

Citius, Altius, Fortius



4 milliard de gobelets plastique consommés en France chaque année
0,1g de plastique en moins = 400 T économisées /an

Comment faire, fonction de :

- Température de la boisson
- Résistance du plastique
- Transmission de la chaleur
- Pression des doigts de l'utilisateur

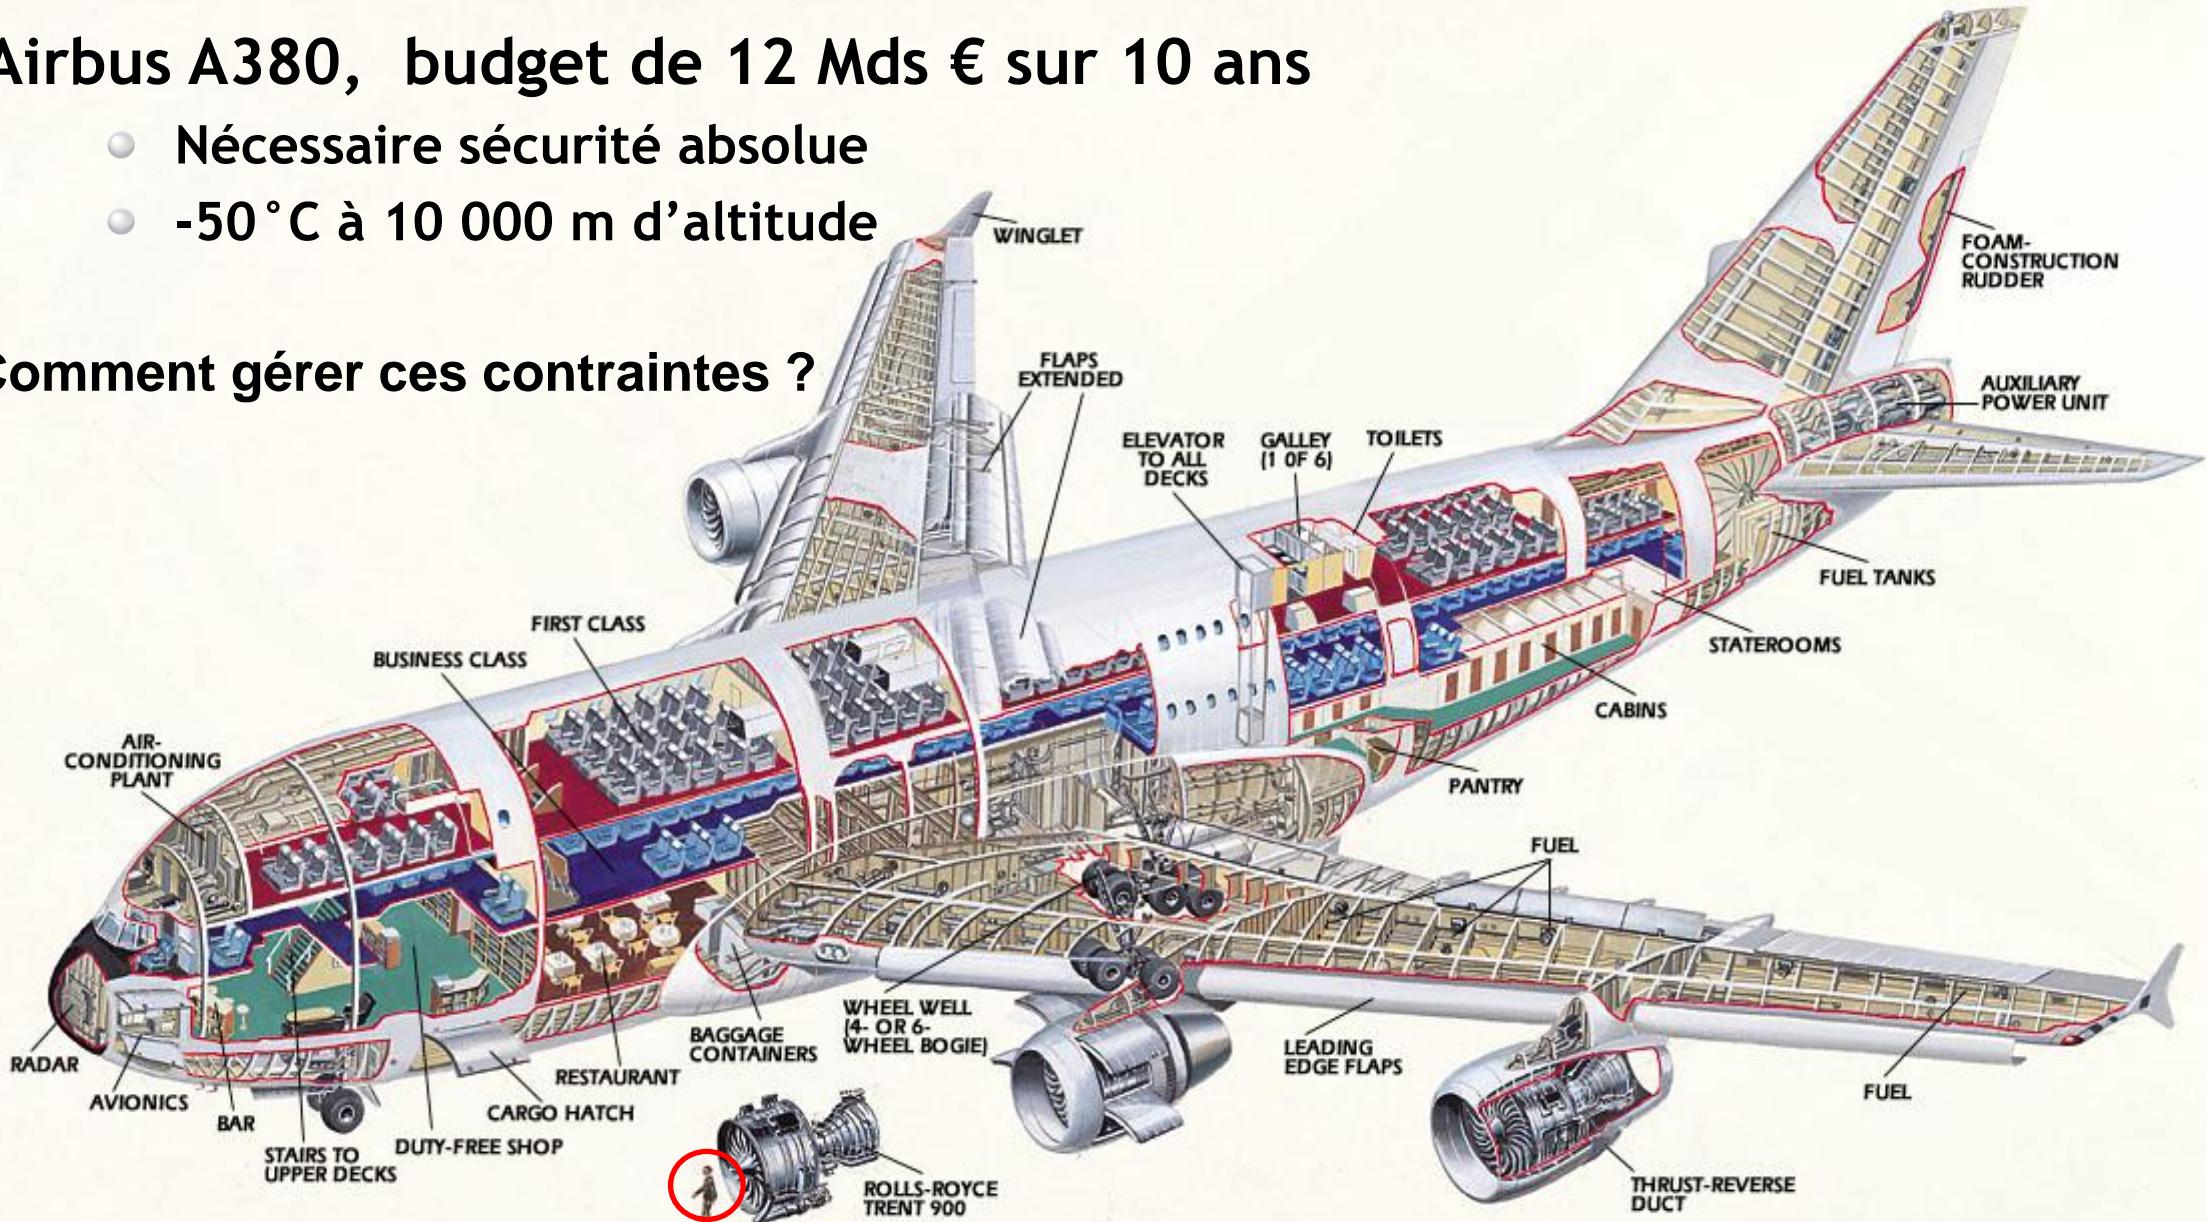


= plusieurs centaines de milliers de combinaisons à gérer

Airbus A380, budget de 12 Mds € sur 10 ans

- Nécessaire sécurité absolue
- -50 °C à 10 000 m d'altitude

Comment gérer ces contraintes ?



La clé est dans la simulation



Objectifs

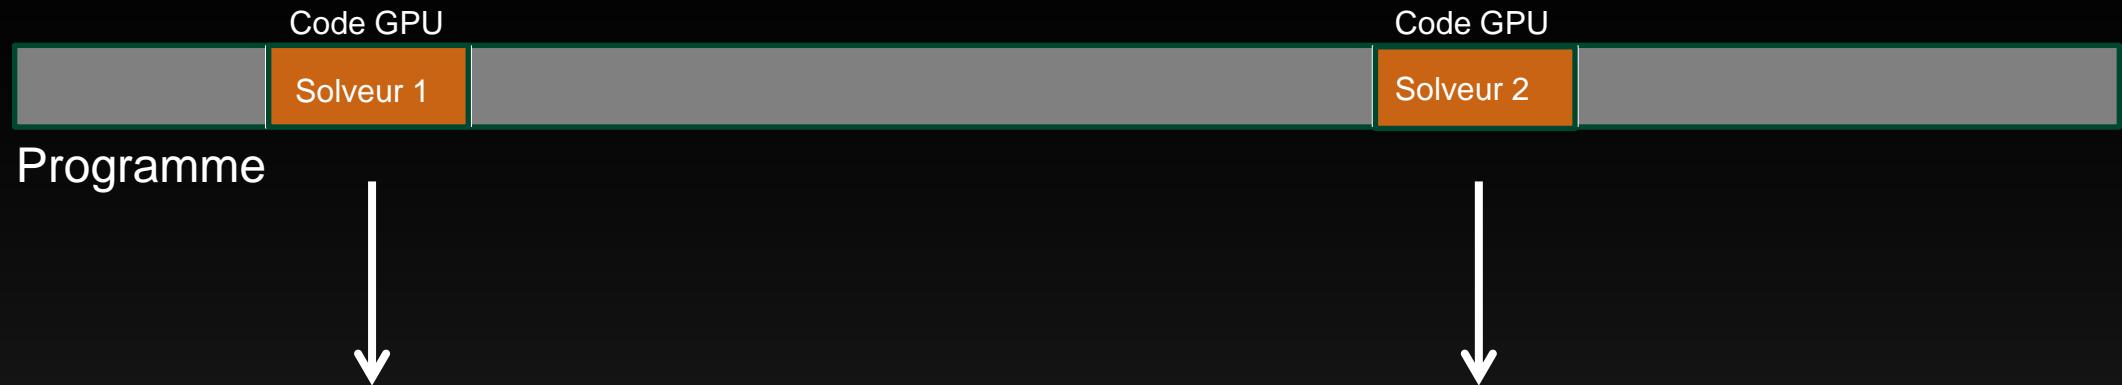
- Réduction temps total d'exécution
- Multiplication des simulations

Si cluster CPU

- Onéreux
- Besoin de place
- Consommation électrique élevée



La nouvelle clé est d'utiliser le GPU



Durée identique mais augmentation des simulations

ou



Application plus rapide

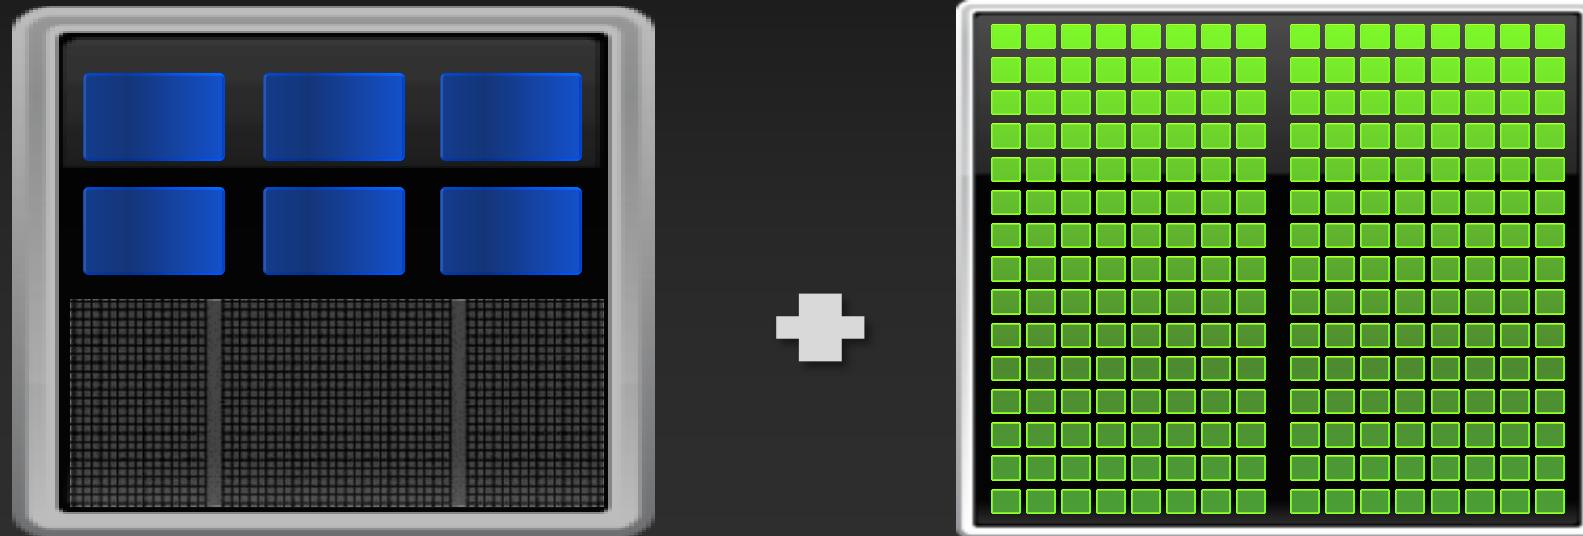
GPU Computing

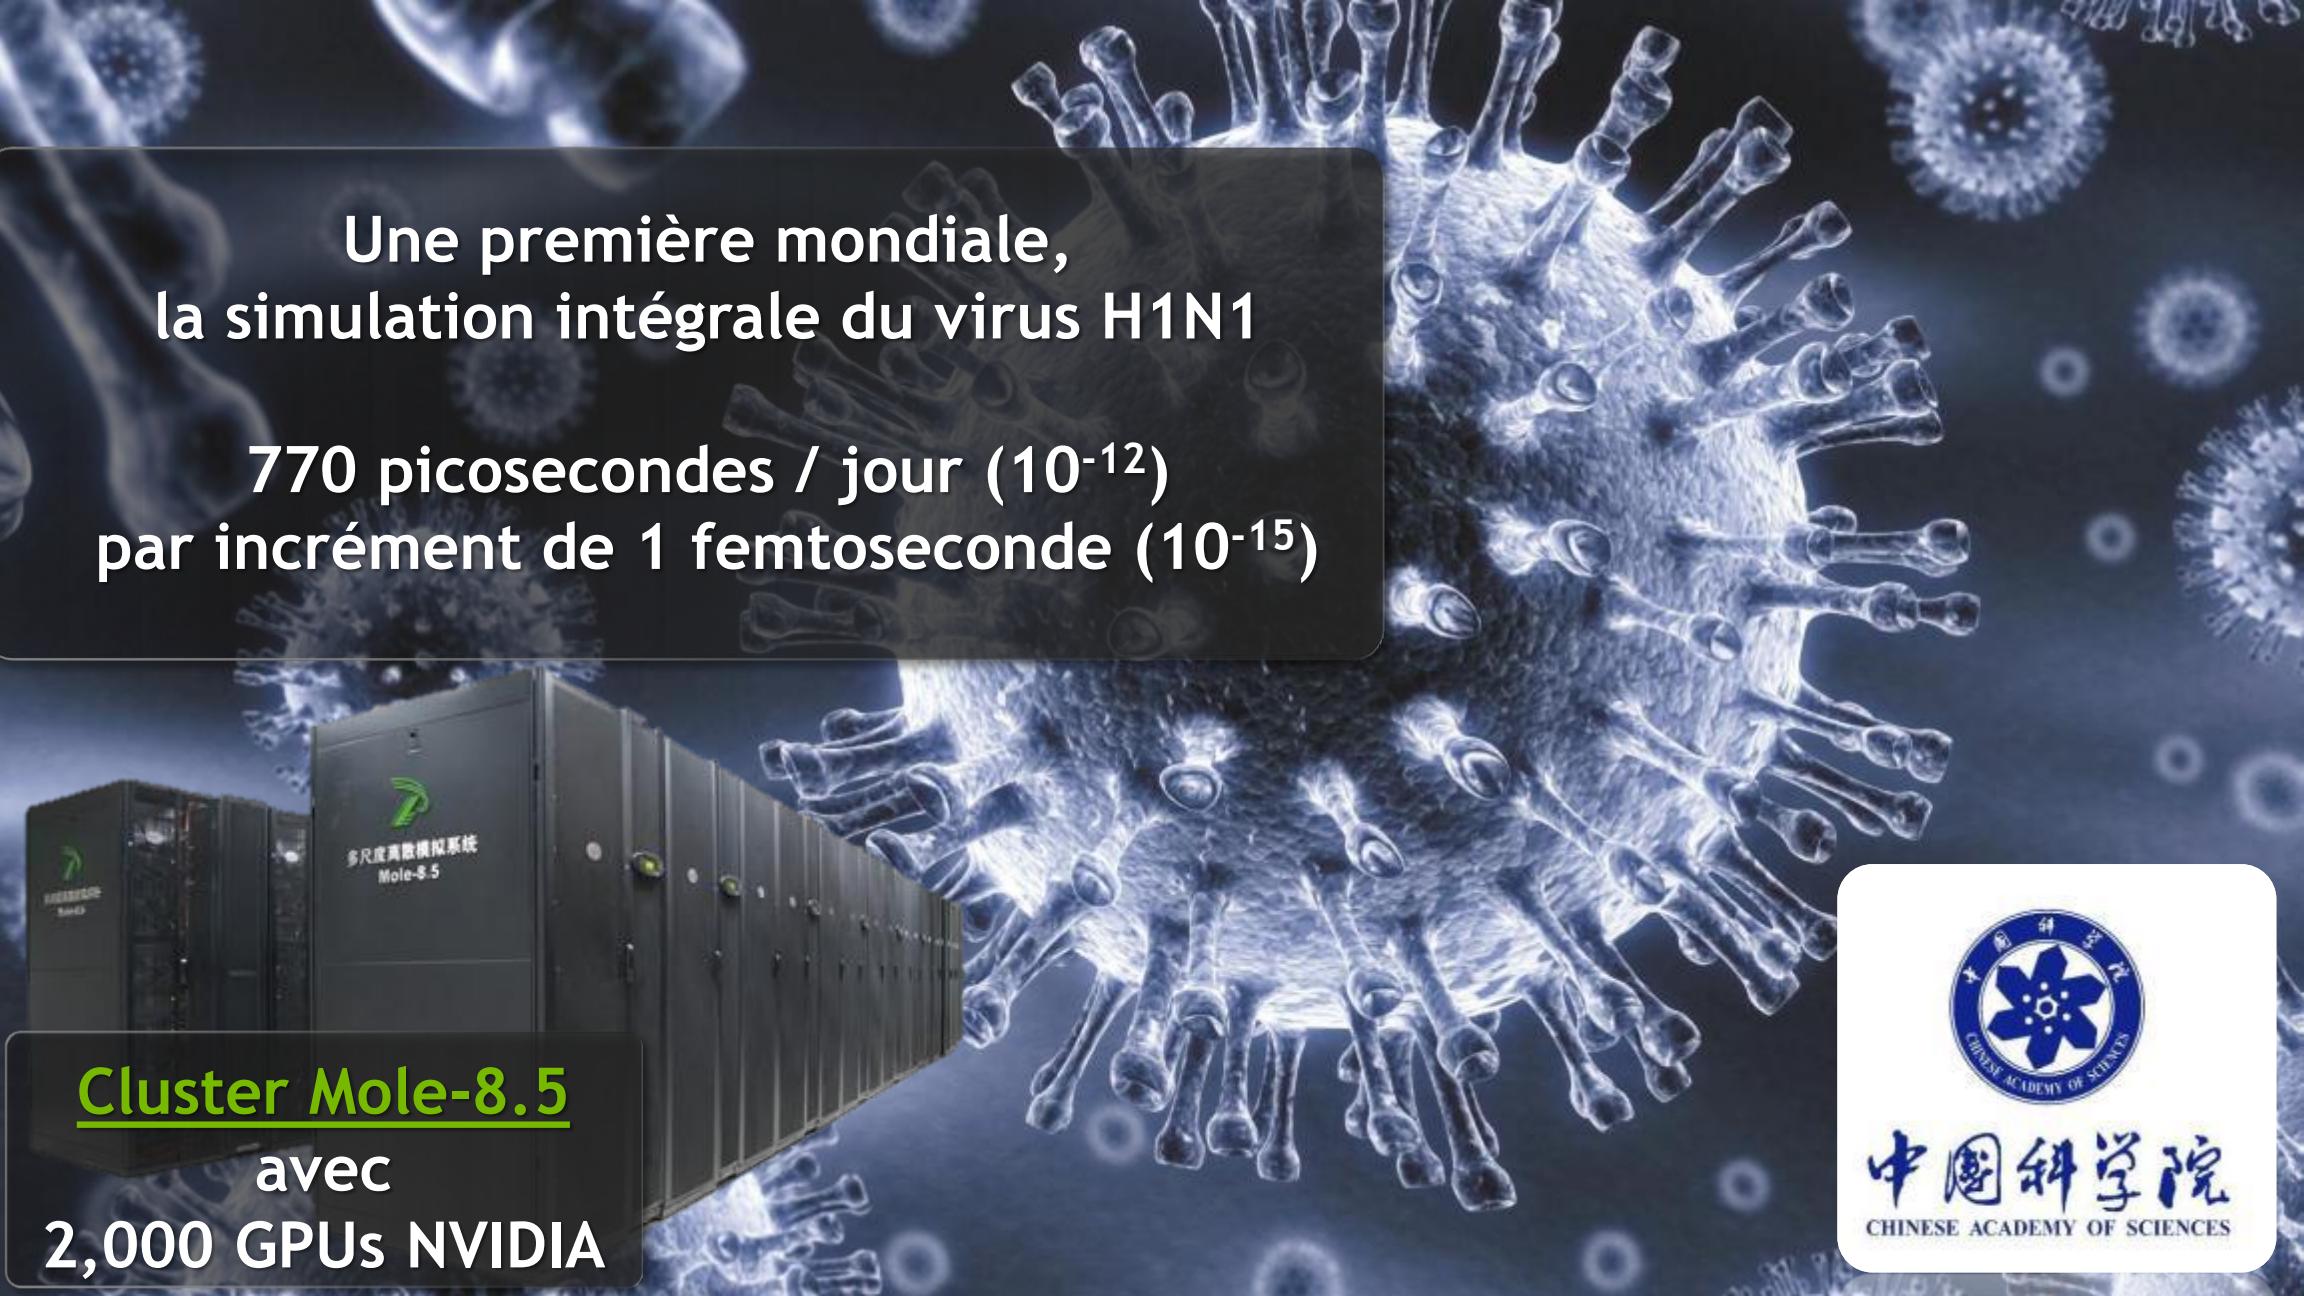
CPU = gare de triage, anticipation maximale

GPU = processeur dédié massivement parallèle

Seule une partie du code est portée sur le GPU

CPU + GPU = optimal pour calcul puis affichage





Une première mondiale,
la simulation intégrale du virus H1N1

770 picoseconds / jour (10^{-12})
par incrément de 1 femtoseconde (10^{-15})



Cluster Mole-8.5
avec
2,000 GPUs NVIDIA



中國科学院
CHINESE ACADEMY OF SCIENCES

1

Applications portées sur GPU

- Molecular Dynamics, Molecular Docking, Quantum Chemistry...

2

Aide au développement

- Formation
- Librairies
- Compilateur à base de directives
- Programmation

3

Un seul code, du portable au cluster

- Portable NVIDIA GeForce & Quadro
- Station de Travail NVIDIA Quadro & Tesla
- Cluster NVIDIA Tesla

4

Valorisation de la recherche

Applications

Jun 2012: A major update patch will shortly be released for AMBER 12.

This will provide support for Kepler I (K10) and GTX6XX card.

It will also include additional performance boosts (expect **75+ns/day** for JAC NVE on GTX680)

Liste des applications

NVIDIA GPU Test Drive

Test en ligne gratuit

AMBER 12 NVIDIA GPU ACCELERATION SUPPORT

[Background](#)	[Authorship & Support](#)	[Features](#)	[Supported GPUs](#)	[System Size Limits](#)
[Accuracy Considerations](#)	[Installation and Testing](#)	[Running GPU Accelerated Simulations](#)		
[Considerations for Maximizing GPU Performance](#)	[Benchmarks](#)			
[Recommended Hardware](#)	[MD SimCluster](#)			
[Return to Main Amber Page](#)				

News and Updates

- Jun 2012: A major update patch will shortly be released for AMBER 12. This will provide support for Kepler I (K10) and GTX6XX cards. It will also include additional performance boosts (expect 75+ns/day for JAC NVE on GTX680)
-  [NVIDIA and Partners Announce MD SimCluster Test Drive Program for AMBER](#)
- [Looking to build your own GPU workstations with 2 GPUs for < \\$2500. The following is a suggested specification by Ross Walker for building your own AMBER GPU machine](#)

Background

This page provides background on running AMBER (PMEMD) with [NVIDIA GPU](#) acceleration.

One of the new features of AMBER 11 was the ability to use NVIDIA GPUs to accelerate PMEMD for both explicit solvent PME and implicit solvent GB simulations. This has been further extended in AMBER 12 and will shortly be improved further with an update patch for AMBER 12. While this GPU acceleration is considered to be production ready, and indeed is widely used, it has not been tested anywhere near as extensively as the CPU code has over the last 30 years. Therefore users should still exercise caution when using this code. The error checking is not as verbose in the GPU code as it is on the CPU. If you encounter problems during a simulation on the GPU you should first try to run the identical simulation on the CPU to ensure that it is not your simulation setup which is causing problems. Feedback and questions should be posted to the [Amber mailing list](#).

New in AMBER 12

AMBER 12 includes full GPU support in PMEMD and is a major update over AMBER 11. Focus for AMBER 12 has been on increasing the features available in the GPU accelerated code. Performance is currently similar to AMBER 11 + latest bugfixes, however, a Kepler and performance patch will be released shortly that will significantly boost AMBER 12



Tesla Bio Workbench



A philanthropic initiative of the NVIDIA Foundation



COMPUTE THE CURE

Find a cure for cancer.
Faster.

www.opengenomicsengine.org

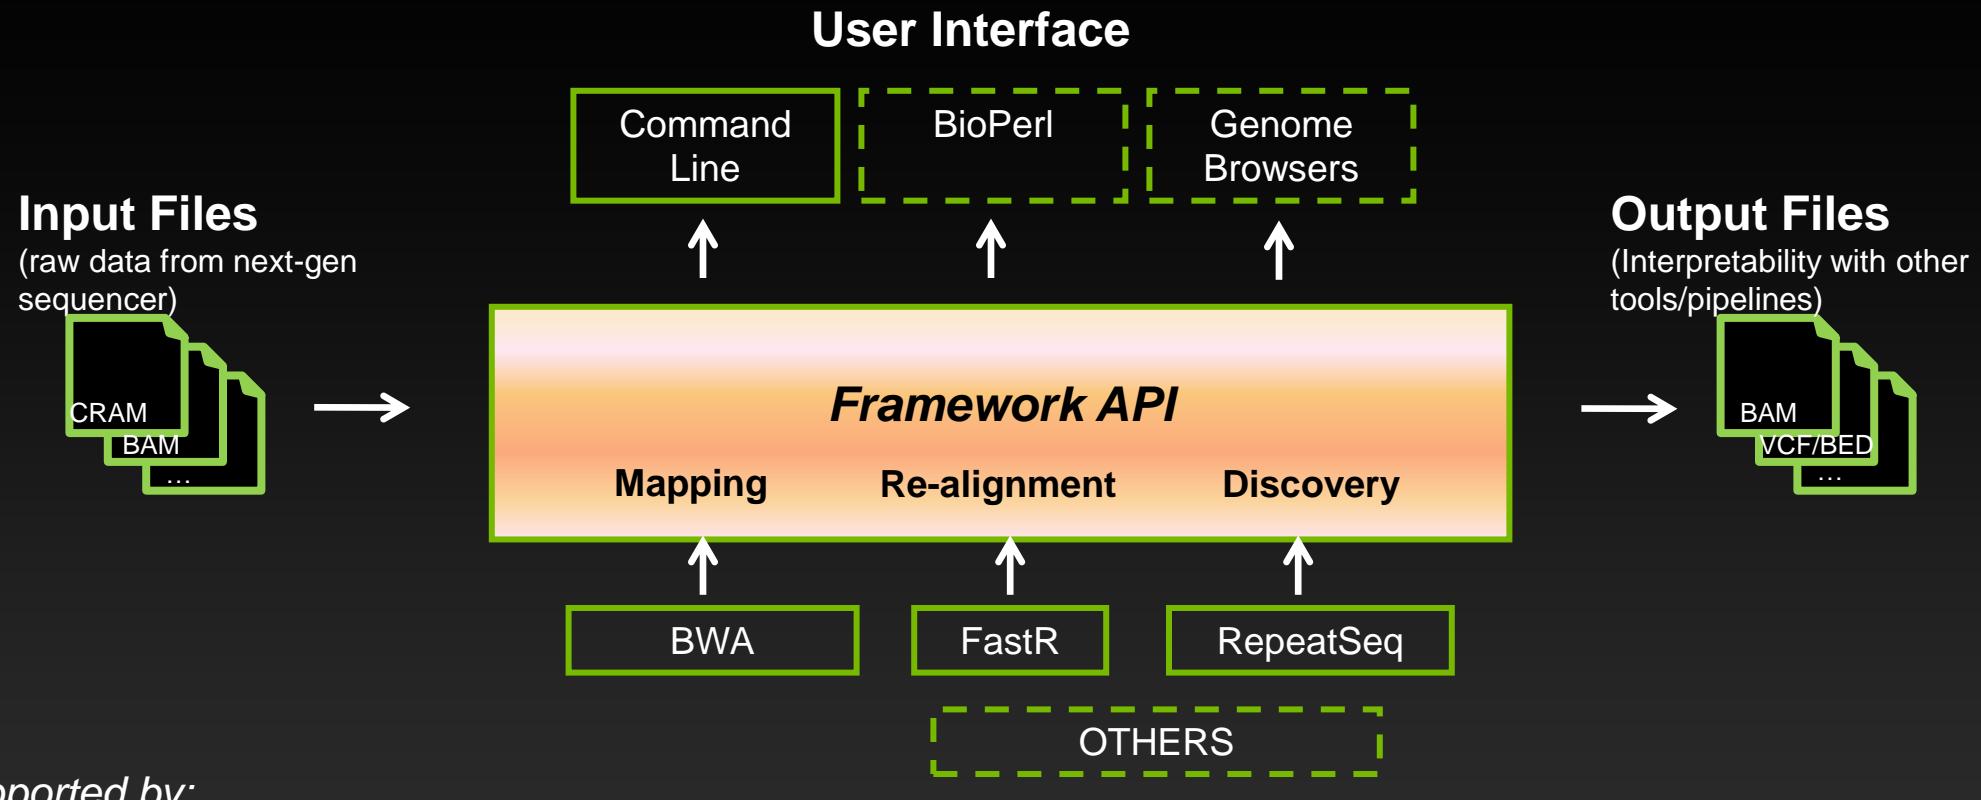
OPEN GENOMICS ENGINE™

- An open source platform that streamlines the comparative analysis of cancer genomes to speed discovery
 - *Cancer scientists can shorten the time required to analyze human genomes in their workflows*
 - *Computational scientists can more easily create algorithms and make them readily available to others*
- Funded through Compute the Cure, a philanthropic initiative of the NVIDIA Foundation
- Learn more at www.opengenomicsengine.org

Supported by:



OPEN GENOMICS ENGINE™



Supported by:

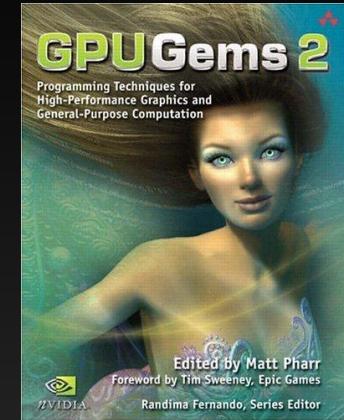
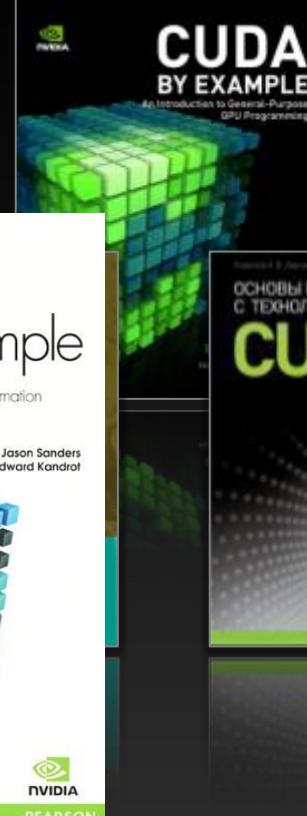
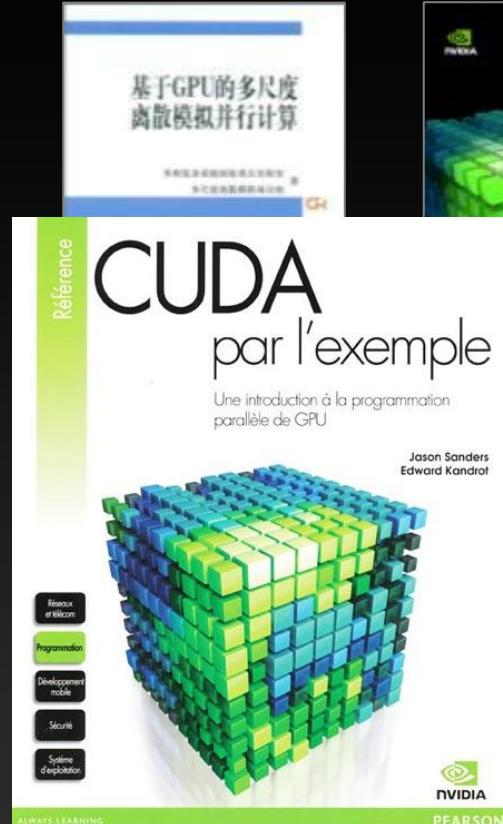


Sequencing Algorithms
(plug-ins)

Phase I

Phase II

2 Se former



Formations en ligne

CUDA Training

This page has online courses to help you get started programming or teaching CUDA as well as links to Universities teaching CUDA.

This page organized into three sections to get you started

- [Introductory CUDA Technical Training Courses](#)
- [CUDA University Courses](#)
- [CUDA Seminars and Tutorials](#)

Introductory CUDA Technical Training Courses

- Volume I: [Introduction to CUDA Programming](#)
 - Exercises (for Linux and Mac)
 - Visual Studio Exercises (for Windows)
 - Instructions for Exercises
- Volume II: [CUDA Case Studies](#)

CUDAcasts - Downloadable CUDA Training Podcasts

- [Introduction to GPU Computing](#)
- [CUDA Programming Model Overview](#)
- [CUDA Programming Basics - Part I](#)
- [CUDA Programming Basics - Part II](#)

[Follow This Link For Additional GPU Computing Online Seminars](#)

CUDA University Courses

University Of Illinois : ECE 498AL

Taught by Professor [Wen-mei W. Hwu](#) and [David Kirk](#), NVIDIA CUDA Scientist.

- [Introduction to GPU Computing](#) (60.2 MB)
- [CUDA Programming Model](#) (75.3 MB)
- [CUDA API](#) (32.4 MB)
- [Simple Matrix Multiplication in CUDA](#) (46.0 MB)
- [CUDA Memory Model](#) (109 MB)
- [Shared Memory Matrix Multiplication](#) (81.4 MB)
- [Additional CUDA API Features](#) (22.4 MB)
- [Useful Information on CUDA Tools](#) (15.7 MB)
- [Threading Hardware](#) (140 MB)
- [Memory Hardware](#) (85.8 MB)
- [Memory Bank Conflicts](#) (115 MB)
- [Parallel Thread Execution](#) (32.6 MB)
- [Control Flow](#) (96.6 MB)
- [Precision](#) (137 MB)

Understanding OpenACC Directives By CAPS

In this webinar you will refine your knowledge of OpenACC directive programming and how you can get better acceleration for real scientific applications.

Presented by CAPS, one of the leading experts in parallel computing.

[Video\(mp4\)](#)

OpenACC For Cray Supercomputers

In this webinar you will get an overview of OpenACC support in Cray's compilers and an insight into one of the most easy to use solutions for high performance parallel computing. Presented by Cray, one of the leading experts in supercomputing

[Video\(mp4\)](#)

OpenACC Acceleration For Real Science - Using CAPS HMPP

In this webinar you will refine your knowledge of OpenACC directive programming and how you can get better acceleration for real scientific applications. Presented by CAPS, one of the leading experts in parallel computing.

[Video\(mp4\)](#)

NVIDIA CUDA Zone

3 points d'entrées

Librairies

- BLAS, FFT, Sparse Matrix...

Directives

- Consortium OpenACC

Language de programmation

- C et C++
- FORTRAN

1

GPU-Accelerated Libraries

Drop in a GPU-accelerated library to replace MKL, IPP, FFTW and other widely-used libraries



C++ Template
Library



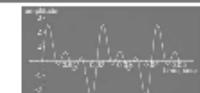
cuBLAS



cuSPARSE



NPP



cuFFT

[More GPU-Accelerated Libaries](#)

2

GPU Directives

Automatically parallelize loops in your Fortran or C code using GPU Directives



- Easy : simply insert hints in your code
- Open : run on either CPU or GPU
- Powerful: tap into the power of GPUs within minutes

[Learn More...](#)

[Learn More About Directives](#)

3

Programming Languages

Develop your own parallel applications and libraries using a programming language you already know



CUDA C/C++
GPU Acceleration for
C and C++ Apps.
[Learn more...](#)



CUDA Fortran
GPU Acceleration for
Fortran Applications
[Learn more...](#)

[More Programming Language Solutions](#)

Un seul code du PC portable au cluster HPC

Prise en main /
Dév. mono GPU
Portable



GeForce / Quadro

Dév. Multi-GPU
Station de Travail



Quadro / Tesla

Déploiement
Centre de Calcul



Tesla

Valorisation de la recherche

Un service de NVIDIA Research

Academic Partners

CUDA Centers of Excellence

CUDA Fellows

CUDA Research Centers

CUDA Teaching Centers

CENTERS & PARTNERS

Academic Partners
CUDA Centers of Excellence
CUDA Fellows
CUDA Research Centers
CUDA Teaching Centers

PROGRAMS

Academic Partnership
CUDA Center of Excellence
CUDA Research Center
CUDA Teaching Center
Graduate Fellowship

ADDITIONAL INFORMATION

CUDA Courses Around the World
CUDA Course Materials
CUDA Education and Training
CUDA Forums
CUDA Zone
Developer Zone
Intern & Coop Programs
Sign Up For Research News
Tegra Prototype Instructions

EXTERNAL LINKS

GPUComputing.net
GPGPU.org

Support for Academics & Researchers

NVIDIA believes in extending our University investment to include support of teaching, research, and advanced education. NVIDIA currently has five programs at varying levels of participation and collaboration.

Academic Partnership Program

The focus of the Academic Partnership Program (formerly Professor Partnership) is to form partnerships with academic researchers worldwide doing research and/or teaching using the power of the GPU. The program is an application process and support can be requested for small scale NVIDIA equipment donations (1-2 boards). We typically review hardware requests on a bi-weekly basis, but with GTC coming up we will be a couple weeks behind. We'll notify you of the results of your submissions as soon as we can.

CUDA Center of Excellence Program

The CUDA Center of Excellence (CCE) program recognizes, rewards, and fosters collaboration with universities at the forefront of massively parallel manycore computing research. Schools identified as CUDA Centers of Excellence have proposed a unique vision for improving the technology and application of parallel computing. The CUDA Center of Excellence program is competitive and prestigious. Universities meeting the requirements are invited to submit proposals to NVIDIA at any time; proposals will be evaluated on a rolling basis.

CUDA Research Center Program

Institutions identified as CUDA Research Centers are doing world-changing research in a particular domain or field by leveraging CUDA and NVIDIA GPUs. Any institution whose primary goals include research can apply for the CUDA Research Center program. Applications are reviewed on a rolling basis, typically quarterly. We are currently reviewing proposals submitted prior to April 11, 2012 and expect to announce the results early June. The next review will be early August 2012, application deadline is July 11, 2012.

CUDA Teaching Center Program

The CUDA Teaching Center Program is designed to support and encourage teaching establishments to include GPU Computing using CUDA C/C++ as part of their course offerings. To help facilitate this teaching effort, NVIDIA may grant equipment, funding and/or course material assistance. This is an application based program and applications are reviewed on a quarterly basis. We are currently reviewing proposals submitted prior to April 11, 2012 and expect to announce the results early June. The next review will be early August 2012, application deadline is July 11, 2012.

Graduate Fellowship Program

The NVIDIA Graduate Fellowship Program provides funding to Ph.D. students who are researching topics that will lead to major advances in a number of fields, and are investigating innovative ways of leveraging the power of the GPU. We select students each year who have the talent, aptitude and initiative to work closely with us early in their careers. Recipients not only receive crucial funding for their research, but are able to conduct groundbreaking work with access to NVIDIA products, technology and NVIDIA engineers. Students should have completed at least one year of their doctoral program with some results of their research. Our application period is now closed and we are reviewing applications, with the results announced mid March 2012.

Le premier CUDA Research Center en France





Pour toute information complémentaire

Jean-Christophe Baratault
jbaratault@nvidia.com
P 06 80 36 84 83

