Réflexions et actions en cours en matière de calcul et de données au CNRS

www.cnrs.fr

M. Daydé

Directeur du Comité d'Orientation pour le Calcul Intensif au CNRS (COCIN)

Délégué Scientifique CNRS / INS2I

Rôle et missions du COCIN

- Créé en Décembre 2010
- Réflexion collective sur les besoins, la structuration et les évolutions en calcul intensif au CNRS
- Prospective sur les besoins des différentes communautés, proposition de maintenance et de développement coordonné des moyens / ressources liées au calcul intensif, en particulier pour l'IDRIS.
- Dix personnalités scientifiques désignées par chacun des instituts du CNRS + Directeur de l'IDRIS + 4 ingénieurs experts
- Le président et directeur désignés par le Président du CNRS

Missions et production du COCIN

Créé en Décembre 2010 : réflexion collective sur les besoins, la structuration et les évolutions de l'écosystème du calcul intensif et des données scientifiques au CNRS

- Composition : 1 représentant / institut + DU IDRIS + 4 experts

Michel Bidoit (INS2I) : Président	B. Jouve(INSHS)	**M. Daydé (INS2I) : Directeur**
D. Girou (IDRIS)	F. Godeferd (INSIS)	Ph. Helluy (INSMI)
S. Lamarre (INEE)	L. Lellouch (INP)	P.-E. Macchi (IN2P3)
JC. Michalski (INSB)	C. Pouchan (INC)	JP Vilotte (INSU)

Experts : D. Bascle (INC), F. Berthou (INP), M. Libes (INSU), V. Miele

Invités : V. Breton (IDGC), O. Porte (DSI), JP Proux (GENCI)

- Production :
  - *Livre Blanc sur le Calcul Intensif au CNRS fin 2012*
  - *Propositions pour une nouvelle stratégie du calcul et des données au CNRS (Décembre 2013)*
  - *Livre blanc sur l'Informatique en appui à la recherche (Mars 2014)*
  - *Evolution des coûts des infrastructures pour le calcul intensif et le traitement des données à grande échelle (2014)*
  - *Pratiques en matière de données dans le périmètre du CNRS*

Mission pour le calcul haute performance, les grilles, le “cloud” et les infrastructures de données scientifiques “massives”

- **Nomination d’un chargé de mission auprès de la Présidence du CNRS : Denis Veynante**
- **Contexte :**

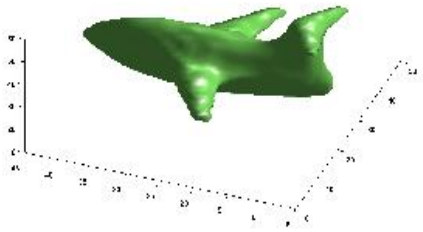

➡ **Développement modélisation / simulation**

➡ **Avancées scientifiques significatives**

‣ *Astrophysique, chimie, climat, mécanique des fluides, physique,...*

➡ **Explosion des besoins de puissance de calcul**

‣ *Machines massivement parallèles*
‣ *Couplages multi-physiques*

➡ **Explosion des volumes de données**

‣ *Simulations, Grands instruments (LHC, téléscopes, satellites, réseaux de capteurs,...), Grandes bases de données*

➡ “**Discipline” transverse**

‣ *Par nature très inter-disciplinaire, Pas de réelle “appropriation” par chaque communauté, “Maturité” très différente selon les disciplines*

Mission pour le calcul haute performance, les grilles, le “cloud” et les infrastructures de données scientifiques “massives”

- *Définition et mise en oeuvre d’une politique globale et cohérente du CNRS sur les objets à coût est significatif (y compris en ETPT)*
- *Pilotage ou participation au pilotage des ressources dédiées*
  - IDRIS, CC-IN2P3
  - France-Grilles
  - Maison de la Simulation
  - ...
- *Cohérence de la politique du CNRS avec les “objets” nationaux (GENCI, Renater, Infranum,...) ou internationaux (PRACE,...)*
- *Aider les DSR à construire une stratégie conjointe avec nos partenaires*

Mission pour le calcul haute performance, les grilles, le “cloud” et les infrastructures de données scientifiques “massives”

➡ **Comité de Pilotage : COCIN (*Comité d’Orientation pour le Calcul INtensif*)**

- *Représentation effective de tous les instituts...*
- *Rôle plus “actif” / relais dans les instituts*

➡ **Périmètre et organisation à préciser :**

- *Réflexions à conduire avec :*
  - *Instituts*
  - *Mission pour l’Interdisciplinarité*
  - *Cellule TGIR*
  - *DASTR*
  - ...

Le paysage national du calcul

Moyens nationaux (Tiers-1)
- Pour le CNRS :
  - IDRIS (UPS, Orsay)
  - CC-IN2P3 (USR, Lyon)

- Périmètre, missions et moyens bien identifiés
- « Provenance » des ressources

Moyens régionaux (Tiers-2) **Action structurante possible**
- Méso-centres :
  - Promotion du calcul intensif
  - Réponse à des besoins puissance de calcul / stockage
- Centres de compétences (maison de la simulation,…)
  - Pas directement opérateurs de moyens de calcul
  - Souvent adossés à un méso-centre

CNRS pas toujours présent dès l'origine (CPER,….) mais sollicité après coup :
- Structures pérennes (USR, UMS)
- Frais de fonctionnement
- Personnel

Calcul + données : coûts croissants

- Calcul intensif :
  - Coordination au moins au sommet de la pyramide du calcul (Tiers-0 / Tiers-1 / Tiers-2, EQUIP@MESO,…)
  - Evolution technologique + consommation énergétique croissante + adaptation des codes + compétences + support
- Données :
  - Explosion des besoins et des demandes non coordonnées (CPER) même si certaines communautés sont structurées (e.g. physique des hautes énergies, sciences de l'univers, bio, …)
  - Conforter compétences + support
- Impact sur l'organisation de la recherche

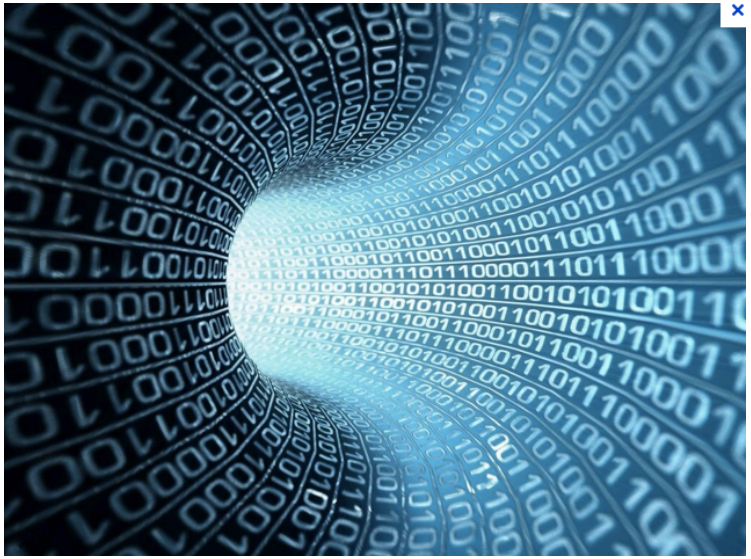

Eléments clés de la stratégie du CNRS

- Ne plus dissocier HPC de l'analyse et valorisation des masses de données issues des simulations numériques (climat, fluides turbulents,…), grands instruments (, LHC, ITER, LSST, LOFAR, plateformes génomiques …) et grands systèmes d'observation au sol (i.e., sismologie et géodésie : RESIF) et dans l'espace (Euclid, WFIRST, GAIA, imagerie et interférométrie)…
- *Calcul + data intensif pas uniquement problème de ressources mais un **changement de paradigme** dans la recherche scientifique :*
  - Plus d'inter/pluridisciplinarité (informatique, maths et autres disciplines),
  - Vision holistique des Infrastructures calcul / données / grands instruments / plateformes expérimentales / systèmes d'observation
- Rationaliser le déploiement des infrastructures, coordonner les demandes
- *En s'appuyant stratégie nationale / de site autour de défis scientifiques et maîtrise des coûts !!!*

ENQUÊTE EN COURS AUTOUR DES PRATIQUES SUR LES DONNÉES

Enquête en cours : pratiques en matière de données au sein du CNRS

- *Objectifs* :
  - Mieux cerner les pratiques et les besoins
  - Promouvoir une synergie transdisciplinaire autour des données
- Enquête démarrée en 2014avec à ce jour retour des instituts :
  - IN2P3
  - INC
  - INP
  - INEE
- Et des centres / instituts
  - CC IN2P3
  - CINES
  - IDRIS
  - Institut Français de bioinformatique (IFB)

INSTITUT DU DÉVELOPPEMENT ET DES RESSOURCES EN INFORMATIQUE SCIENTIFIQUE

Bases de données sur la Biodiversité, Ecologie, Environnements et Sociétés

Infrastructures de données

- Communautés scientifiques avec des besoins / compétences bien établies :
  - Sciences de l'univers (OSU, observatoires virtuels, ….)
  - Physique des hautes énergies (grille WLCG)
  - Biologie (RENABI, France Génomique, IFB)
  - …
- Reste des besoins immenses plus ou moins émergeants et un besoin de structuration nationale et au niveau des sites

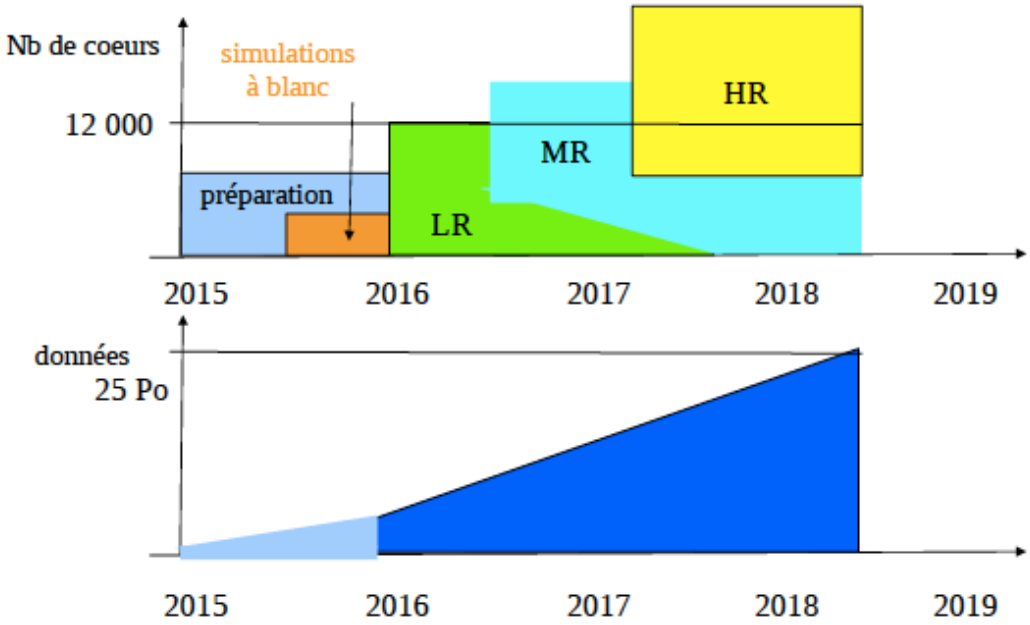

Demandes IPSL pour CMIP6 2015-2018 (JL Dufresne) : 100 millions d'heures pendant 3 ans (12,000 cœurs par an), 25 Po stockage

E-Biothon: an experimental platform for BioInformatics

M. Daydé, B. Depardon,A. Franc, JF Gibrat, R Guillier, Y Karami, C Pérez, F Suter, B Taddese, M Chabbert, S Thérond

Context

- Recent technological advances, such as high-throughput sequencers, enable biology researchers to have access to massive amounts of raw data (petabytes of data are generated every year) on the composition of viruses, bacteria, plants and animals (including human species).
- Analyzing this data to take advantage of the information available is a crucial task requiring very large amount of computer processing capacity.
- Parallel computing is then becoming more and more central in advancing the research in life sciences
- CNRS, IBM, Inria, the Institut Francais de Bioinformatique and SysFera have joined forces to give researchers access to the E-Biothon Cloud platform, hosted at IDRIS (CNRS national supercomuting center). It provides parallel computing resources and help the life science community to prepare their codes for the next generation of massively parallel computers
- Some of the challenges arising now in life sciences are related to data management, parallelism and new algorithms,….

Bioinformatics: the data deluge

- Increased amount and size of data
- Modelization in biology
  - Lot of teams working on mathematical models and software that model phenomena and analyze data

- From IEEE Spectrum, nearly 2000 genome sequencers in operation in 2013, producing 15 PB
- With desktop-class whole genome sequencing solutions, dramatic growth of sequencers expected.
- Annual sequencing-related data growth expected to grow beyond 3.3 ExaBytes by 2018.
- Data-driven medecine requires data sets far greater than DNA sequence output.

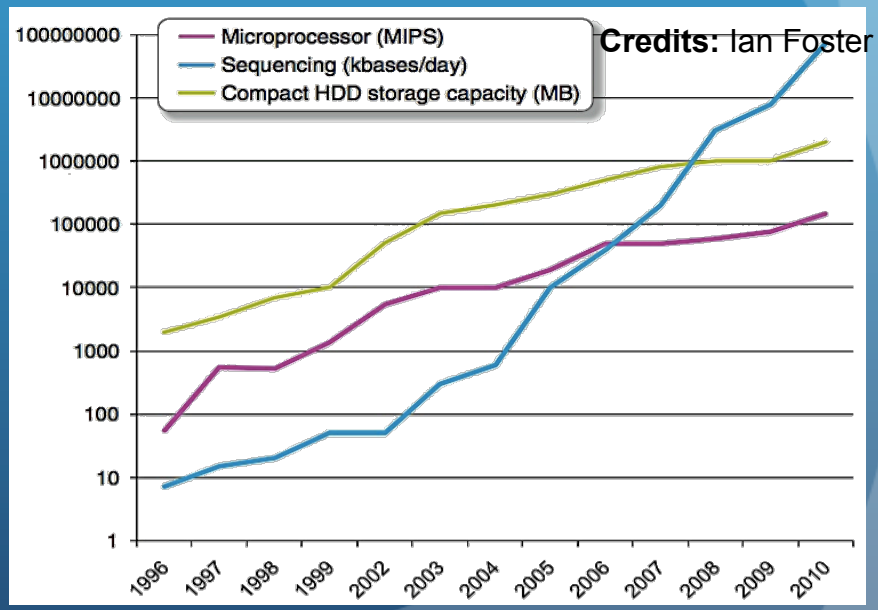

Credits: Ian Foster

Back to the DECRYPTHON 2004-2012

Goals:

- Speeding up research on genetic and rare diseases
- Make transparent the use of distributed computing resources to users

Using:

- A computing grid involving 6 sites (Bordeaux, Lille, ENS Lyon, Paris 6, Orsay and UPMC)
- With DIET middleware and a web portal

E-Biothon Platform and tools

The BlueGene/P computer

- 4 racks of BLueGene/P
  - Peak performance 56 Tflops
  - Each rack has 1024 nodes with 4 cores
  - 200 TB of storage
- Two operation modes : standard one (HPC) and High Throughtput Computing (HTC)

SysFera-DS Intuitive Web Interface

- Brings HPC & Cloud environment to local desktop
- Runs non-interactive and interactive graphical HPC applications
- Manages & visualizes remote Big Data
- Manages projects & access rights
- Follows application & resource usage

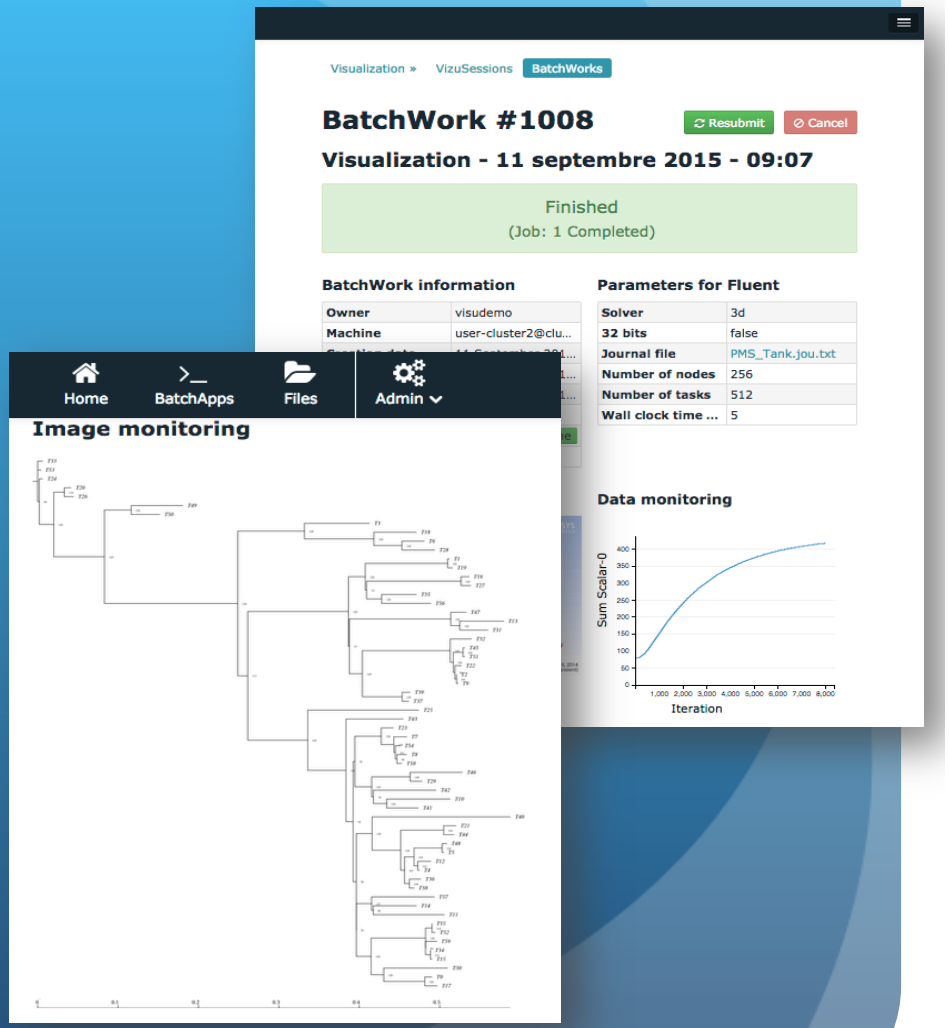

Job Submission

- Submission form specific per application
- Single or multi-step jobs
- Output preview

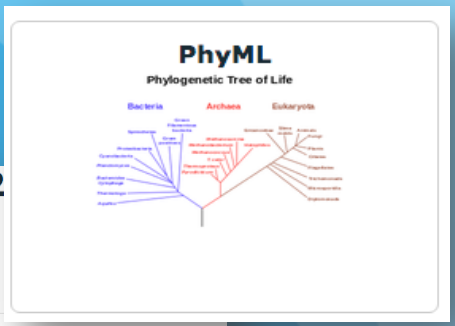

New BatchWork (step 2

Selected application: PhyML

Input Data

Sequence file * — Please select a sequence file

Data type * — DNA / Amino-acids

Sequence file type * — interleaved

Number of datasets * — 1

Substitution model

Substitution model * — GTR

Equilibrium frequencies * — optimized / empirical

Transition/transversion ratio — 0,00 — Fixed / Estimated

Proportion of invariable sites — 0,00 — Fixed / Estimated

Number of substitution rate categories * — 4

namd13 - 31 juillet 2015 - 11:46

Finished
(Jobs: 3 Completed)

BatchWork information

Owner	rguillier
Machine	phyml01@babel
Creation date	31 July 2015 - 11:46:42
Start date	31 July 2015 - 11:48:34
End date	31 July 2015 - 12:00:29
Machine	babel
Submission script	view
Progression	3 / 3 job(s) done

Parameters for namd13-job-multistep

Input Tarball	test.tar
Time Limit	10
CPU Number	1024
Initial Step	0
Last Step	2

There are 3 jobs attached to this BatchWork

Completed — Job #2544

Completed — Job #2545

Completed — Job #2546

Remote file management

Users can access their remote files on Babel

- Upload/download
- Preview results
- Copy/move/delete
- Rights

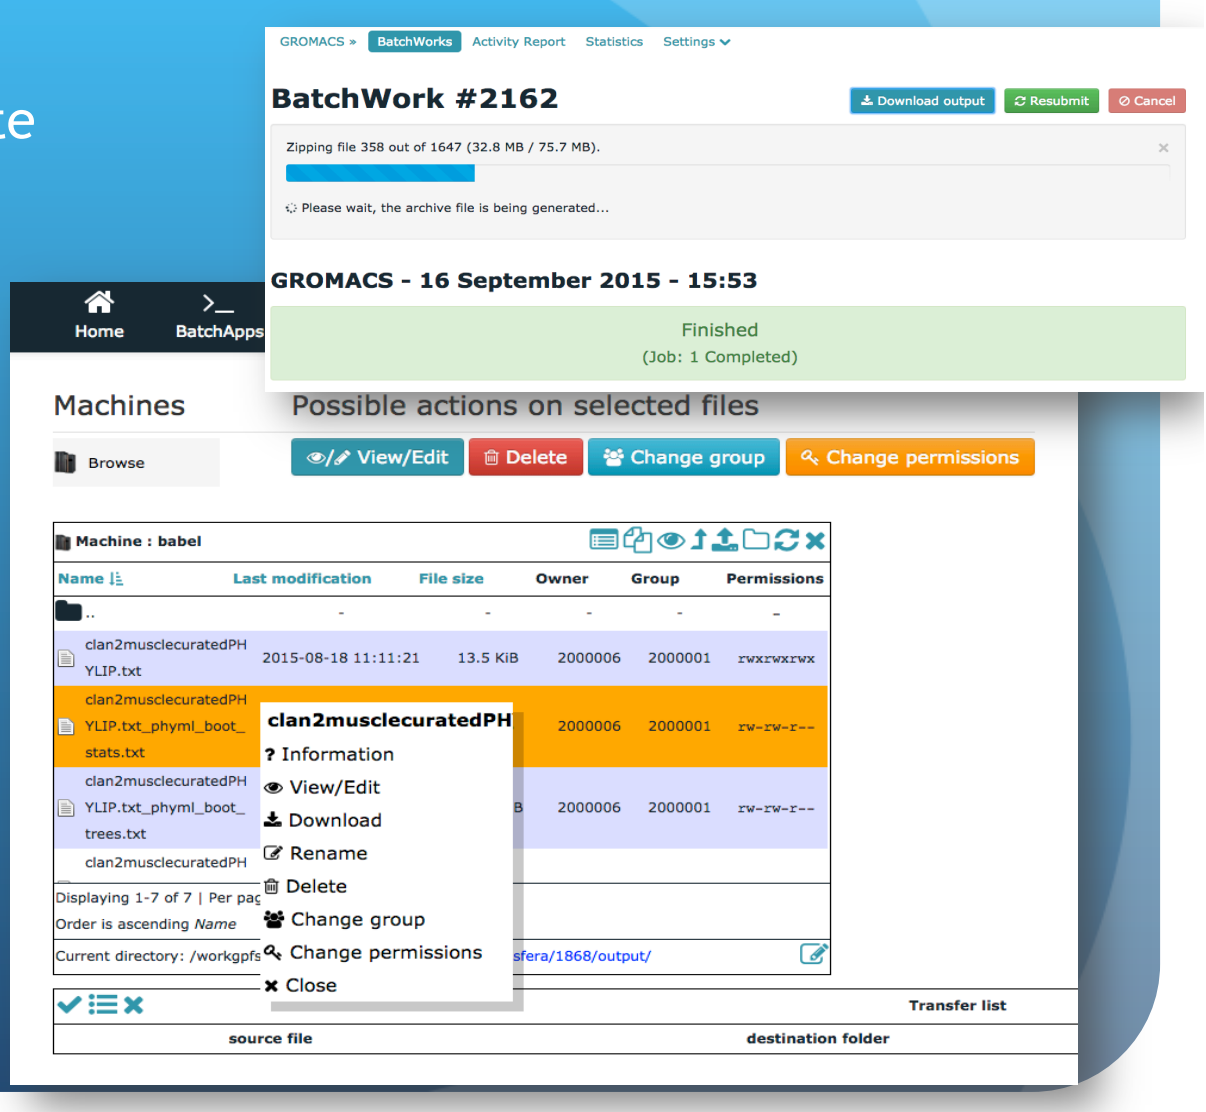

Architecture

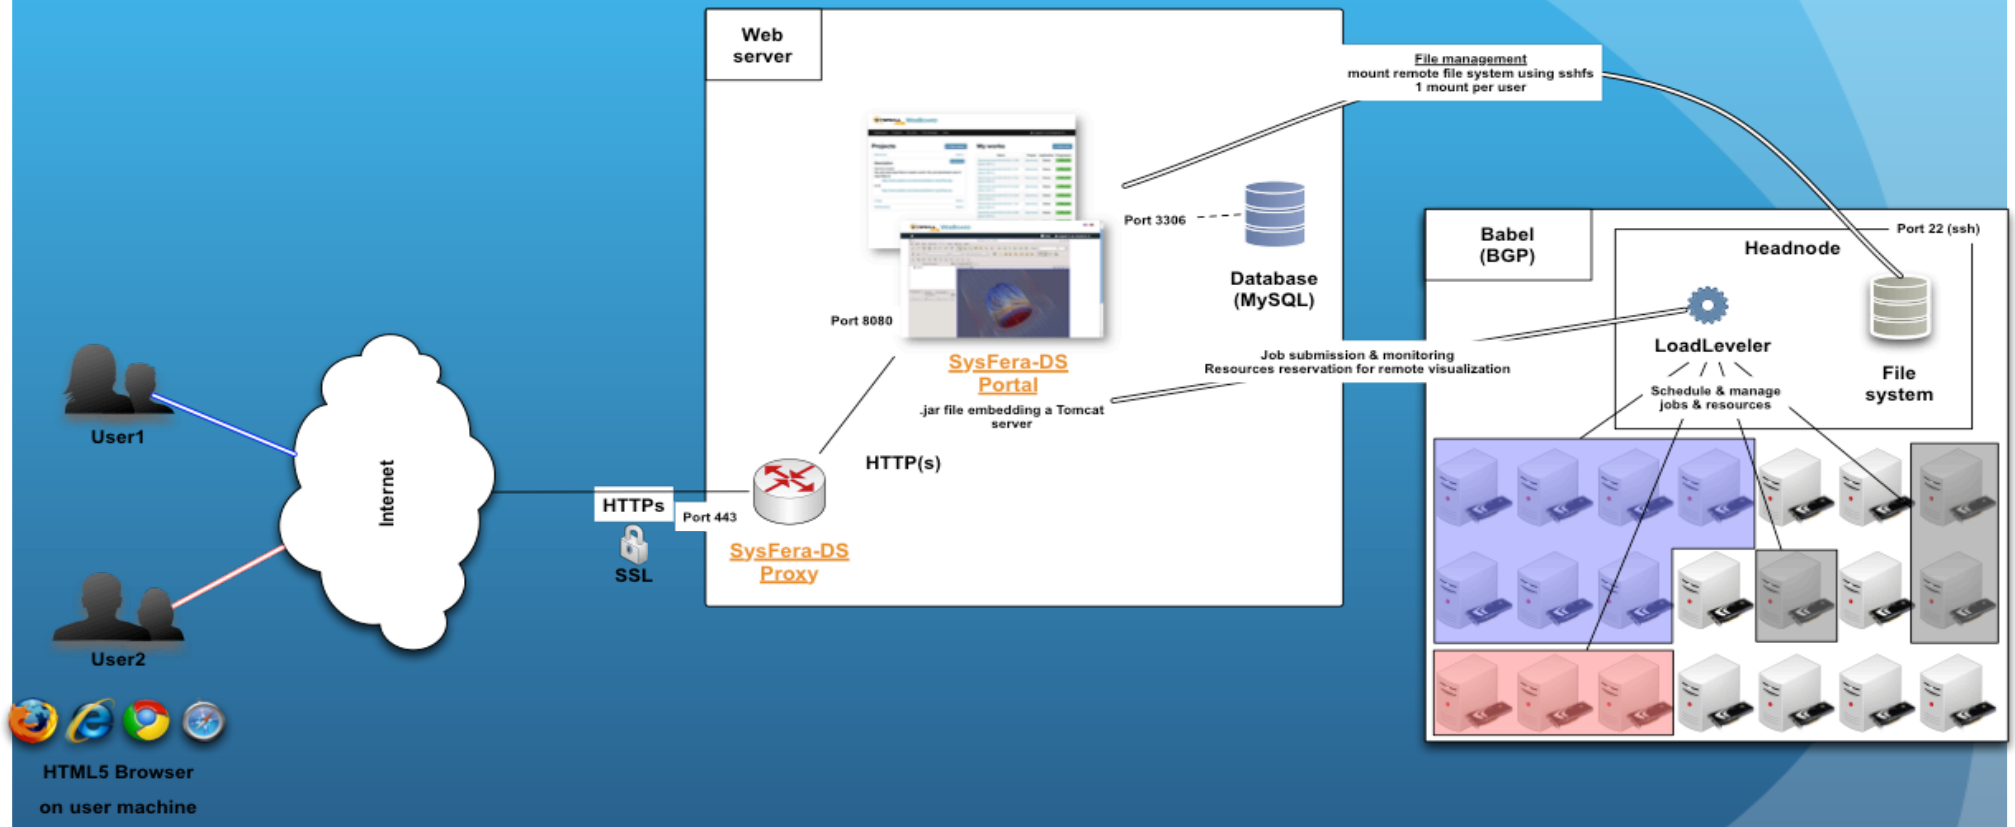

Applications

PhyML

A simple, fast, and accurate algorithm to estimate large phylogenies by maximum likelihood (ML)

- Inputs :
  - Data : a file containing DNA or protein sequences.
  - An evolutionary model (chosen by the user).
- Principle :
  - Compute the probability to observe the data given the evolutionary model and a phylogenetic tree.
- Complexity problem :
  - For *n* sequences, the number of possible trees is $O(n^n)$.
  - Cannot compute the likelihood for all trees.
- PhyML algorithm :
  - Based on a hill-climbing approach, which modifies the tree so as to maximize the likelihood.
- PhyML citations :
  - PhyML paper is the most cited in ecology-environment since 2007 (see Science Watch).
  - It was cited more than 10 000 times.
- Some application fields :
  - Genomics : to predict gene function and identify therapeutic targets.
  - Medicine : to trace the origin of epidemics and prevent new pandemics.
  - Ecology : to inventory the biodiversity and to preserve environment.

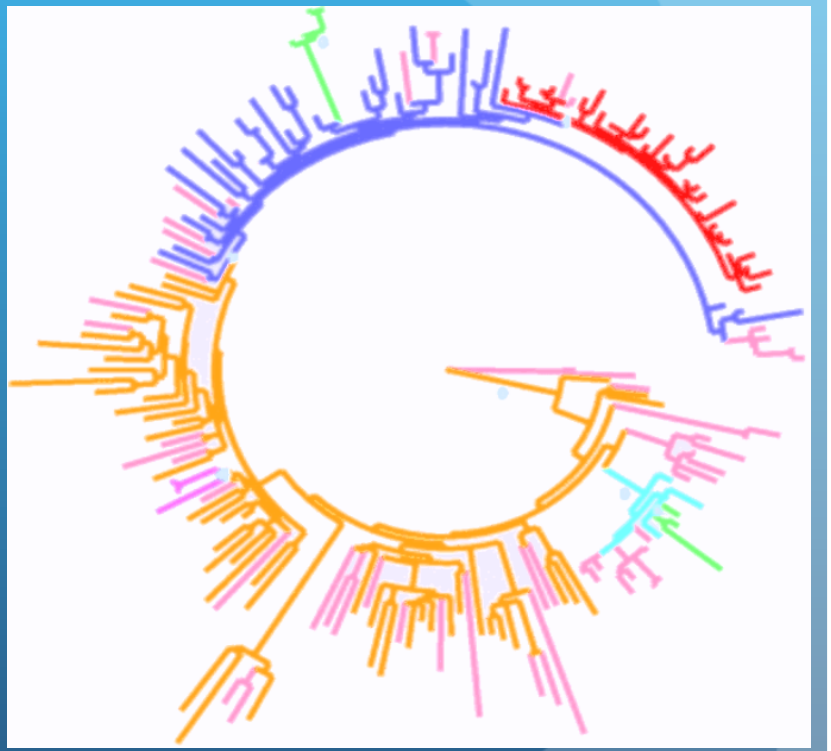

Applications (con't)

- NAMD
  - Parallel molecular dynamics code for high-performance simulation of large biomolecular systems: originally developed at University of Illinois.
  - Designed to scale on a large number of cores (over 500,000) and to process complex structures at atomic-level detail, such as the HIV capsid that contains more than 1,300 proteins and 64 million atoms.
  - Widely used software in the field of molecular dynamics, it represents 40% of the monthly computing time E-Biothon platform.
- LAMMPS
  - Massively parallel simulation tool for the movement of molecules
  - Developed by the Sandia National Laboratories.
  - Designed to efficiently compute Newton's equations of motion for collections of atoms, molecules, or macroscopic particles that interact via short- or long-range forces with a variety of initial and/or boundary conditions.
- GROMACS
  - A molecular dynamics package primarily designed for simulations of proteins, lipids and nucleic acids that have a lot of complicated bonded interactions.
  - Initially developed by the Biophysical Chemistry department of University of Groningen.
  - Able to simulate the Newtonian equations of motion for systems with hundreds to millions of particles using CPU or GPU.

Use of the platform

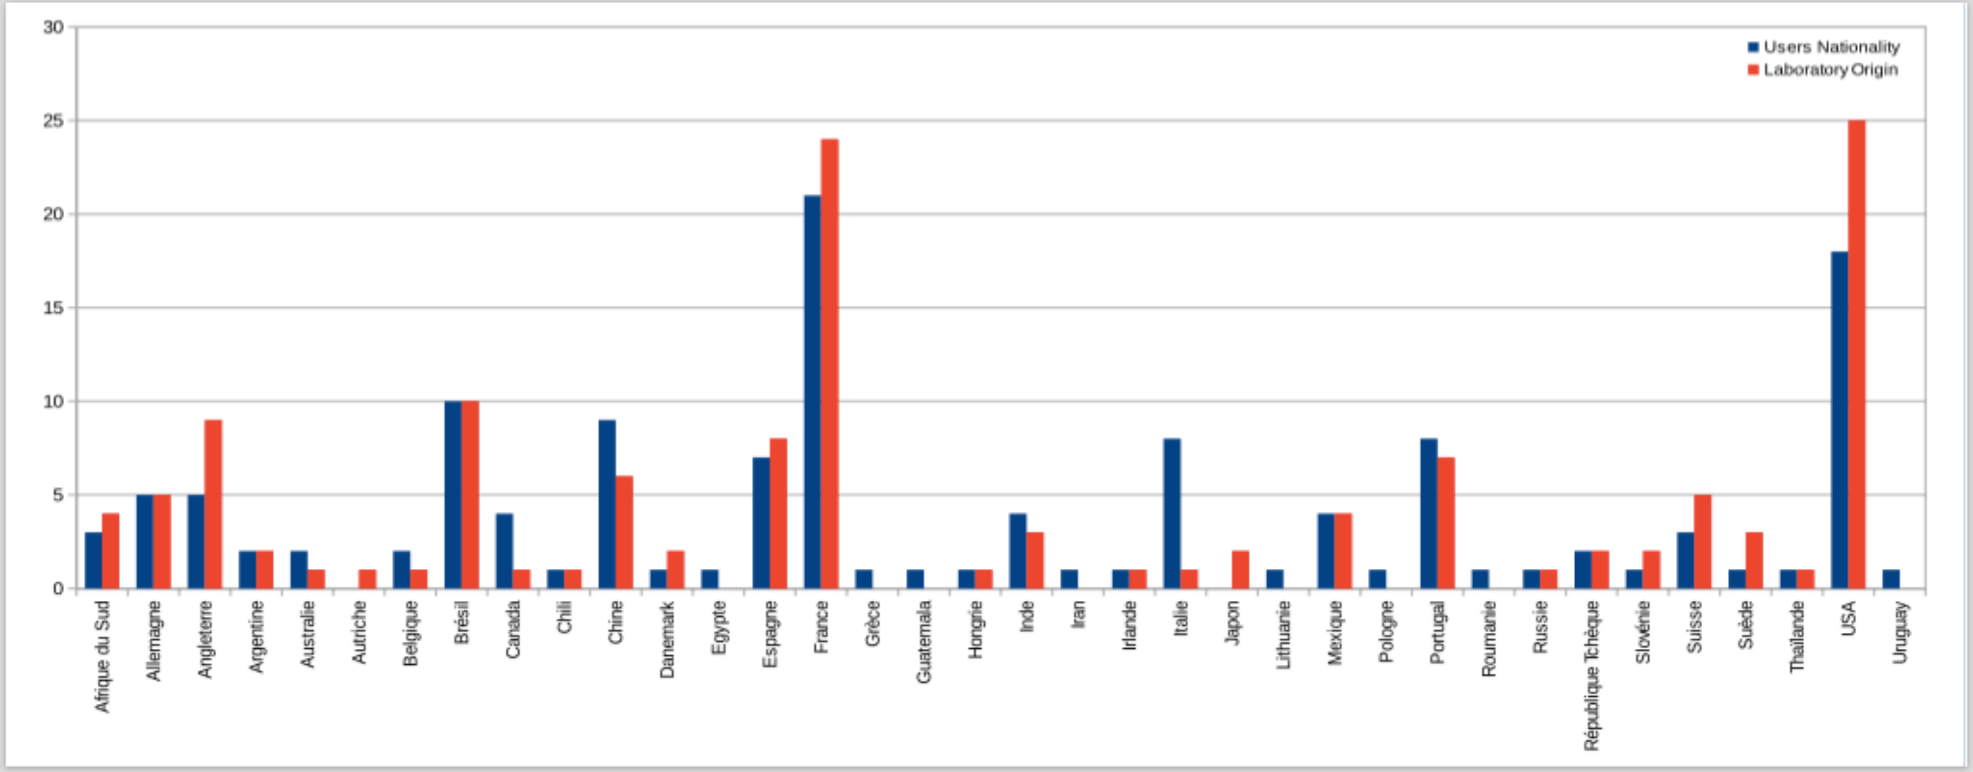

Use of the platform (con’t)

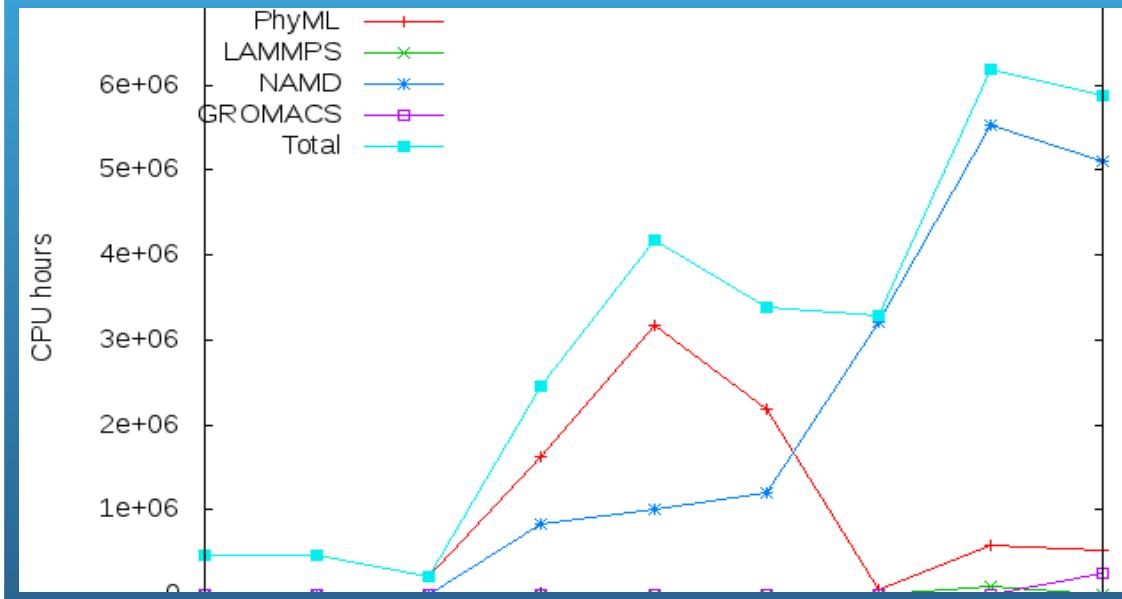

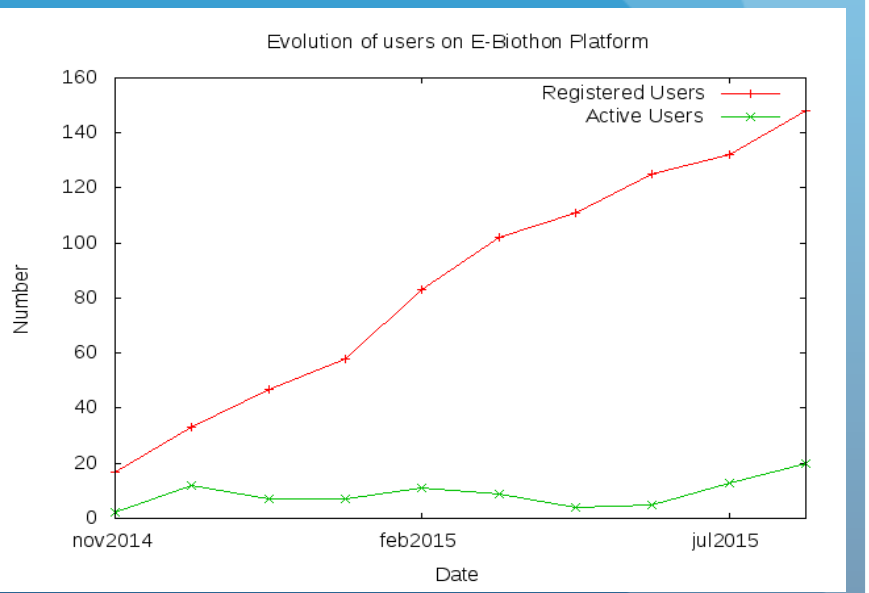

Experiments and first results

Correlated Motions in Distinctive Conformational States of the Chemokine Receptor CXCR4

B. Taddese and M. Chabbert

UMR CNRS 6214 – INSERM 1083, University of Angers, France

CXC chemokine receptor 4 (CXCR4)

- G protein coupled receptor
- A single ligand: **CXCL12**
- Involved in inflammation, chemotaxis, neural development and cancer
- co-receptor for HIV-1 viral entry.

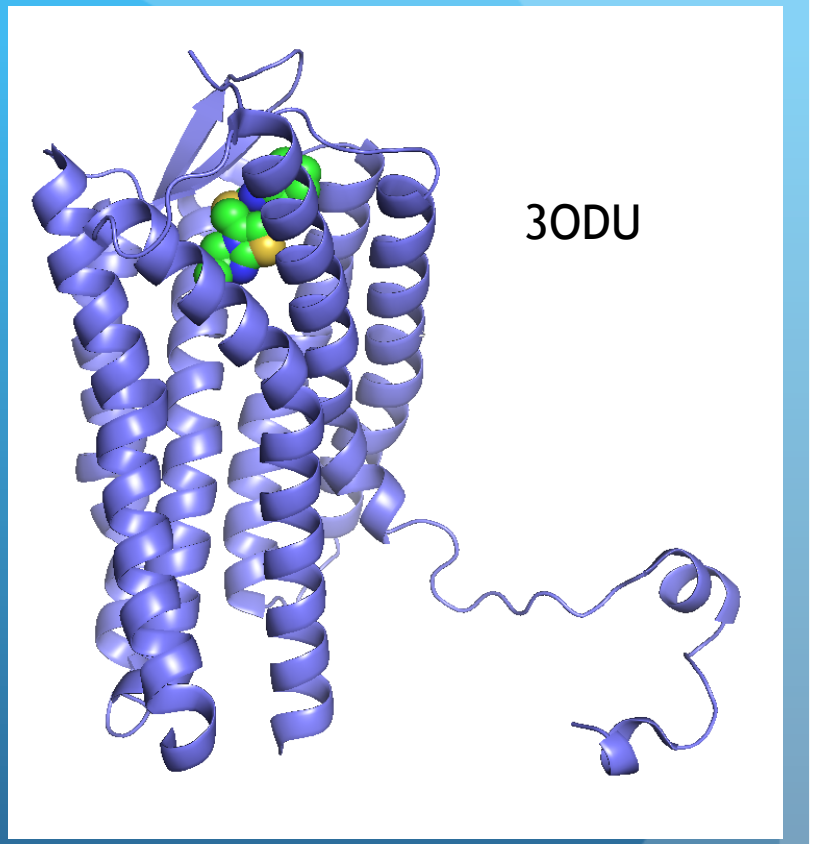

Important drug target

Resolved crystal structure

Prototype of chemotaxic receptors

MD System preparation

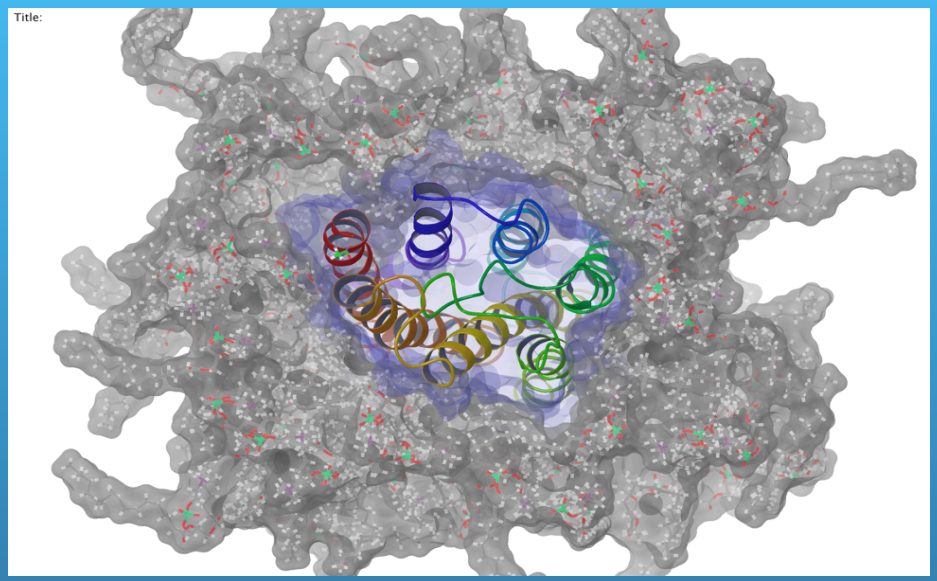

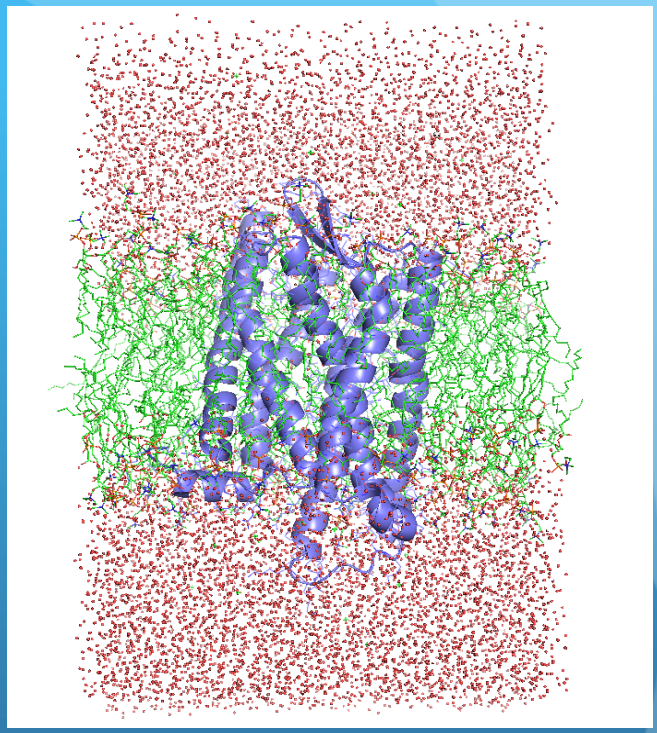

- Crystal structure of inactive CXCR4 (Wu et al., 2010)
- Models inserted into 120 POPC lipids
- Hydrated using TIP3 water models
- Cl^{2-} ions were introduced to neutralize the net positive charge
- Total of ~80000 atoms

How to access longer timescales with atomistic details

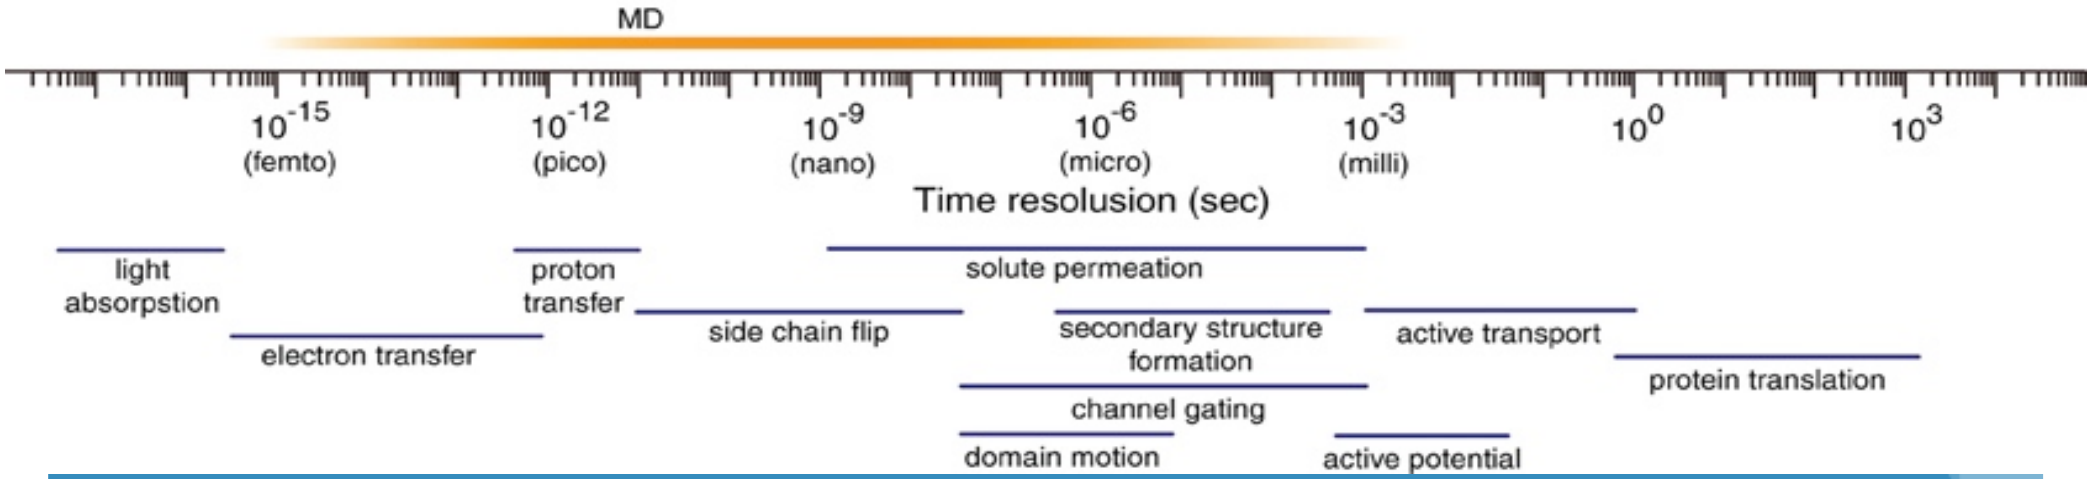

- Molecular dynamics (MD) can only reach nanoseconds to microseconds
- Activation GPCRs on millisecond timescales
- Sampling of conformational space can be improved
  - Increase computational power
  - algorithmic improvements
  - methodological developments: **accelerated dynamics**

Accelerated MD

Enhancement of sampling by

- adding a positive boost potential to the potential energy surface, effectively decreasing the energy barriers
- thus accelerating transitions between the low-energy states.

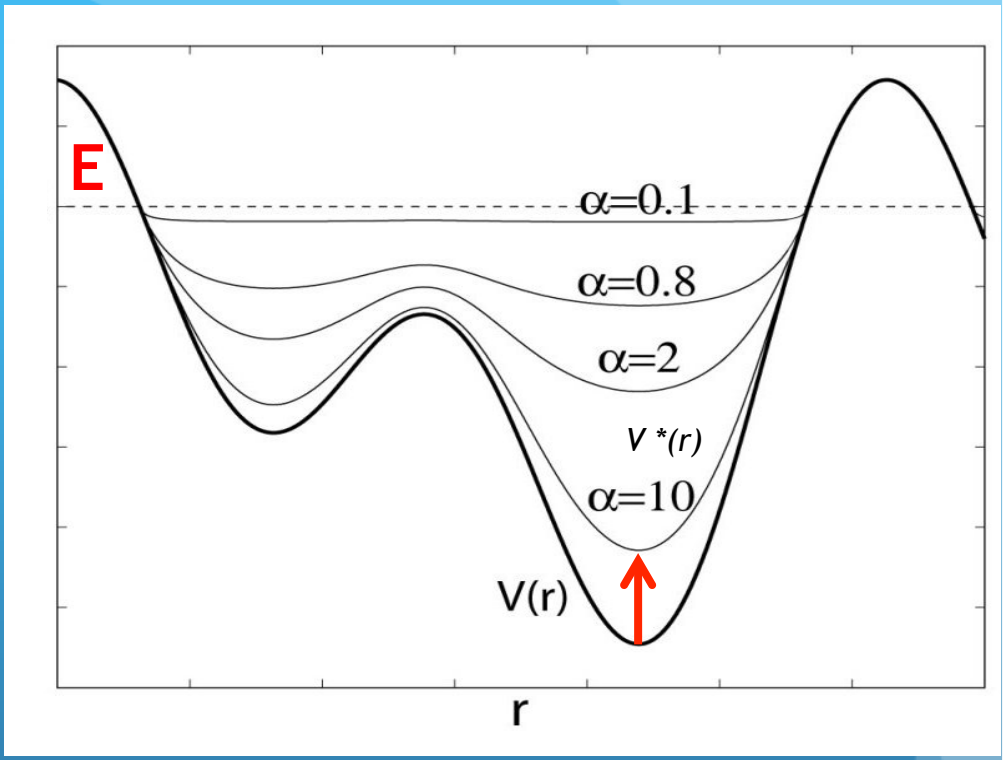

- ***V(r)*** **-** original potential
- **E** - reference energy: **calculated from cMD**
- ***V *(r)*** - modified potential
- **α** is the acceleration factor establishing the shape of the modified potential

Advantages:

- Simplicity: only two parameters are required: **E** and **α**
- It maintains the approximate shape of the underlying (or “unaccelerated”) free energy landscape
- It does not require the definition of a “reaction coordinate”

aMD allows observation of conformational transitions

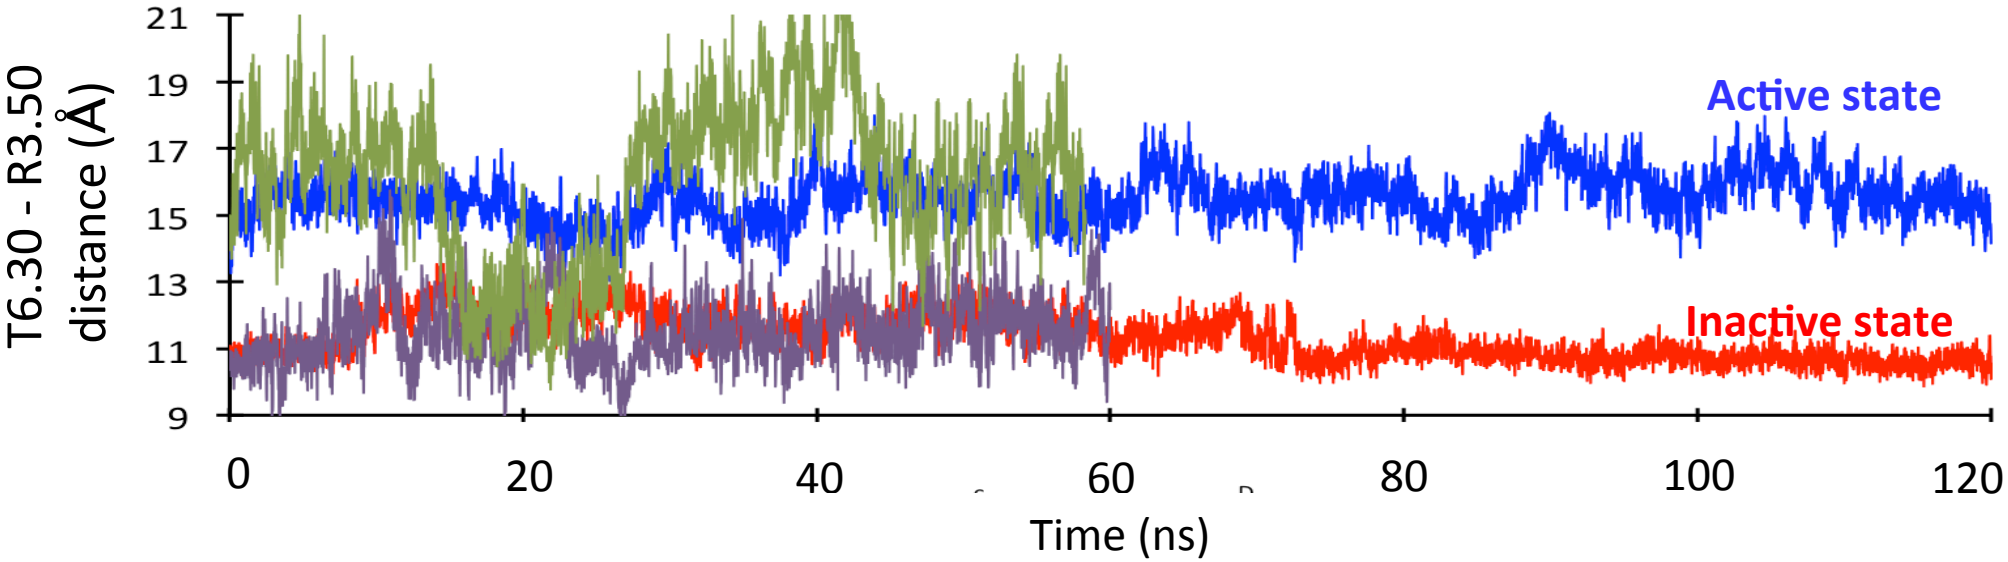

In the **aMD active model run**, **inward** movements of TM6 towards TM3, resembling GPCR inactivation.

Inversely, in the **aMD inactive model run**, **outward** movement of TM6.

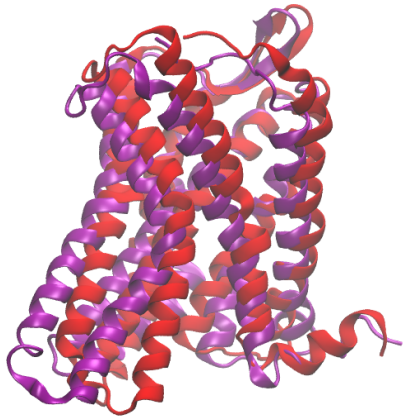

cMD inactive
aMD inactive

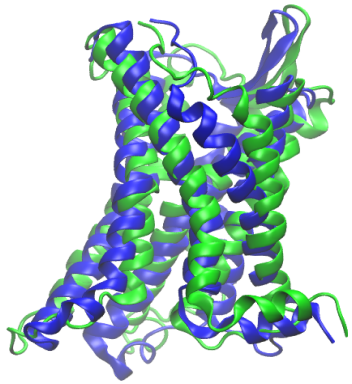

cMD active
aMD active

Sampling of the conformational space

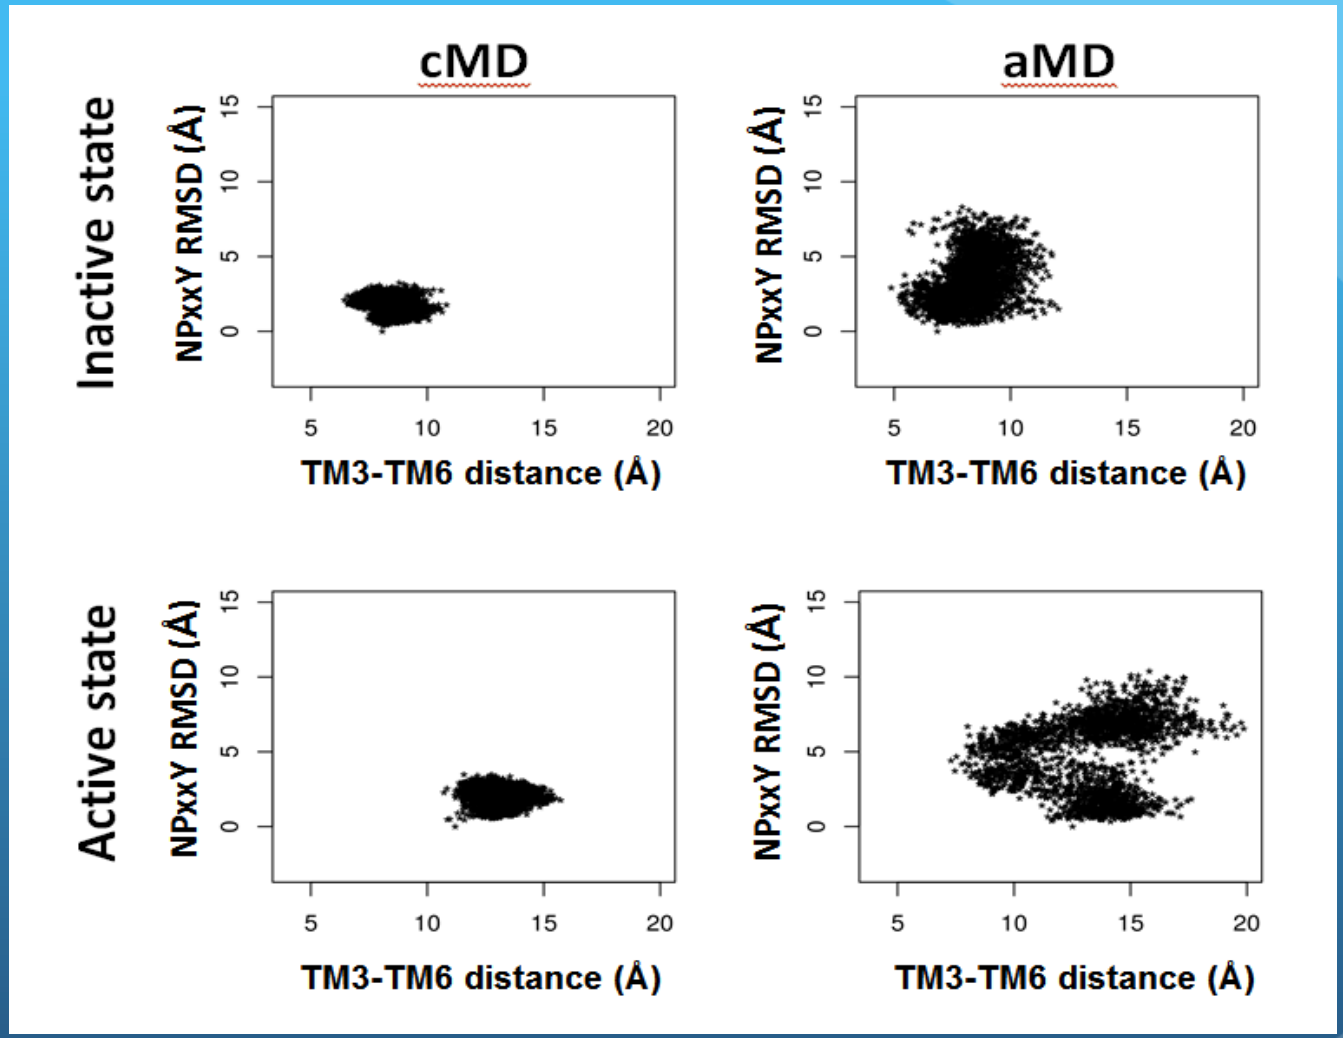

Conclusions

- aMD is an efficient method to increase conformational sampling
- aMD increases the probability of sampling large amplitude conformational transitions by several orders of magnitude

Prospectives

- Analysis of the conformational transition pathway(s)
- Search of biased ligands towards specific conformational states

Biodiversiton project

Alain Franc
INRA - UMR BioGeCo /
Inria - Pleiade Team

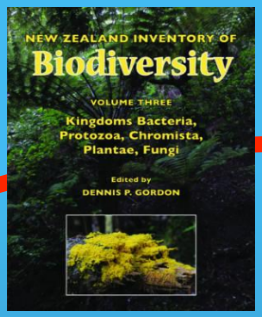

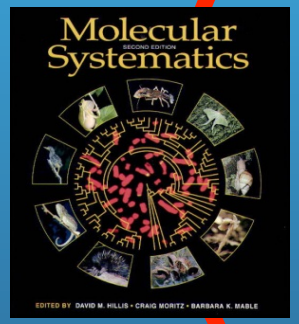

+

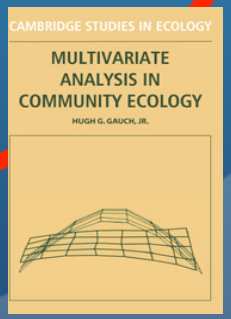

Taxonomy on Edit distance

Definition: The edit distance between two strings is defined as the minimum number of edits needed to transform one string into the other, with the allowable edit operations being insertion, deletion, or substitution of a single character.

kitten → **s**itten (substitution of 'k' with 's')
sitt**e**n → sitt**i**n (substitution of 'e' with 'i')
sittin → sitting (insert 'g' at the end).

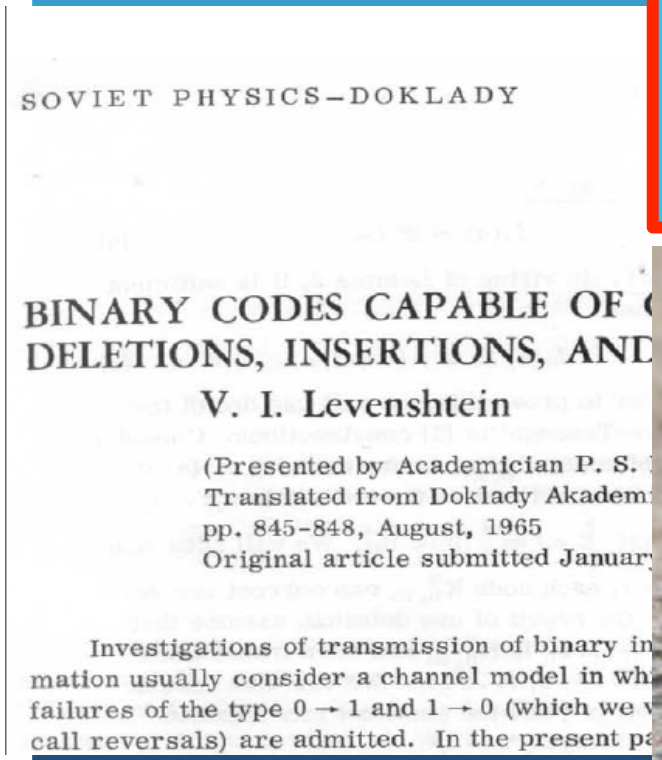

SOVIET PHYSICS–DOKLADY

BINARY CODES CAPABLE OF
DELETIONS, INSERTIONS, AND

V. I. Levenshtein

(Presented by Academician P. S.
Translated from Doklady Akadem
pp. 845–848, August, 1965
Original article submitted Januar

Investigations of transmission of binary in
mation usually consider a channel model in wh
failures of the type 0 → 1 and 1 → 0 (which we
call reversals) are admitted. In the present pa

A taxonomic annotation pipeline

Goal : a set of references (from NCBI, local, … (10^3 reads typically)
a set of queries NGS (10^5 reads typically)
computes the distance between each pair (reference × query)

The work flow Hyperparallelization of distances computations (all ≈ 10^8)
with MPI
on a BlueGene /P, and next /Q

Then annotation for each read
look at nearby references
if all ref ∈ same taxon
⇒ informative read
⇒ taxon present
else: disregarded

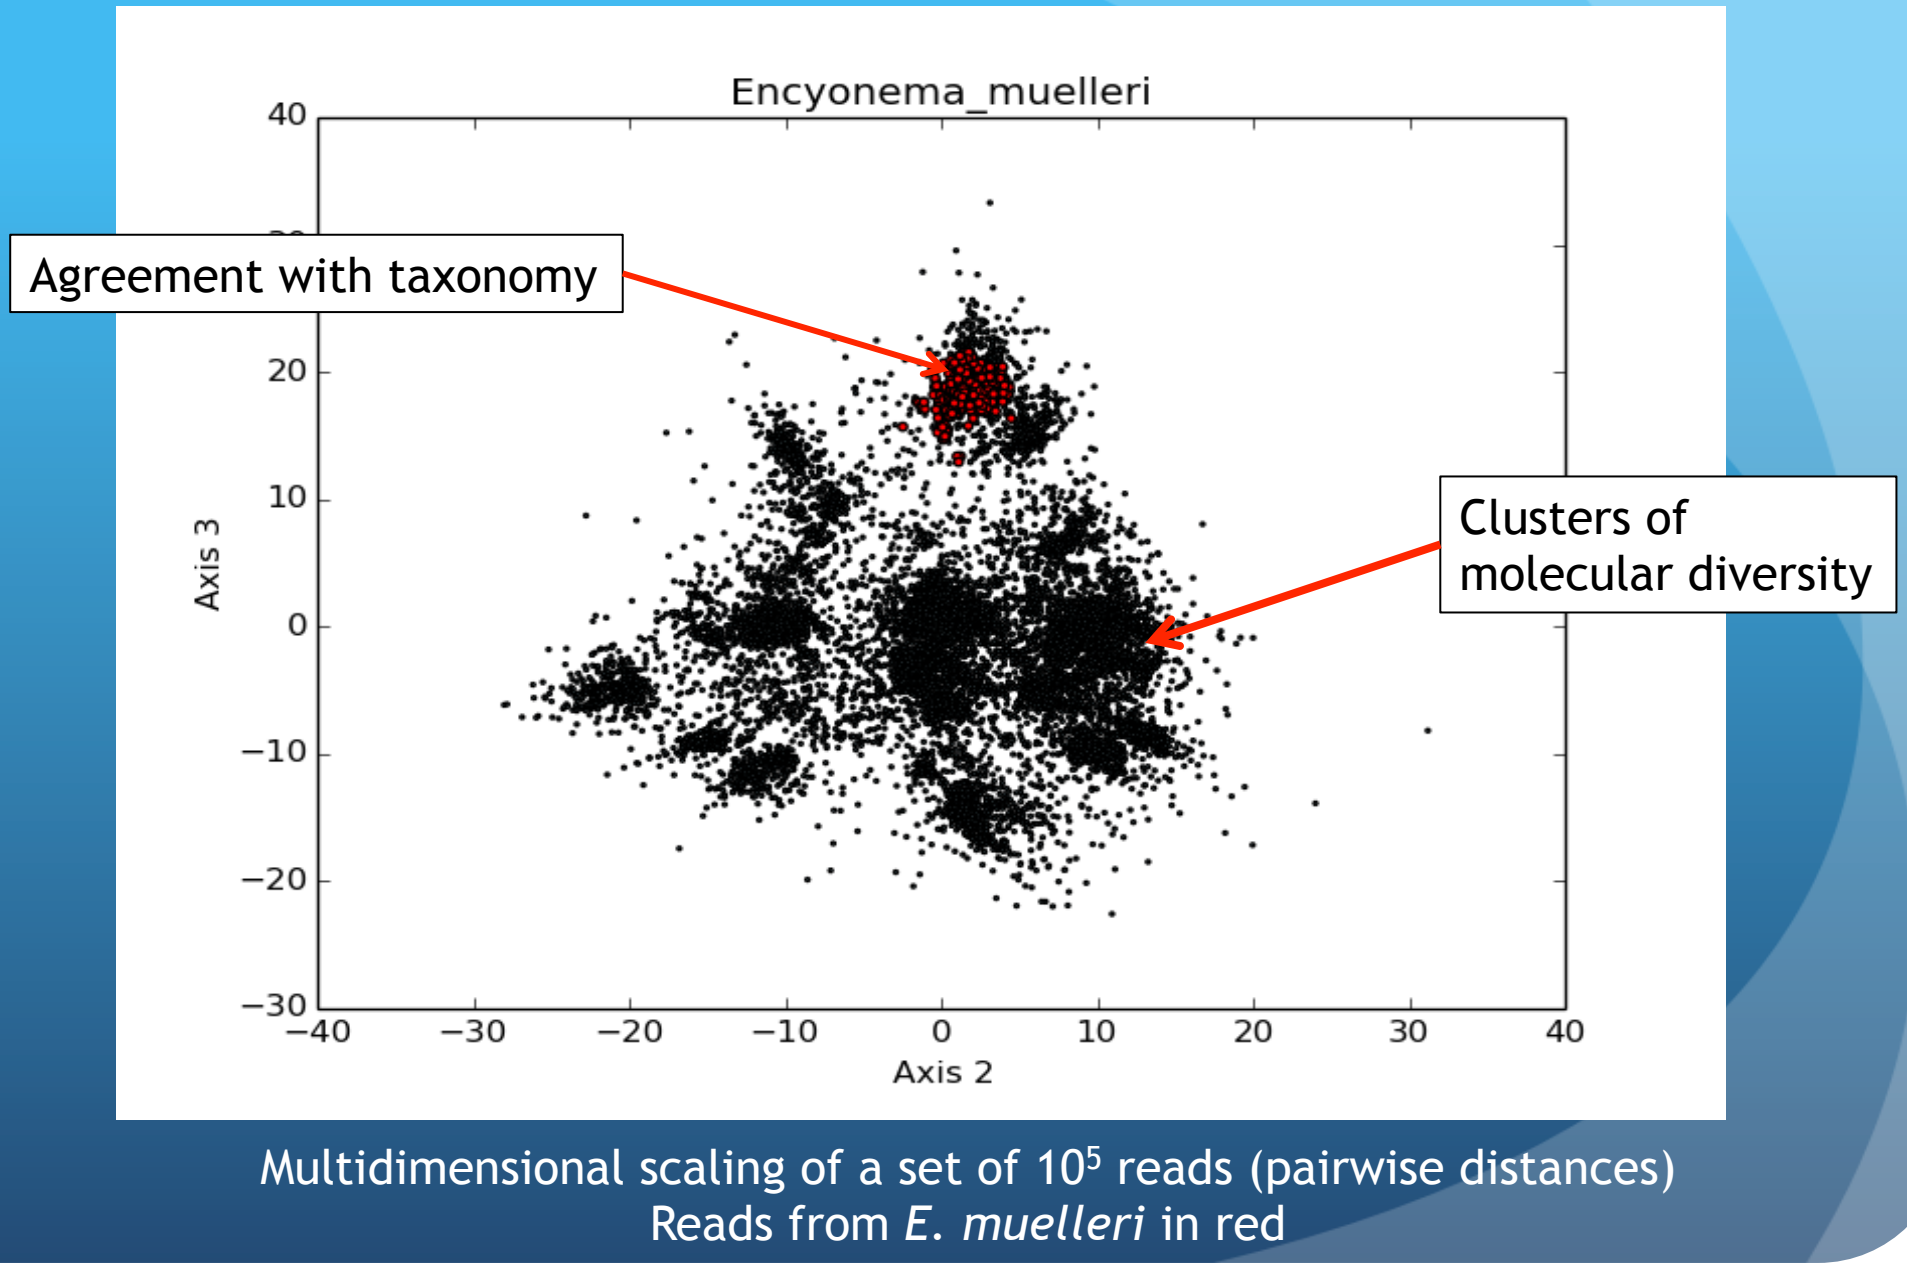

Multidimensional scaling of a set of 10^5 reads (pairwise distances)
Reads from *E. muelleri* in red

Insyght: a tool for bacterial comparative genomics

Jean-François Gibrat, INRA, UR1404, Unité Mathématiques et Informatique Appliquées du Génome à l'Environnement

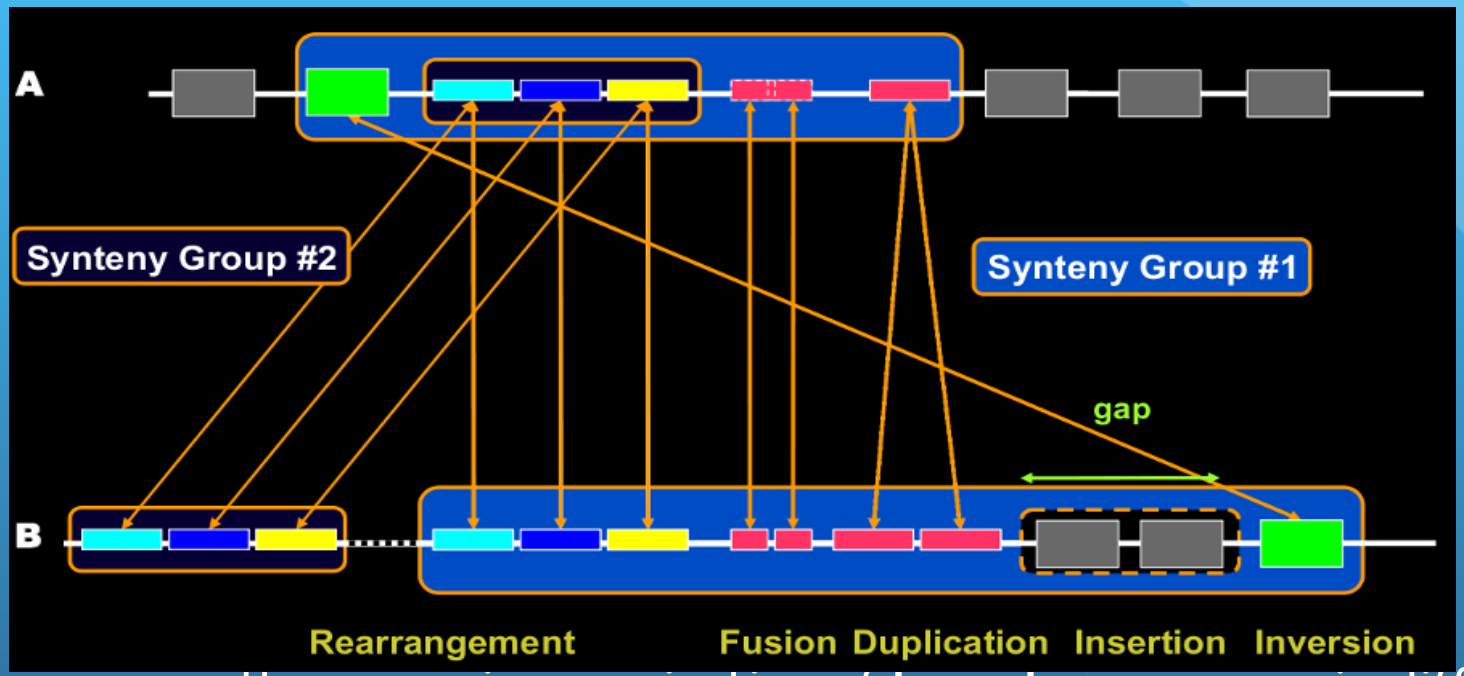

- Goal : compare the genomic organization of **homologous** genes in different bacterial genomes.
- This provides information about the gene function, the evolution of genomes (fusion, duplications, insertions, deletions).
- Comparative genomics is the **cornerstone** of genome analysis strategies.

Insyght components

- Insyght consists of:
  - A relational database that contains the data to visualize
    - Primary data that are parsed from files
    - Secondary and tertiary data that are computed from cross comparisons of the primary data
  - A pipeline of Perl programs to populate the database.
    - Heavy computations are sent to a supercomputer.
  - A web application (AJAX, HTML5 canvas) to visualize the data that tightly integrates 3 complementary views.
    - a table for browsing among homologs
    - a comparator of orthologs' functional annotations
    - a genomic organization view that improve the legibility of genomic rearrangements and distinctive loci
- http://genome.jouy.inra.fr/Insyght

Insyght computational loads

- Cross comparison of ~**2700** bacterial genomes, each having about **4000** genes on average,
  - first using BLAST (a well-known software for comparing genes)
  - then with an in-house software based on dynamic programming using the BLAST results.
- **>3.5 M** genome comparisons, each involving **16 M** genes
- BLAST jobs generated **1.2 TB** of raw, compressed data
- The relational database has a total size of **3.5 TB**, the largest table having ~**6 billion** rows and occupying ~**1 TB** of disk space (~**2 TB** with indexes).
- Visualization start-up time and most loading times take a few seconds even for whole genomes and multiple comparisons

Conclusion

- The E-Biothon platform is open to the life science community.
- It is still evolving for incorporating new applications.
- The success of this project is not only coming from the availability of a parallel computing platform since the user support provided by all the partners of the project (CNRS, IBM, IDRIS, Inria, Institut Francais de Bioinformatique and SysFera) is crucial for efficiently deploying new applications and managing the platform.
- Thanks to CNRS, IBM, IDRIS, Inria, l'Institut Francais de Bioinformatique [14] and SysFera [26] without which it would not have been possible to setup the project.
- We are also very grateful to all the people that have been experimenting the platform since the beginning of the project.

Conclusion

- Calcul intensif / données : *grand instrument scientifique pluridisciplinaire*, catalyseur de nouvelles connaissances scientifiques

- Besoins calcul / stockage toujours **en forte croissance** (d'où l'**explosion des demandes notée au niveau du CPER)**

- F**acteur majeur** de la dérive des coûts informatiques : **foisonnement d'infrastructures de calcul et de données au niveau local** (i.e. Tiers-3) aggravé par **l'augmentation des demandes non-coordonnées + morcellement** et de **désorganisation** des **infrastructures de données**

- *Stratégie du CNRS : coordination / rationalisation des investissements aux niveaux site / national autour de défis scientifiques avec l'ensemble des acteurs concernés*