

XDEM: from HPC to the Cloud

Xavier Besseron

LuXDEM Research Centre

<http://luxdem.uni.lu/>

BusinessClouds 2017

January 24, 2017

Outline

- 1 What is XDEM?
- 2 XDEM for HPC
- 3 XDEM in the Cloud: XDEMaas
- 4 Cloud vs In-House HPC: A Cost Study using XDEM
- 5 Conclusion

eXtended Discrete Element Method

What is XDEM?


What is XDEM?

XDEM software is **multi-physics simulation** toolbox modeling granular materials and processes




Particle Motion

Snow, Sand, ...



Chemical Reactions


Coke, Iron ore, Biomass, ...




+ coupling external libraries:

CFD with OpenFoam, **FEM** with Diffpack


Examples




Tire on Snow
XDEM-FEM coupling



Heat transfer to the walls
of a rotary furnace
XDEM-CFD coupling



Charge/Discharge of hoppers




Impacts on an Elastic Membrane
XDEM-FEM coupling

XDEM-HPC:

Parallel Implementation of XDEM

XDEM needs HPC!



- 427,878 particles
- 6,742 triangles
- Impact
- Rolling friction


Hopper charge

- 15 s of simulation
- 92 hours with 120 cores
- Est. seq. time > 4 months

Hopper discharge

- 18 s of simulation
- 120 hours with 144 cores
- Est. seq. time > 6 months

XDEM needs HPC!



- 427,878 particles
- 6,742 triangles
- Impact
- Rolling friction


Hopper charge

- 15 s of simulation
- 92 hours with 120 cores
- Est. seq. time > 4 months

Hopper discharge

- 18 s of simulation
- 120 hours with 144 cores
- Est. seq. time > 6 months

XDEM needs HPC!



- 427,878 particles
- 6,742 triangles
- Impact
- Rolling friction

Hopper charge

- 15 s of simulation
- 92 hours with 120 cores
- Est. seq. time > 4 months

Hopper discharge

- 18 s of simulation
- 120 hours with 144 cores
- Est. seq. time > 6 months

Parallel XDEM simulation for HPC

XDEM is computation intensive and uses a lot of memory




⇒ Parallel and distributed implementation

- Aggregates memory of many computing nodes
- Benefits from speedup of many computing cores

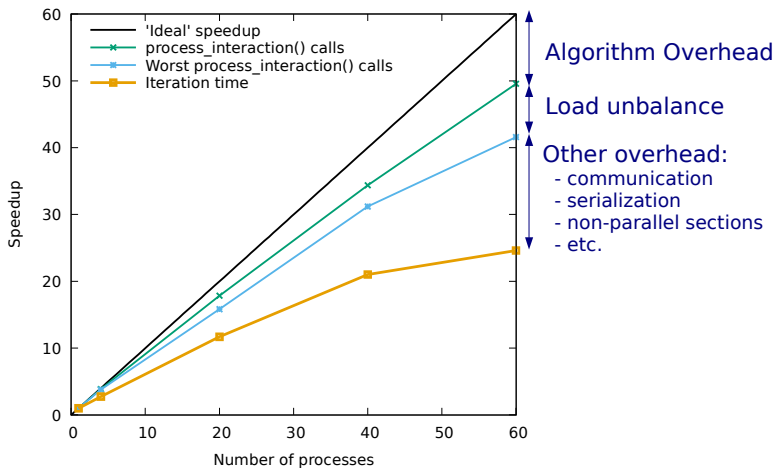
Periodic Load balancing

- based on Orthogonal Recursive Bisection (ORB)
- adapt to changes in particle distribution

Communications based on MPI



XDEM Performance (Work-in-Progress)



XDEM-HPC: Work-in-Progress / Future Work

Better Load balancing algorithms

- Zoltan, RIB, RCB

Hybrid MPI/OpenMP

- \Rightarrow Avoid intra-node communications

GPU-based collision detection

- using a third-party library (Bullet)

XDEM in the Cloud

XDEMaaS: XDEM as a Service

XDEMaaS: XDEM as a Service

What? → Full XDEM Workflow

- Pre-processing with the XDEM GUI
- XDEM parallel execution
- Post-processing with ParaView

For Who?

- Team members
- Interns, Partners
- ...

Why? → Hide technical difficulties

- Easy installation and usage
- Controlled and well-tested environment
- No data transfer, data lives 'in the Cloud'
- Can accommodate the different requirements:
 - Large memory, many computing nodes
 - Graphical session with GPU rendering

XDEMaaS: XDEM as a Service

How? → Using UL HPC infrastructure

- XCS portal web interface
- Remote visualization based on VirtualGL/TurboVNC
- Pre-configured jobs for XDEM tasks
 - Visualization jobs run on GPU nodes
 - Parallel jobs on Gaia cluster




XDEMaaS: XDEM as a Service

Welcome to

Bull's Extreme Computing Studio

[Hide](#) ✕

Bull is offering to speed your forecasting and simulations with its Extreme Computing solutions. Leveraging on the expertise acquired while developing some of the world's most powerful supercomputers. Bull has designed its bull-x range specifically for Extreme Computing. Modular architecture, standard components, centralized administration: Bull offers solutions to meet every type of need, from departmental server to world-class supercomputer.



Applications shortcuts



Abaqus



MATLAB



Mathematica



Paraview



R



Upload input data

You can add your private data which you can store on HPC Clusters.



Submit job

You can submit a new job.



Monitor jobs

You can monitor execution of your jobs



Check results

You can check your job results



New visualization session

You can create new visualization session.



Monitor visualization sessions

You can monitor your visualization sessions

XCS web portal

XDEMaaS: XDEM as a Service

Case: fdirs/besseron/Upload/XDEM_examples/RoundHopper_4279_Hooke/RoundHopper_4279_Hooke-36x36x36.h5

File name: Particle

	X (m)	Y (m)	Z (m)	size_a (m)	size_b (m)	size_c (m)	size_d (m)	s
1	-0.0441024	-0.0217636	-0.470414	0.0008270000108	0	0	0	0
2	-0.00601662	-0.0276361	-0.483994	0.0008270000108	0	0	0	0
3	0.0557749	-0.0183344	-0.467443	0.0008270000108	0	0	0	0
4	-0.00239369	0.00762572	-0.498204	0.0008270000108	0	0	0	0
5	-0.0098871	0.0300033	-0.488639	0.0008270000108	0	0	0	0
6	-0.00333732	0.0134881	-0.483597	0.0008270000108	0	0	0	0


#particles: -2147483648


Messages


- WARNING in 'hopper': Dynamics/initial conditions (translation, rotation) will be used for boundary motion! NOTE: 'Boundary_Taylor' is currently selected as Dynamics/Integration scheme but no Dynamics/Time dependent trajectory option is enabled.
- WARNING in 'wall': Dynamics/initial conditions (translation, rotation) will be used for boundary motion! NOTE: 'Boundary_Taylor' is currently selected as Dynamics/Integration scheme but no Dynamics/Time dependent trajectory option is enabled.

Pre-processing using the XDEM GUI via TurboVNC

XDEMaas: XDEM as a Service

XDEM-dev Parallel Driver

XDEM-dev Reconstructpar

XDEM-dev XDMF Converter

Available clusters

Gaia

The name of your job (optional)

XDEM


Fill in the job submission form

Version

20160928-NewDesignDev

Choose input file:

/users/xbessonon/Upload/XDEM_examples/RoundHopper_4279_Hooke-36x36x36.h5

 Show or hide file tree

Write sub-shape output?

☐

Number of cores per node

12

Number of nodes

5

**Walltime in hhh:mm:ss format,
e.g. 2:00:00 for 2 hours job**

120:00:00

Save parameters as default

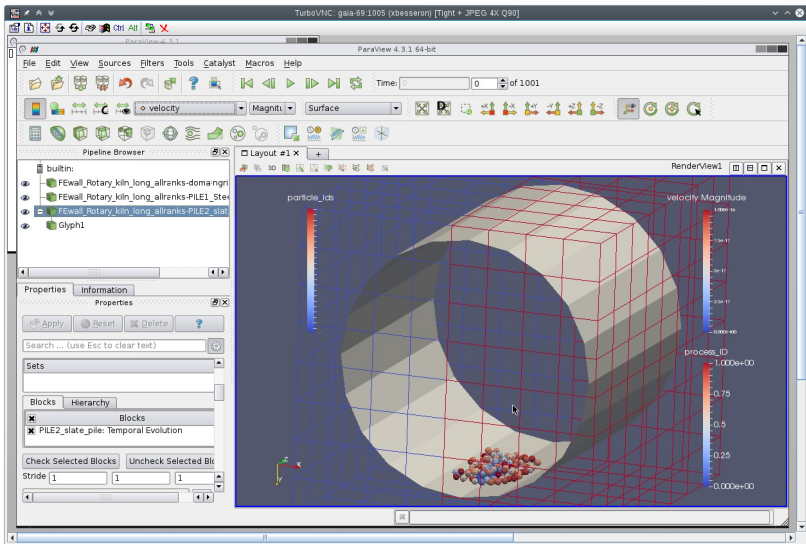
☐ Save input parameters as default for this project

Cancel

Submit

Submission of an XDEM parallel job using the XCS webportal

XDEMaaS: XDEM as a Service



Post-processing with ParaView via TurboVNC

XDEMaaS: Future Work

Support for transparent checkpoint/restart

- Hide all the technical issues

Interactive control and monitoring of XDEM simulations

- Display some metrics in real time
 - iteration time, maximum velocity, total kinetic energy, ...
- Change some parameters, eg simulation end time
- Trigger actions: force checkpoint, write current output

Real time visualization

- Using ParaView Catalyst?

Cloud vs In-House HPC: A Cost Study using XDEM

Cost Analysis of an HPC platform [1]

Amortized CAPEX + OPEX \rightarrow TCO (Total Cost of Ownership)

CAPital EXpenses

- machines (*4 years*)
- servers (*3 years*)
- storage (*5 years*)
- interconnect (*8 years*)
- room equipment (*15 years*)
- building (*30 years*)

OPerating EXpenses

- manpower
- energy (power and cooling)
- support
- software licenses

\rightarrow Cost are calculated assuming a full utilization of the platform

[1] J. Emeras, S. Varrette, and P. Bouvry. "Amazon Elastic Compute Cloud (EC2) vs. in-House HPC Platform: a Cost Analysis". In: [CLOUD 2016](#)

Cost Analysis of UL HPC platform [1]

Chaos cluster

Node	Cores	Memory	GPUs	Nodes	CPU family	Disk	GFLOPS	Hourly Cost
d-cluster1	12	24 GB	0	16	Westmere	250 GB	108.48	\$ 0.439
e-cluster1	16	32 GB	0	16	Sandybridge	250 GB	281.60	\$ 0.433
h-cluster1	12	24 GB	0	32	Westmere	250 GB	108.48	\$ 0.428
r-cluster1	32	1024 GB	0	1	Nehalem	250 GB	289.28	\$ 1.814
s-cluster1	16	32 GB	0	16	Sandybridge	250 GB	281.60	\$ 0.433

Gaia cluster

Node	Cores	Memory	GPUs	Nodes	CPU family	Disk	GFLOPS	Hourly Cost
gaia-[1-60]	12	48 GB	0	60	Westmere	256 GB	108.48	\$ 0.453
gaia-[123-154]	12	48 GB	0	32	Westmere	256 GB	147.36	\$ 0.344
gaia-[61-62]	12	24 GB	1792	2	Westmere	256 GB	108.48	\$ 0.641
gaia-[63-72]	12	24 GB	10240	10	Westmere	256 GB	108.48	\$ 0.599
gaia-[75-79]	16	64 GB	12480	5	Sandybridge	256 GB	281.60	\$ 0.577
gaia-[83-122]	12	48 GB	0	40	Westmere	256 GB	140.64	\$ 0.344
gaia-73	160	1024 GB	0	1	Sandybridge	256 GB	2,560.00	\$ 2.649
gaia-74	32	1024 GB	0	1	Sandybridge	256 GB	614.40	\$ 1.516
gaia-[155-178]	24	128 GB	0	24	Haswell	160 GB	983.04	\$ 0.458

[1] J. Emeras, S. Varrette, and P. Bouvry. "Amazon Elastic Compute Cloud (EC2) vs. in-House HPC Platform: a Cost Analysis". In: [CLOUD 2016](#)

Experimental Setup: Execution Platforms


Amazon EC2: c4.8xlarge instance

- Intel Haswell Xeon E5-2666v3 @ 2.9 Ghz
 - 36 virtual cores
 - 60 GBytes memory
 - 10 Gigabit Ethernet (10 Gb/s)
- Hourly cost: \$ 1.906 (on-demand) or \$ 1.2182 (reserved full upfront)

UL HPC cluster: Dell FC430 nodes

- 2 × Intel Haswell Xeon E5-2680v3 @ 2.5 Ghz
 - 24 cores
 - 120 GBytes memory
 - InfiniBand QDR (40 Gb/s)
- Hourly cost: \$ 0.458

Experimental Setup: XDEM testcase & Benchmark



Discharge of granular materials

- 427,878 spherical particles
- 6,742 triangles
- Impact & Rolling friction
- 15 s simulation
- 750,000 timesteps with $\Delta t = 2e^{-5}$ s
- 30 GBytes output
- Decomposition in 1,470,000 cells
- ORB partitioning
- Sequential time: 3.5 months

⇒ Extracted XDEM benchmark

- 100 representative timesteps

Performance evaluation

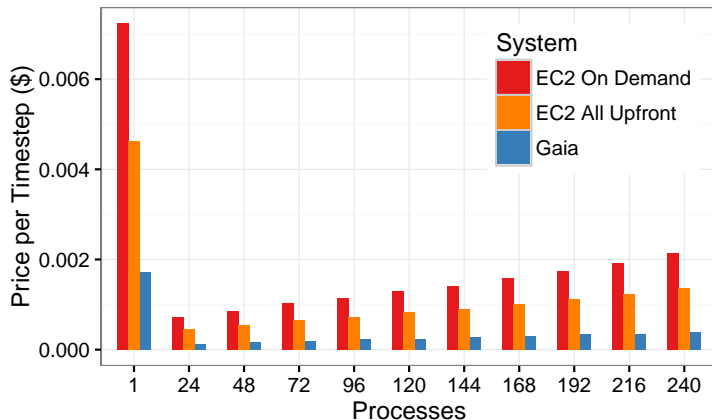
Average execution time of one XDEM timestep

	Timestep execution time	
	Gaia Dell FC430	EC2 c4.8xlarge
Sequential (1 core)	13.55 s	13.66 s
Parallel 24 cores on 1 node	0.99 s	1.36 s
Parallel with 240 cores	0.31 s	0.40 s

- ⇒ Similar sequential performance
- ⇒ 30% overhead with intra and inter-node communication

Taking into account performance and cost

Calculating Price per Timestep



- ⇒ Amazon EC2 is 3 times to 5 times more expensive
- ⇒ Sequential execution still requires to pay for one full node

Cost of running the full XDEM testcase

Extrapolated duration and price

- Full testcase with 750,000 timesteps
- Trade-off between deadline and cost

	Cheapest (24 processes)		Fastest (240 processes)	
Gaia Dell FC430	8.61 days	\$ 94.64	2.69 days	\$ 295.29
c4.8xlarge on-demand	11.81 days	\$ 540.55	3.50 days	\$ 1,602.84
c4.8xlarge reserved	11.81 days	\$ 345.48	3.50 days	\$ 1,024.44

Conclusions

XDEM needs HPC

- Important requirements in computing power and memory

XDEM can provide Cloud services: XDEMaaS

- Accommodate the various requirements of XDEM workflow
- Hide many technical issues

XDEM can use Cloud platforms (IaaS, PaaS)

- Can be expensive!
- XDEM is an HPC application: performance is not yet satisfactory
- Solutions are on the way: containers, SR-IOV-based MPI, ...

Thank you!

Any question?

LuXDEM Research Centre
<http://luxdem.uni.lu/>