Questions ?