IBM Deep Computing

Blue Gene/L Enabling Breakthrough Science - Computational Biology and Beyond

Gyan Bhanot

IBM TJ Watson Research Center
&
Institute for Advanced Study, Princeton

Outline

- Introduction to players
- Application which already Performe on Blue Gene
- Work in progress on Blue Gene
- The Genographic Project – a possible BG project
- Remarks

‘A Mathematician is a device to turn coffee into Theorems’

Paul Erdös

‘A Physicist is a device to turn Biology into Philosophy’

Anonymous

Biologists/Chemists Beware

- A physicist, mathematician and engineer are hunting deer. How do they identify it?
- Physicist: observed that it behaved in a deer-like manner, so it must be a deer.
- Mathematician: asked the physicist what it was, thereby reducing it to a previously solved case.
- Engineer: was in the woods to hunt deer, therefore it must be a deer.

Biology of Transcriptional Regulation
Boston University – Department of Bioinformatics

- Transcription Factors (TFs) bind DNA upstream of a gene and promote or inhibit RNA transcription

- Genes bound by the same TF can be co-regulated

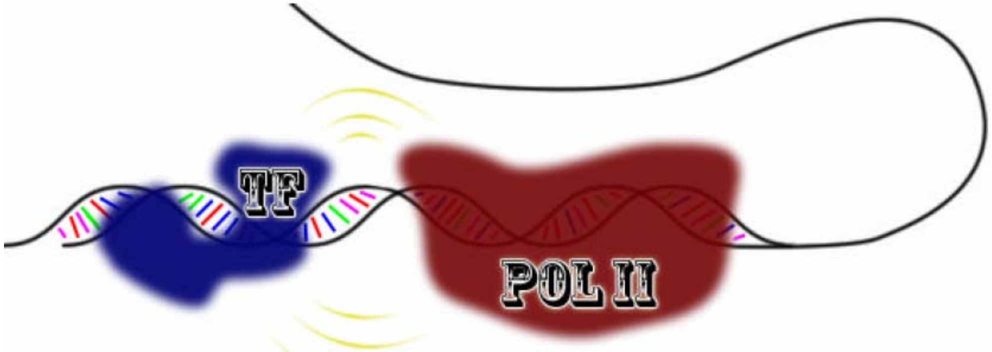

Goal

- **Identify both the TFs and the places they bind (i.e. the genes they regulate)**

- **Identify sets of gene regulated by the same TF**

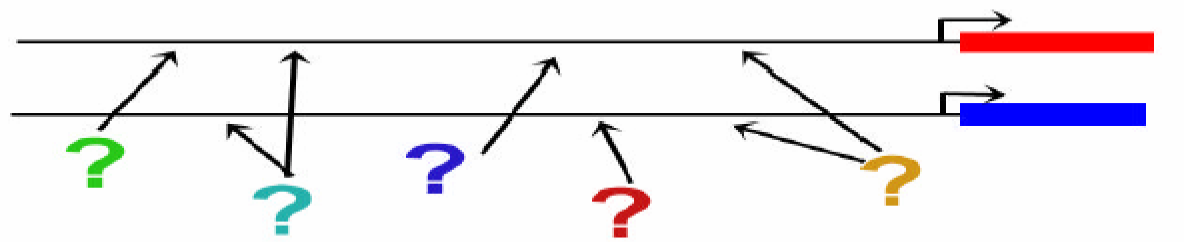

Gibbs Sampling

- Pick initial positions in promoter hypothesized to contain a common binding site

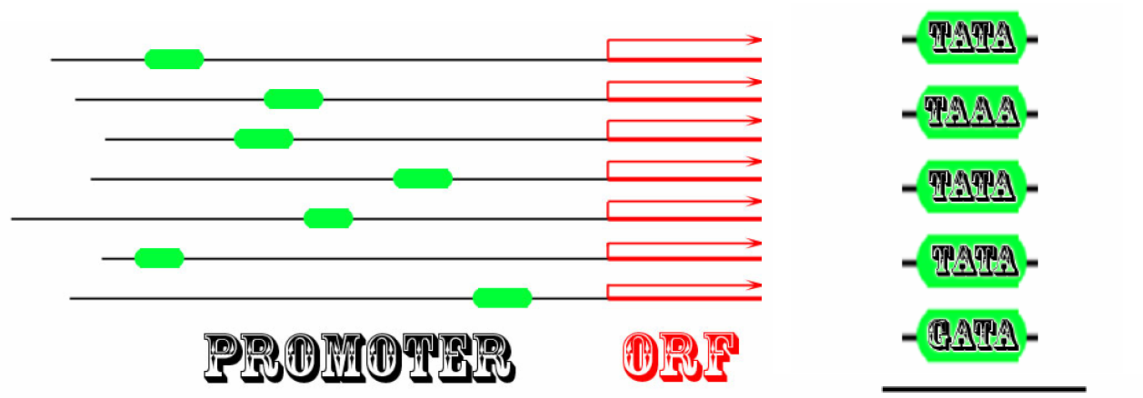

- Calculate a PWM

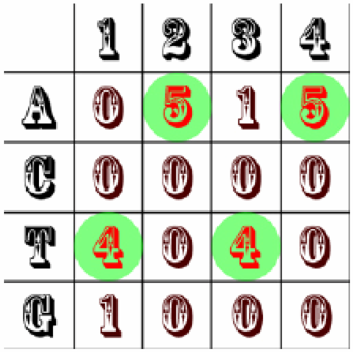

	1	2	3	4
A	0	5	1	5
C	0	0	0	0
T	4	0	4	0
G	1	0	0	0

- In general, the score we want to optimize is the conservation of the PWM

- Iteratively update positions, optimizing PWM

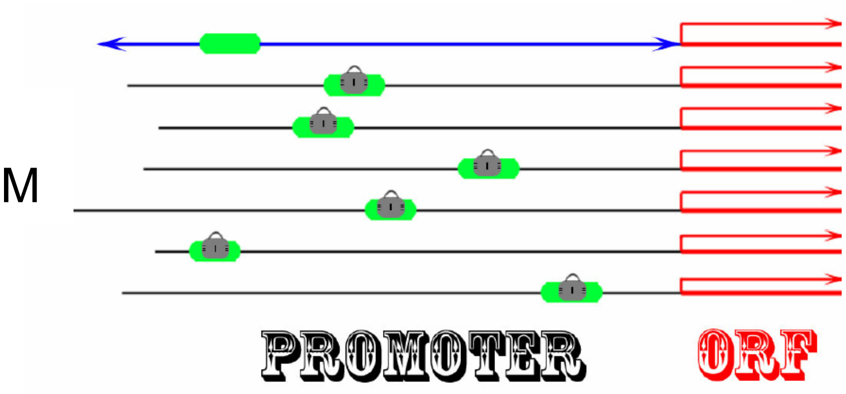

GibTigs BlueGene/L Implementation

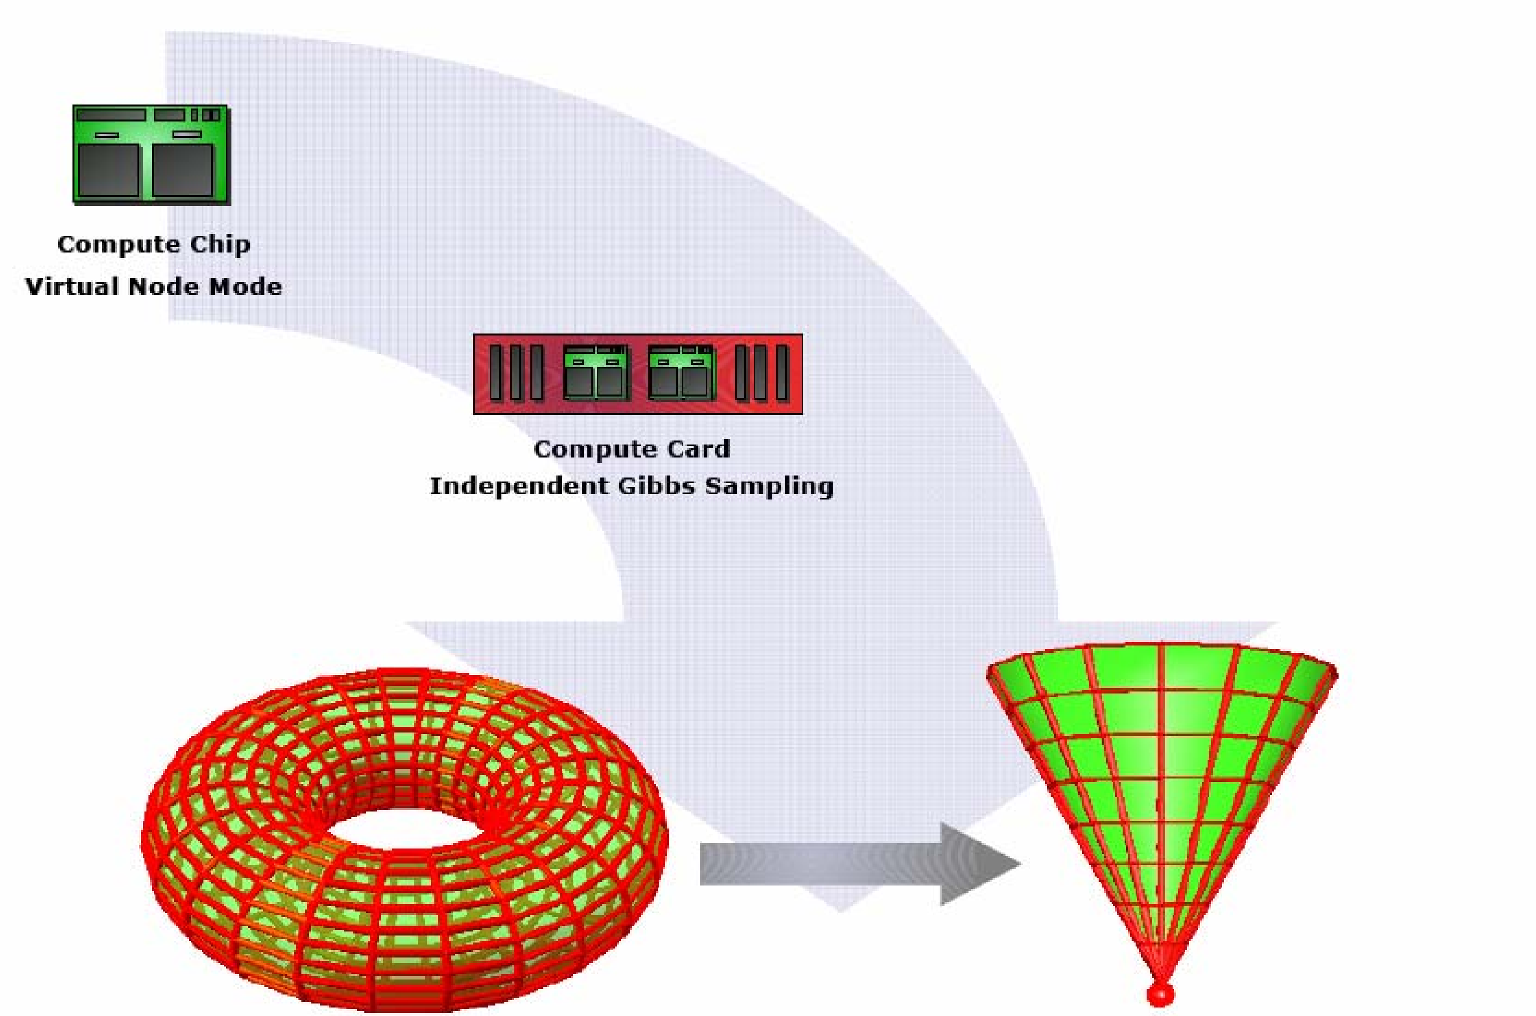

Compute Chip
Virtual Node Mode

Compute Card
Independent Gibbs Sampling

BlueGene/L Rack (Torus Network):
Massively Parallel independent Gibbs Samplings

BlueGene/L Rack (Tree Network):
Data Collection and Clustering

Blue Gene/L Gives Insight into Genomic Regulation

- Redeployment of GibTigs on BlueGene/L systems redefines Boston University's Biomolecular Systems Lab development cycle.
  - One BG/L rack = **~ twenty fold increase** in compute power.
  - GibTigs showing linear scalability up through 2048 CPUs, one BG/L rack in virtual-node mode - sped our development cycle from a few runs a week to many runs a day.
    - Doing so has **enabled large scale parameter searches**, and regular production grade performance evaluations.
    - The results are **drastic improvements in sensitivity** and specificity of our algorithm, none of which would have been possible at our previous development pace.
    - Rather than making conservative modifications to GibTigs, the power of BlueGene has given us the **freedom to take risks in trying new ideas**, many of which have failed, but some of which have provided new insight and new power to GibTigs.
  - **GibTigs has recently proved to be**, according to published measures, **the most powerful predictor of DNA transcription regulatory sites to date**.
- "Having achieved a performance milestone, the scalability of BlueGene/L has encouraged us to think even bigger in our research. "

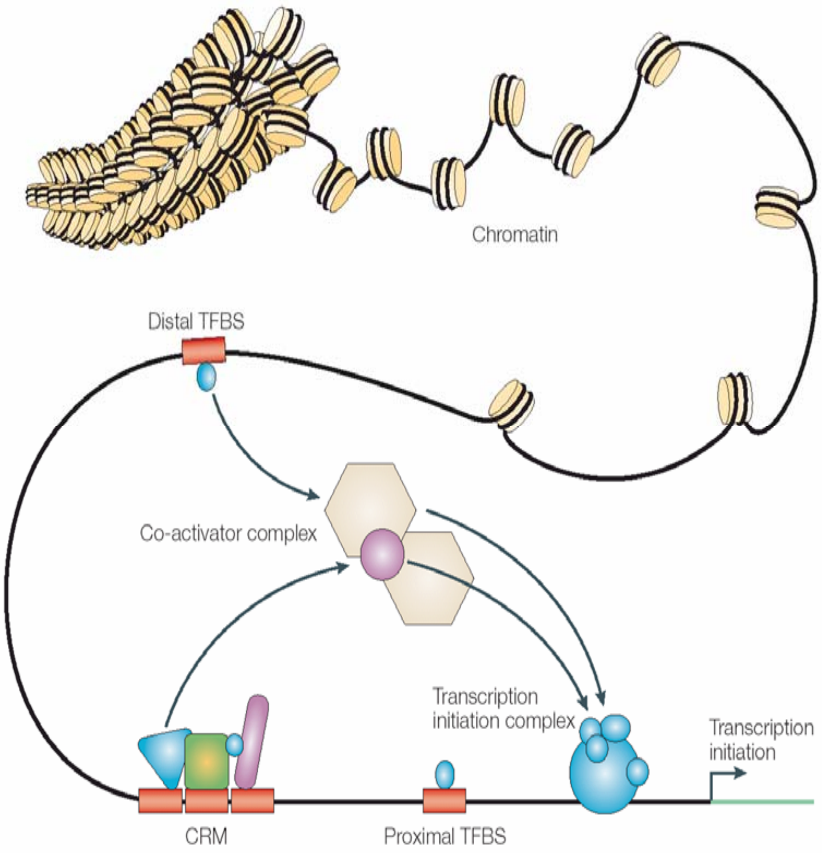

Maize Genome Assembly: Srinivas Aluru - Iowa State

- Maize (i.e. corn) is an economically important crop in the US
- Best studied model organism for the cereal crops.
- Just as the human genome project will intensify upcoming medical advances, cereal genomes (rice and maize) will help improve worldwide food production.
- Maize genome is comparable in size to the human genome (2.5 GB) but is highly repetitive (65-80%). Less than 10-15% is gene space.

Genome Shotgun Assembly

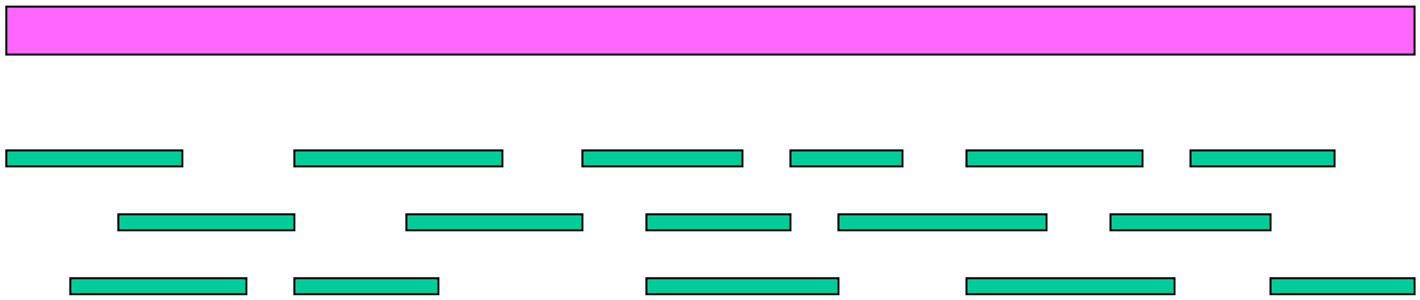

- Identify pairs of fragments that have good exact match (promising pairs).
- Perform alignments on promising pairs.
- Use results to perform assembly.

Genes Are Not Uniformly Sampled

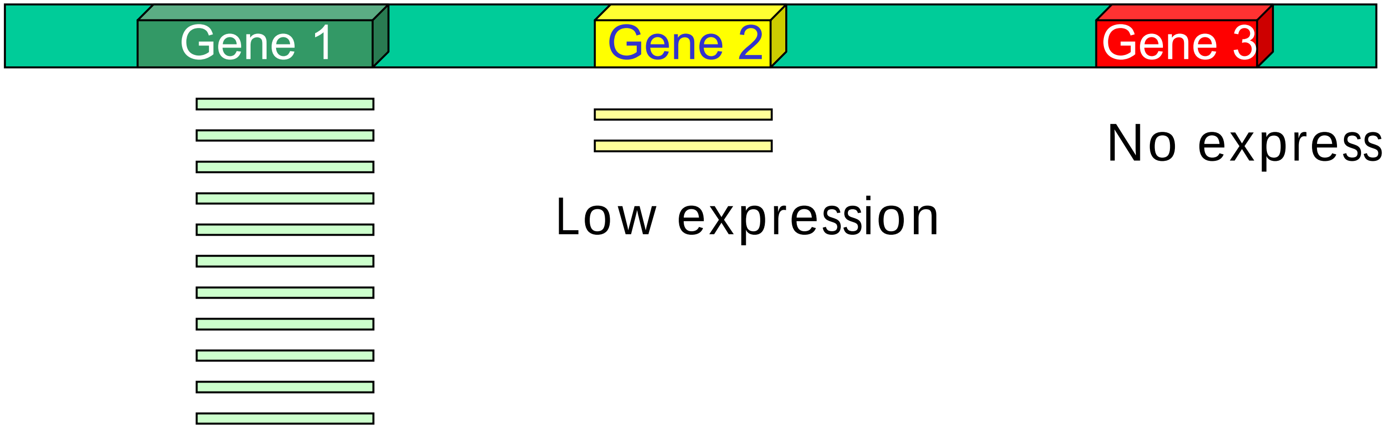

No expression

Low expression

High expression

Random vs. **Biased Sampling**

Uniform Sampling

Nonuniform Sampling

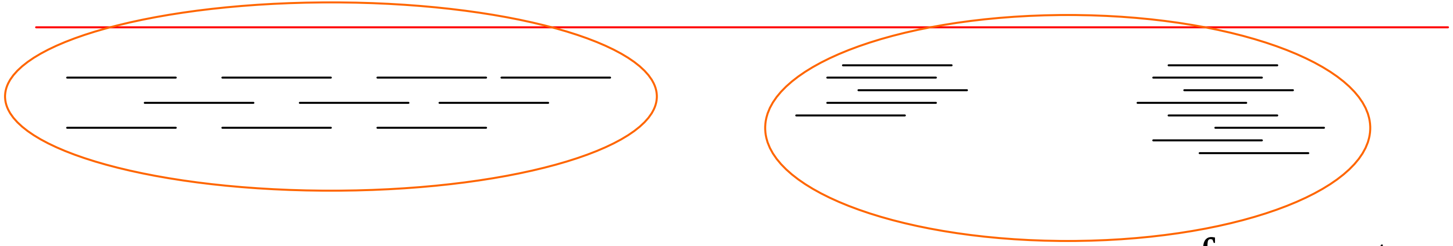

fragments

- Uniform sampling – $O(n)$ promising pairs
- Non-uniform sampling – $O(n^2)$ promising pairs

How to run on a massively parallel system

- Reduce space requirement from quadratic to linear.
- Generate promising pairs in decreasing order of maximal common substring length.
- Constant time per generation of a pairwise maximal common substring.
- Significantly reduce number of alignments without affecting quality.
- parallel processing –
  - reduce run-time;
  - increase available memory.

Clustering Methodology

- Two sequences fall in the same cluster if there is a chain of overlaps that leads from one sequence to the other.
  - Initially, treat each sequence as a cluster by itself.
  - If two sequences from two different clusters show significant overlap, merge the clusters.
  - Use union-find data structure.

PaCE Software Architecture

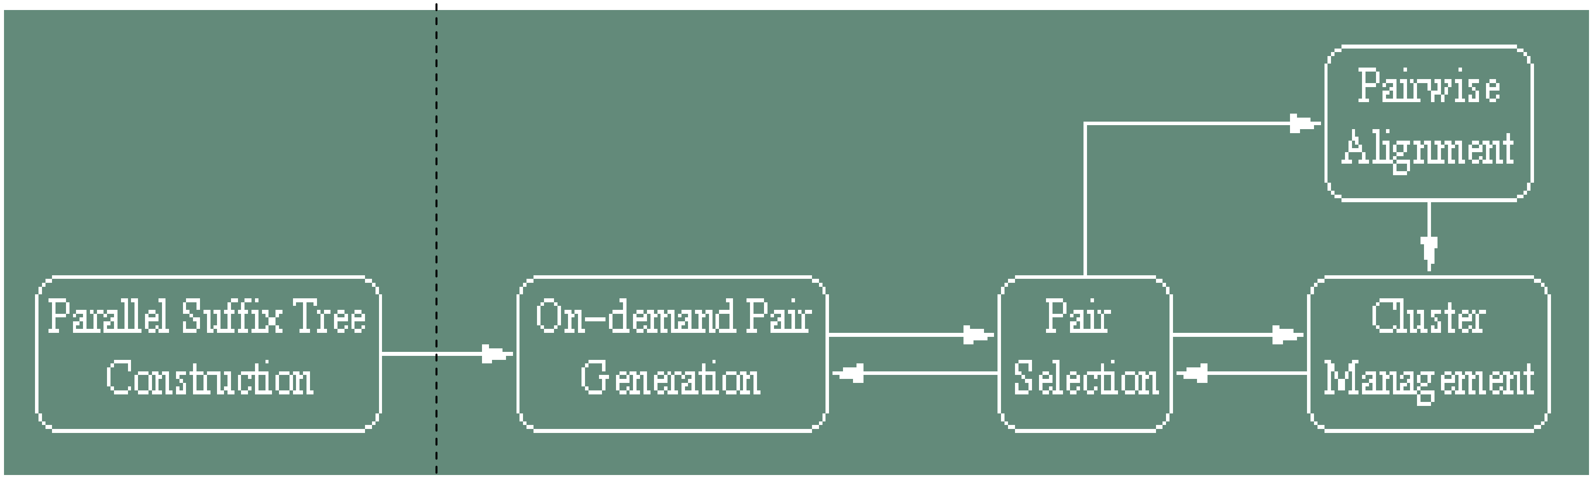

Construction/ Preprocessing Phase

Parallel Clustering Phase

Generalized Suffix Tree (GST) Construction on BlueGene/L

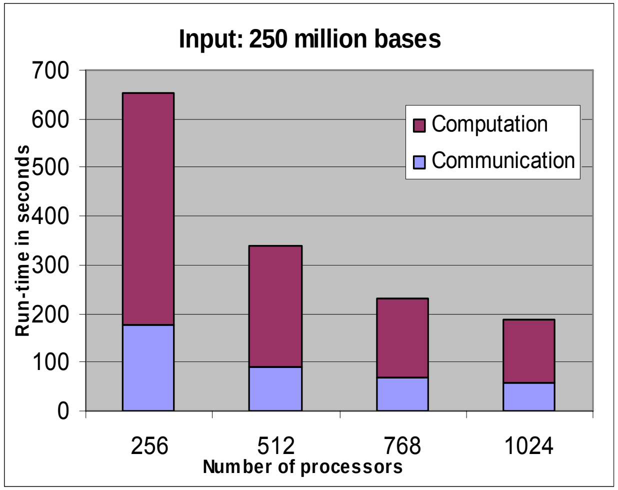

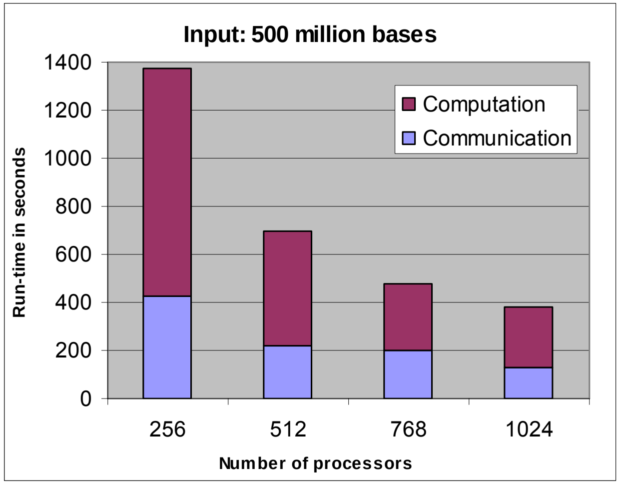

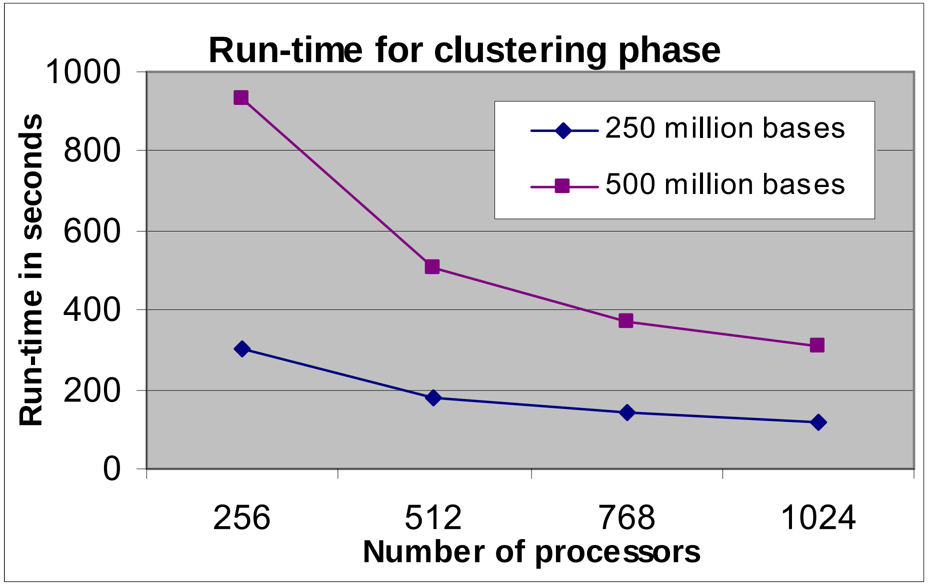

Latest Assembly

Stage	Run time
Preprocessing	~ 20 min
Masking repeats	~ 40 min
Clustering (8,192 BlueGene/L nodes)	~ 1 hr 15 min
Assembly (64 processor xSeries Myrinet cluster)	~9 hr
Total	~ 12 hours

Input: 3,202,268 maize genomic fragments (resulted in 1,607,364 fragments after cleaning)

Blue Matter: Lipid Bilayers

- **Lipids provide the environment for membrane proteins and enable critical functions including cell signalling and cell division.**
- **Studying lipids is crucial to understanding diseases related to these proteins, including muscular dystrophy and Alzheimer's.**
- **One third of all proteins in the human body -- and half of all drug targets -- are membrane proteins**

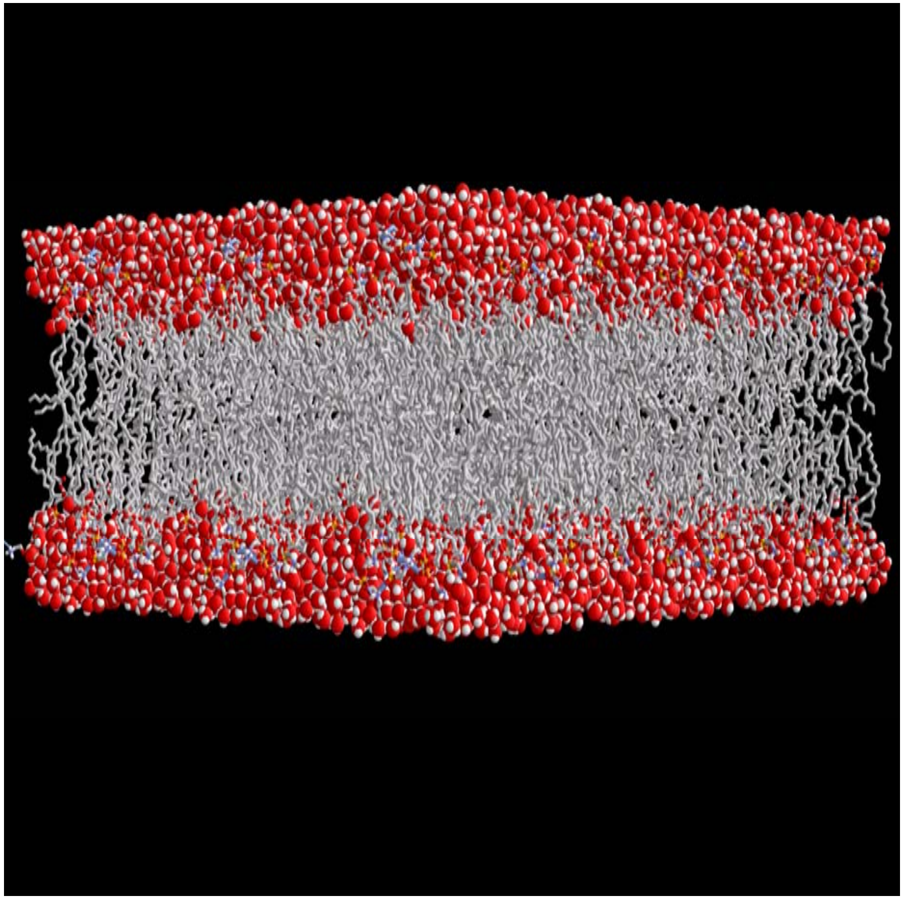

Blue Matter: RHODOPSIN GPCR

Diseases associated with malfunction of GPCRs are:

- **Congestive Heart Failure**
- **Hypertension & Stroke**
- **Cancer**
- **Ulcers**
- **Allergies**
- **Asthma**
- **Anxiety**
- **Psychosis**
- **Migraines**
- **Parkinson’s Disease**

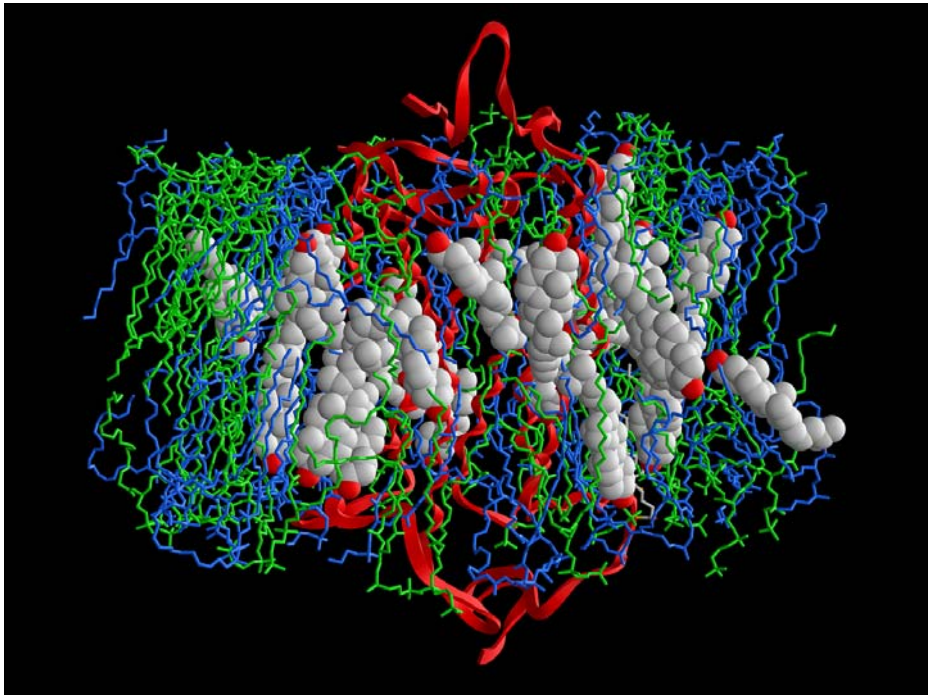

Membrane Proteins
Cell Signaling, Ion/Nutrient Transport, Targets of Many Drugs

Blue Matter on BG/L vs. NAMD on PSC Lemieux

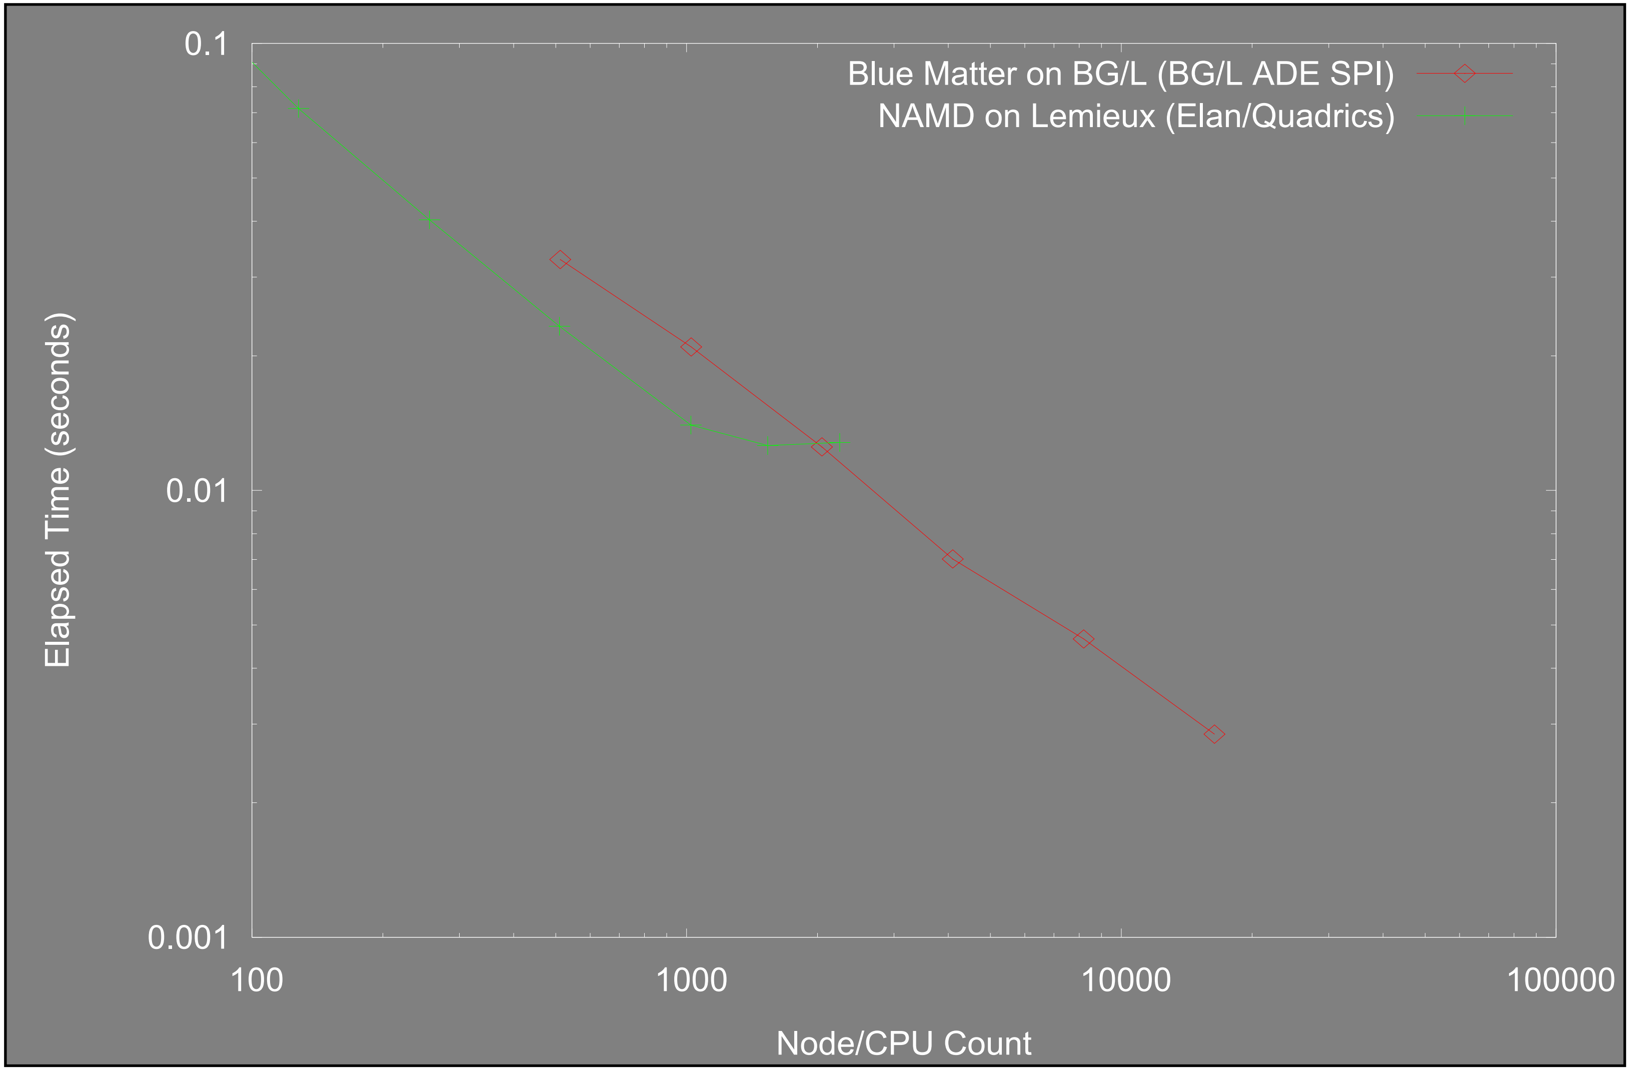

3D-Fast Fourier Transform

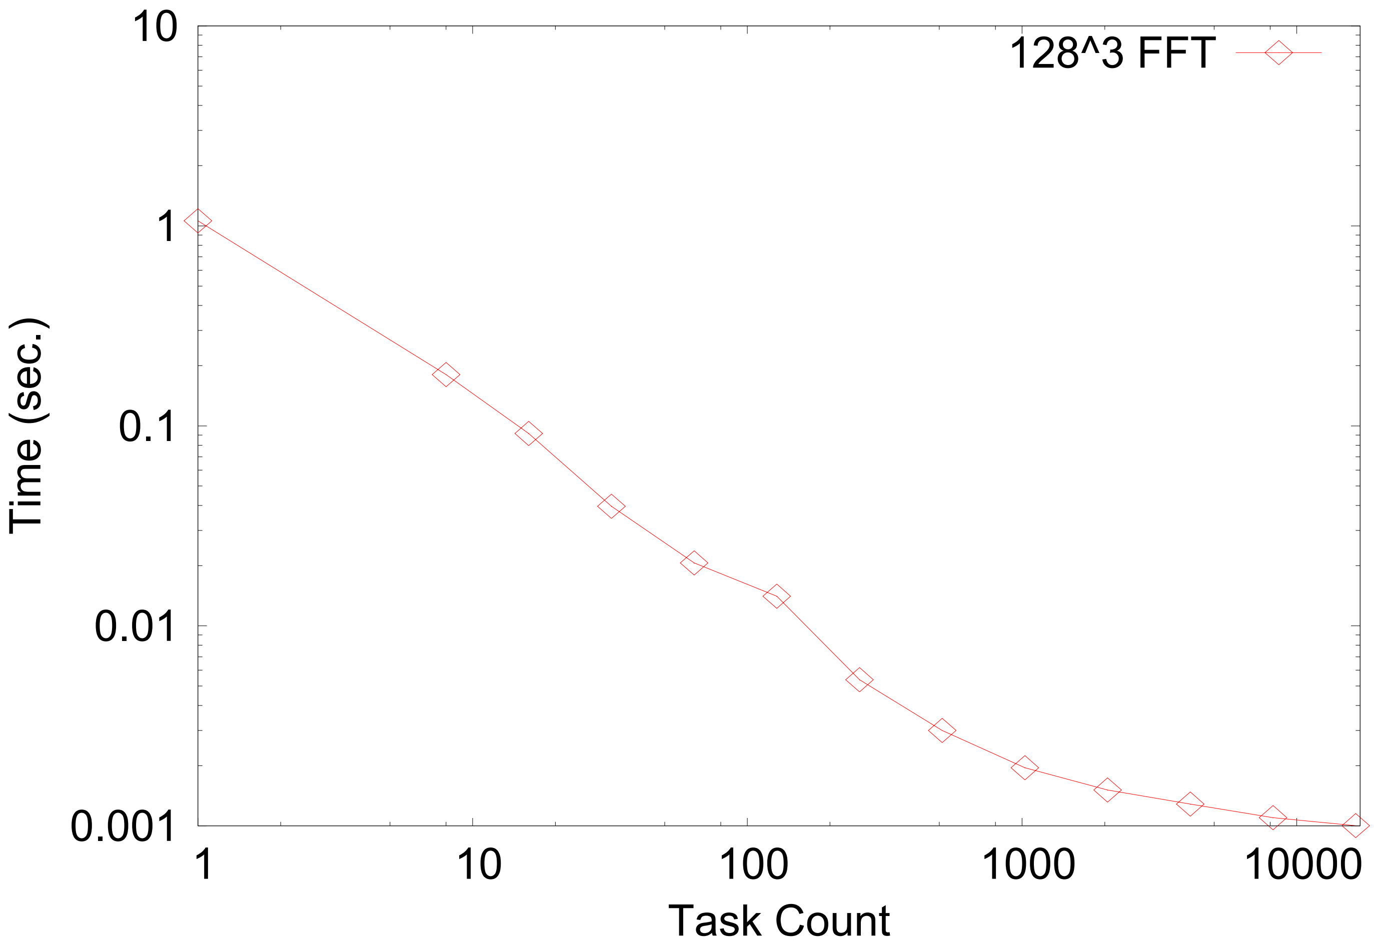

Applications Performance on Blue Gene

- **BG is first HPC system to break barrier of 100+ TeraFlop/s sustained performance on real applications (Molecular Dynamics)**
  - **ddcMD – 101.5 TeraFlop/s (*7 hrs of Uranium atoms on 64 racks*)**
  - **CPMD – 110.4 TeraFlop/s**
- **Several other applications have achieved two orders of magnitude or more higher performance than previously possible – successful scaling achieved from 1K to 100K processors**
- **Gordon Bell Prize competition at SC 2005**
  - **4 of 6 finalists based on Blue Gene**
  - **LLNL/IBM team won for "100+ TFlop Solidification Simulations on Blue Gene/L"**
  - **AIST also captured Best Technical Paper**

ddcMD

Classical MD

2005 Gordon Bell Prize Winner

- **Scalable, general purpose code for performing classical molecular dynamics (MD) simulations using highly accurate MGPT potentials**
- **MGPT semi-empirical potentials, based on a rigorous expansion of many body terms in the total energy, are needed in to quantitatively investigate dynamic behavior of transitions metals and actinides**
- **Visualization of important scientific findings already achieved on BG/L: Molten Ta at 5000K demonstrates solidification during isothermal compression to 250 GPa**

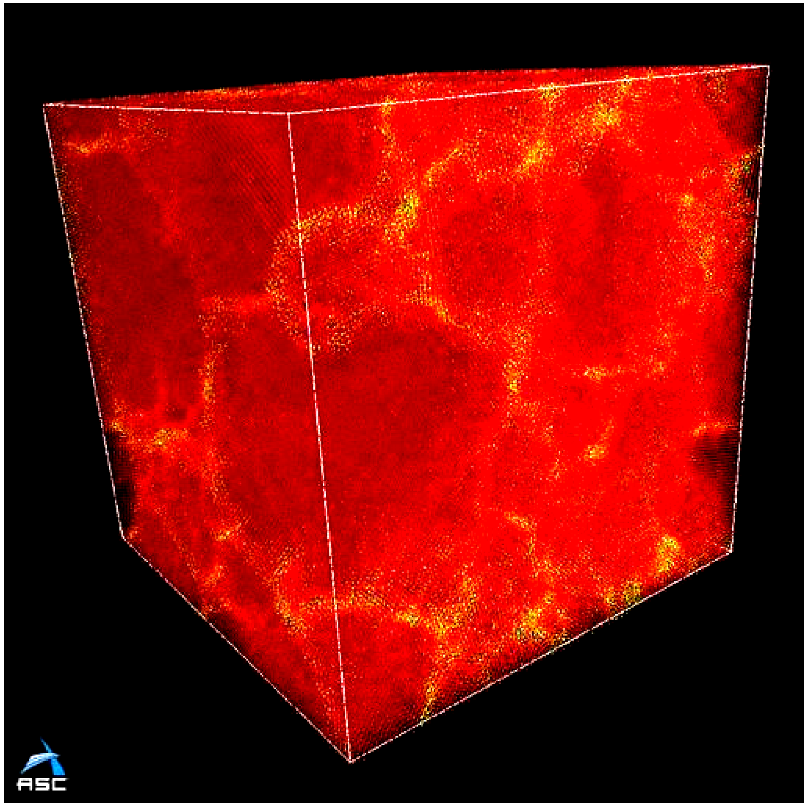

524 million atom simulations on 64K nodes are orders of magnitude larger than any previously attempted runs; *superb strong and weak scaling* expected for full machine - (“very impressive machine” says PI)

Excellent scaling of ddcMD on BG/L supports solidification understanding

- Nucleation is initiated at multiple independent sites in each sample cell
- Growth of solid grains initiates independently, but soon leads to grain boundaries which span the simulation cell
- 101.5 TF on 64 racks
- The ddcMD team is currently using 131,072 CPUs of BG/L for unprecedented
- five hundred million atom MGPT simulations

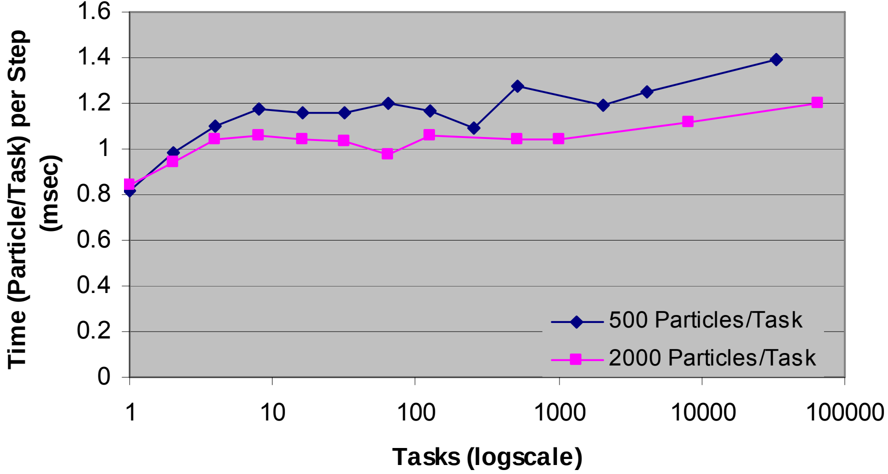

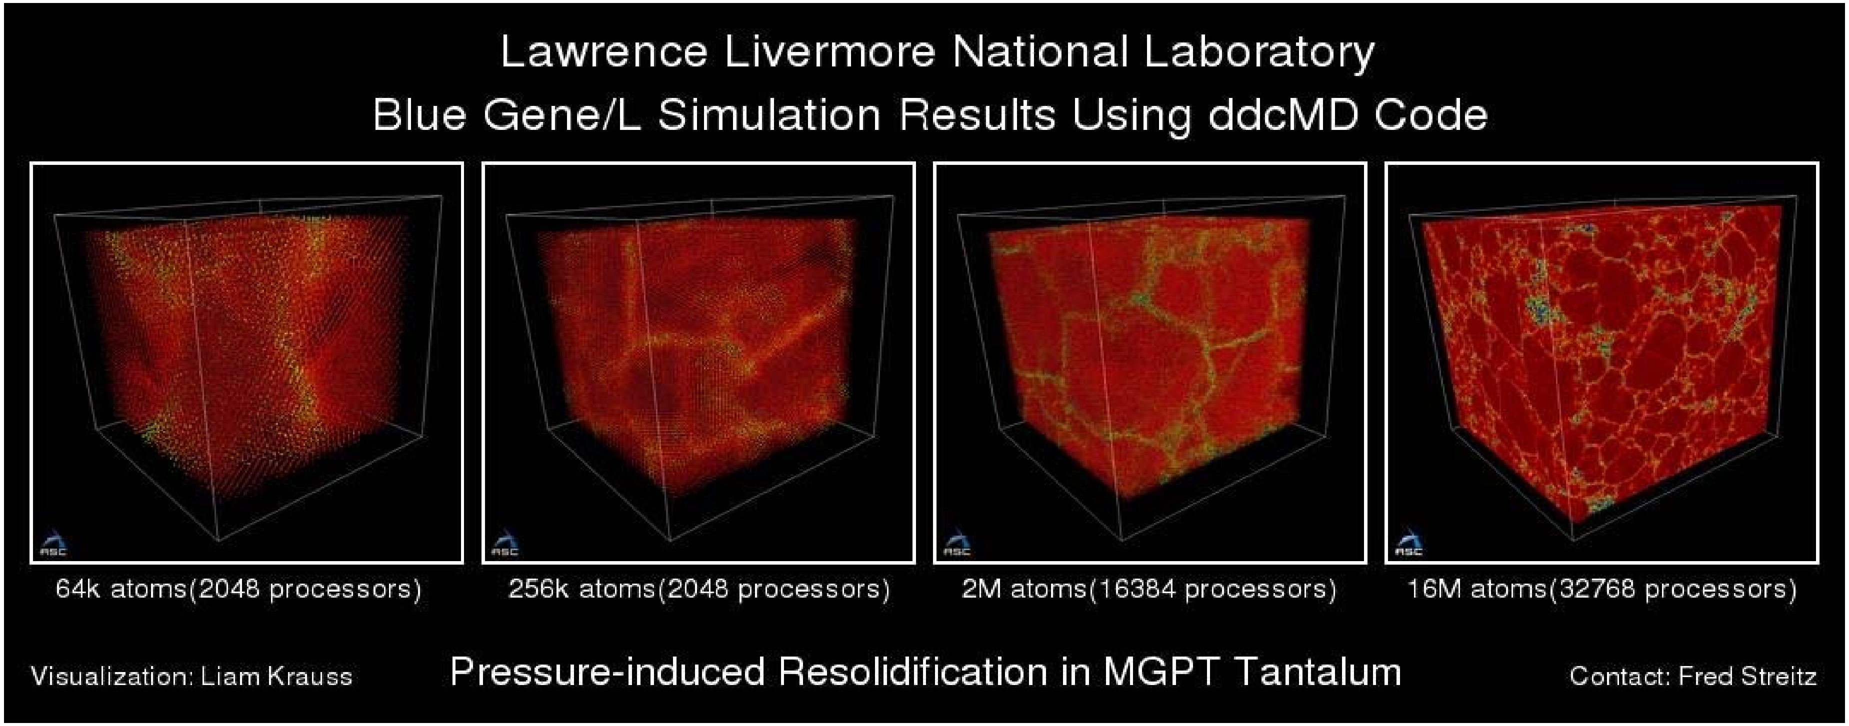

Lawrence Livermore National Laboratory

Blue Gene/L Simulation Results Using ddcMD Code

64k atoms(2048 processors) | 256k atoms(2048 processors) | 2M atoms(16384 processors) | 16M atoms(32768 processors)

Visualization: Liam Krauss

Pressure-induced Resolidification in MGPT Tantalum

Contact: Fred Streitz

Performance of ddcMD on Blue Gene

Weak scaling: MGPT Uranium and Tantalum

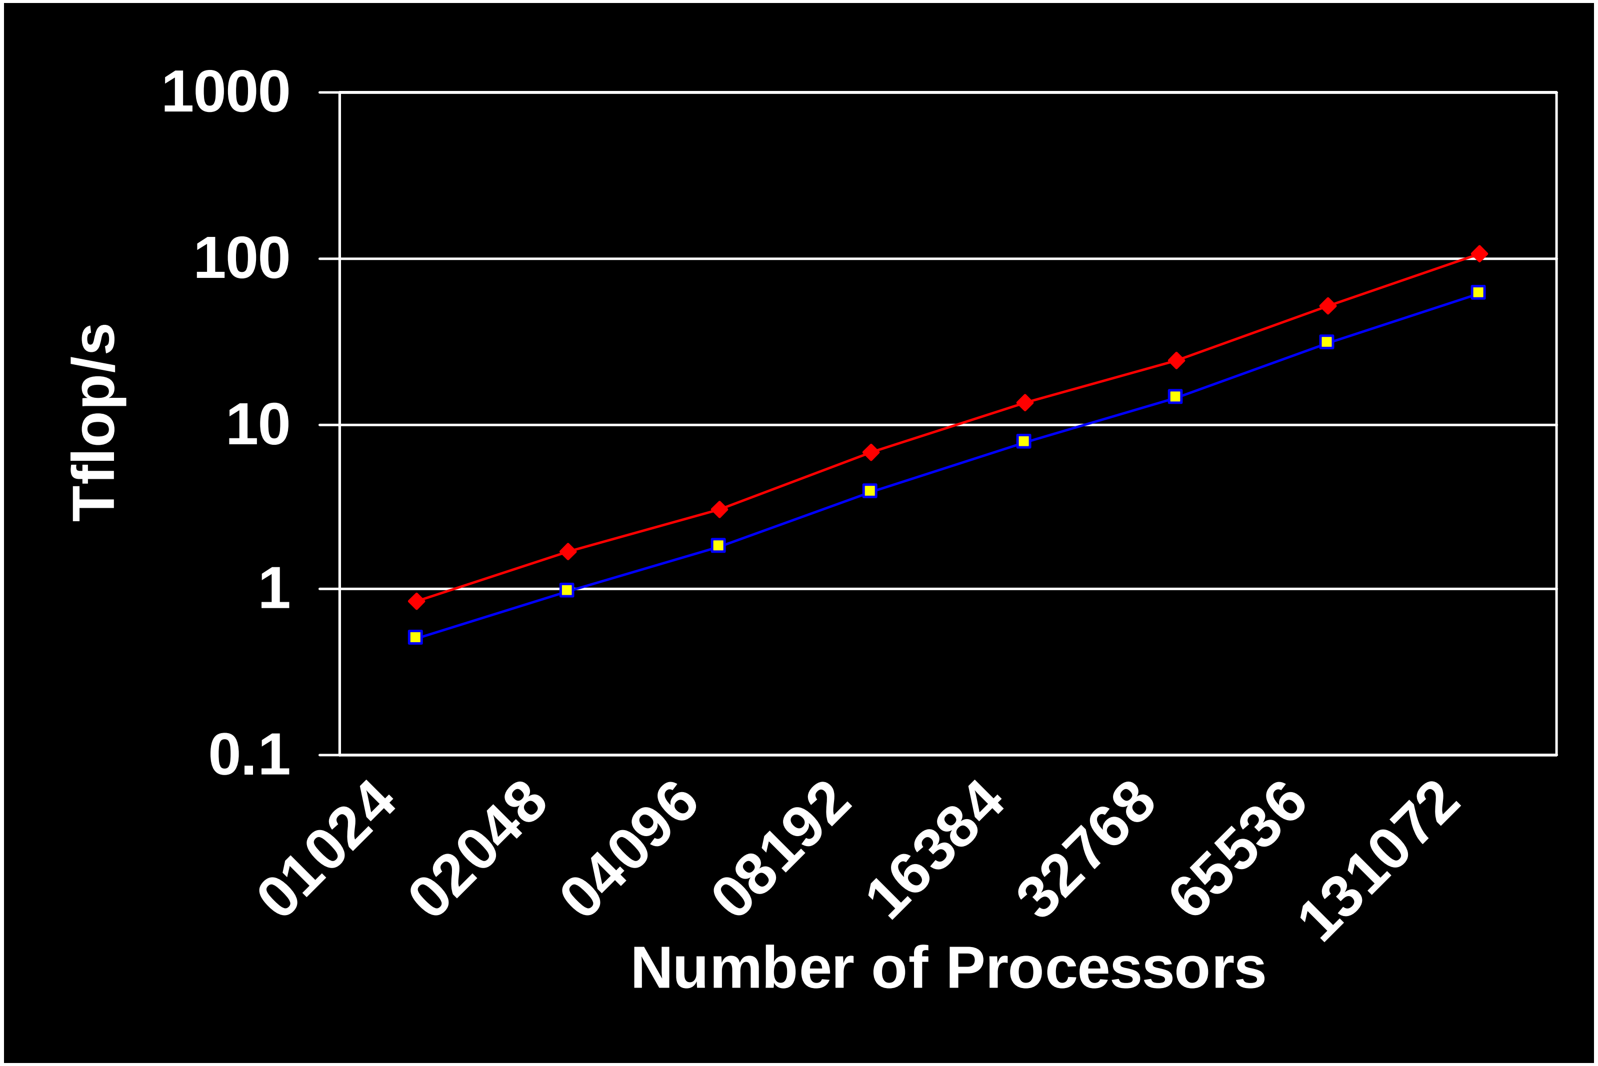

ddcMD Simulation Results

(a) 64K atoms, (b) 256 K atoms, (c) 2,048,000 atoms, (d) 16,384,000 atoms

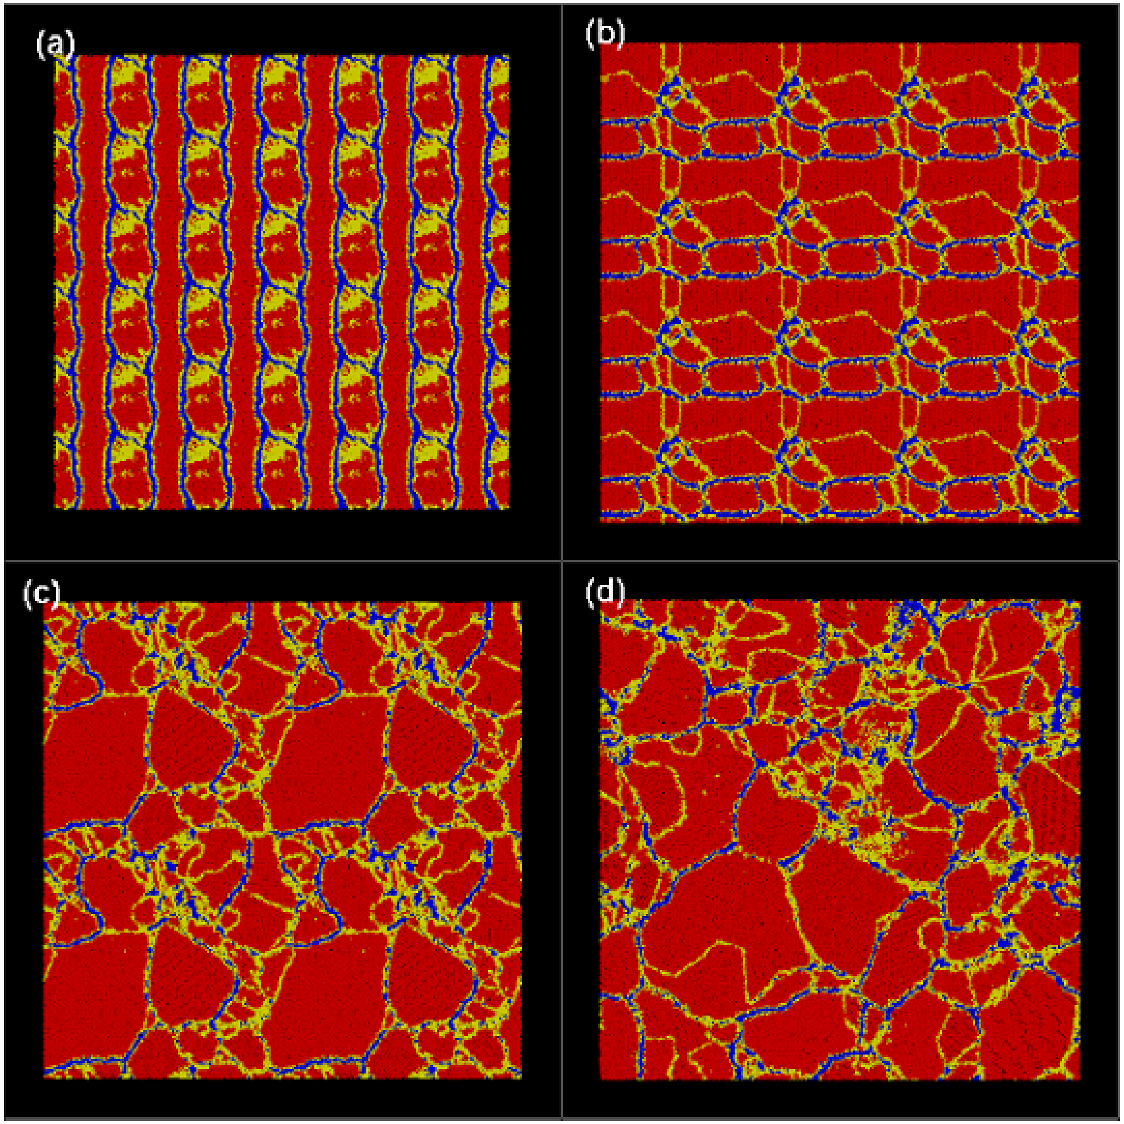

Scaling ddcMD up to 131,072 CPUs

… but allows unprecedented scaling of size or time

- **Weak scaling is virtually flat across the entire machine - enables simulation of tens of billions of atoms (roughly a cubic micron of material)**
- **Strong scaling shows speedup down to 8 atoms/CPU - enables simulations involving millions of steps (typically ns of simulated time)**

Weak Scaling (Ta and U)

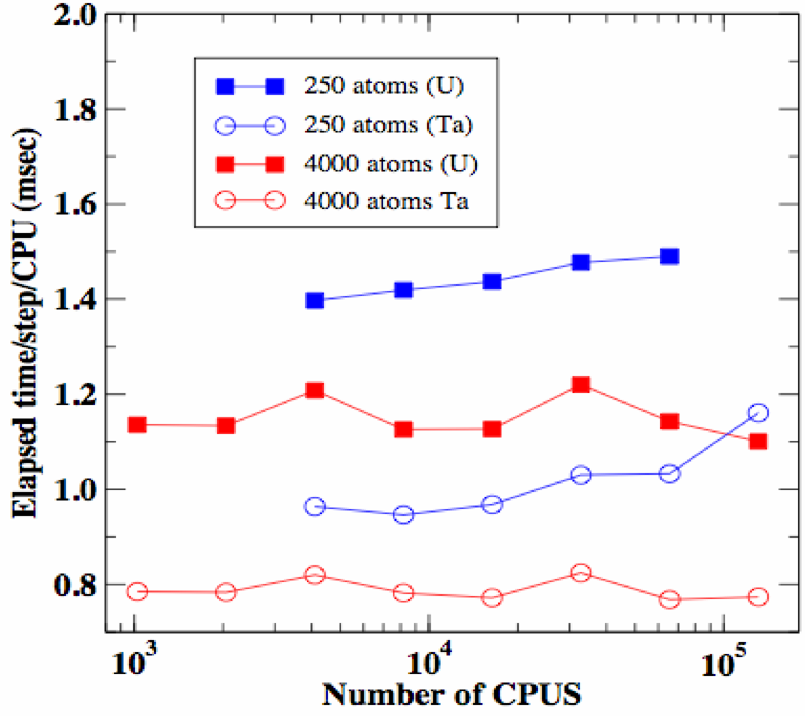

Strong Scaling (Ta and U)

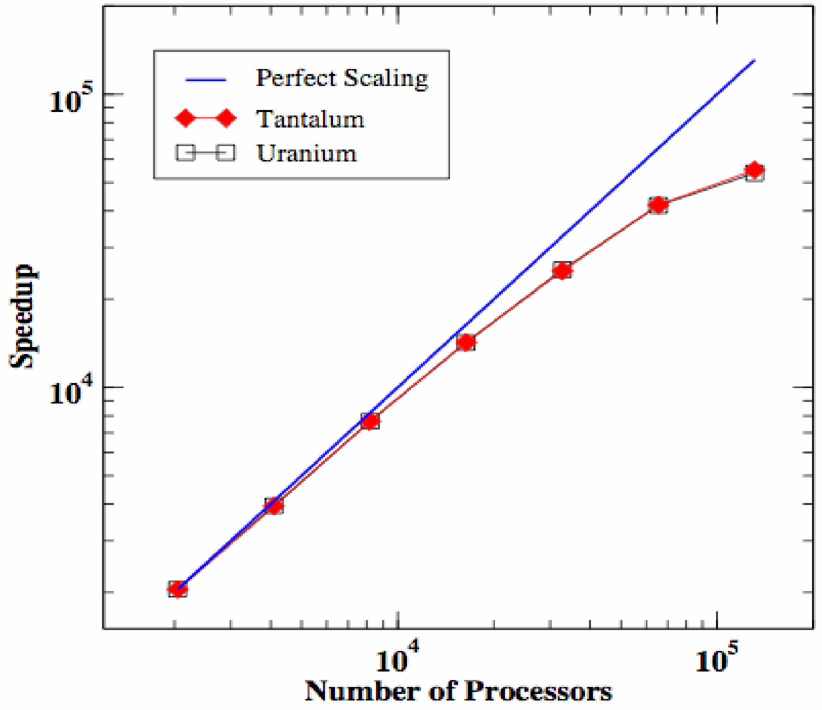

CPMD

Alessandro Curioni, Salomon Billeter, Wanda Andreoni

CPMD Performance on BG/L

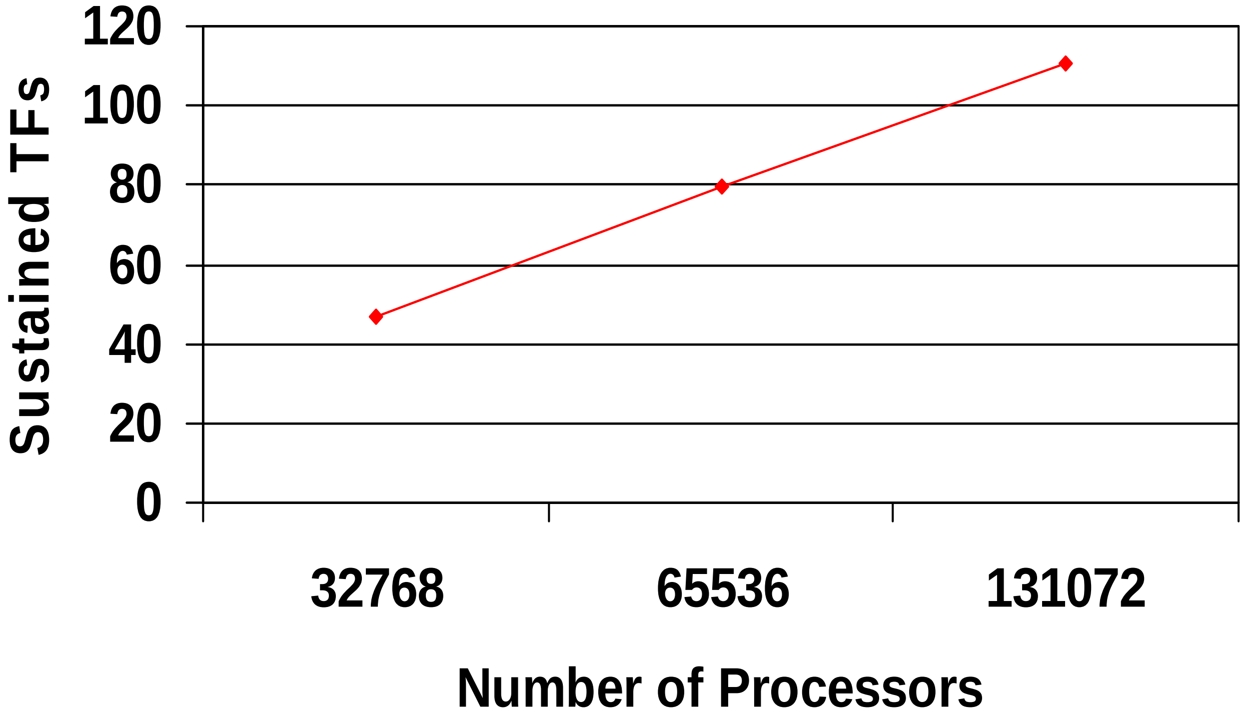

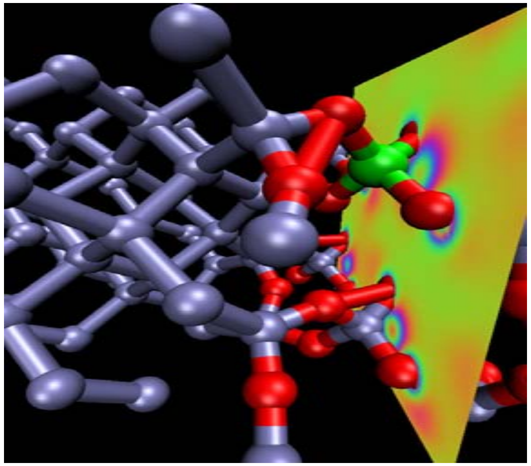

Developed at IBM Zurich from Car Parinello code
Uses Plane Wave Basis functions, FFT, MPI_Collectives
Ongoing project : IBM/LLNL Pd:H (~900 atoms) Hydrogen Storage
Achieved 110.4 Teraflop/s sustained on 64 racks BG/L (excellent strong scaling)

EPFL

IBM's Blue Gene supercomputer allows a quantum leap in the level of detail at which the brain can be modeled. The time has come to begin assimilating the wealth of data that has accumulated over the past century and begin building biologically accurate models of the brain to aid our understanding of brain function and dysfunction.

Henry Markram, Laboratory of Neural Microcircuitry, Brain Mind Institute
Ecole Polytechnique Fédérale de Lausanne
Switzerland

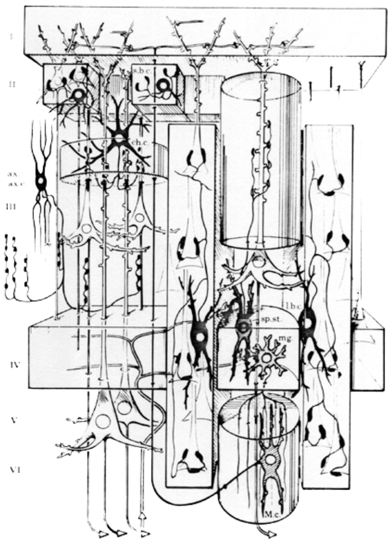

Blue Gene award-winning performance

- **3 in Top10 (#1 and #2)**
- **7 in Top50 (11-50)**
- **9 in Top100 (51-100)**
- **19 in Top500**

www.top500.org

Moore's Law Prevails

Performance improvement causes about 200 systems (40%) to fall out of the ranking every six months.

HPC Challenge Benchmarks	IBM Blue Gene 65536 nodes 64 racks	Cray XT3 (Opteron) 5200 nodes 56 racks
HPL (TFLOP/s) – Linpack TPP benchmark which measures floating point rate of execution for solving linear system of equations	**259.21**	**20.53**
RANDOMACCESS (GUP/s) – measures rate of integer random updates of memory	**35.46**	**0.69 (7.69 on Cray X1E)**
FETE (GFLOP/s) – measures floating point rate of execution of double precision complex one-dimensional Discrete Fourier Transform (DFT)	**2311.09**	**905.57**
STREAM (GB/s) – simple synthetic that measures sustainable memory bandwidth and corresponding computation rate for simple vector kernel	**160,064**	**26,021**

The Real Question

- What can you do with 130K processors? (8K, 16K, 32K)
  - Really BIG problems – Maybe
  - Same problems but much finer resolution, refinements, larger searches in shorter time – Maybe
  - Explore parameters – large parameter space – Maybe
- BUT
  - Perhaps need to rethink the problem
  - Most parallel programs are Single Program Multiple Data
- **What if**
  - Multiple Programs Multiple Data - - Systems of Complex Systems interacting?
  - Handle multi-scale, multi-physics - - Biology is multi-scale?

BlueGene/L will allow overlapping evaluation of models for the first time

[redacted]

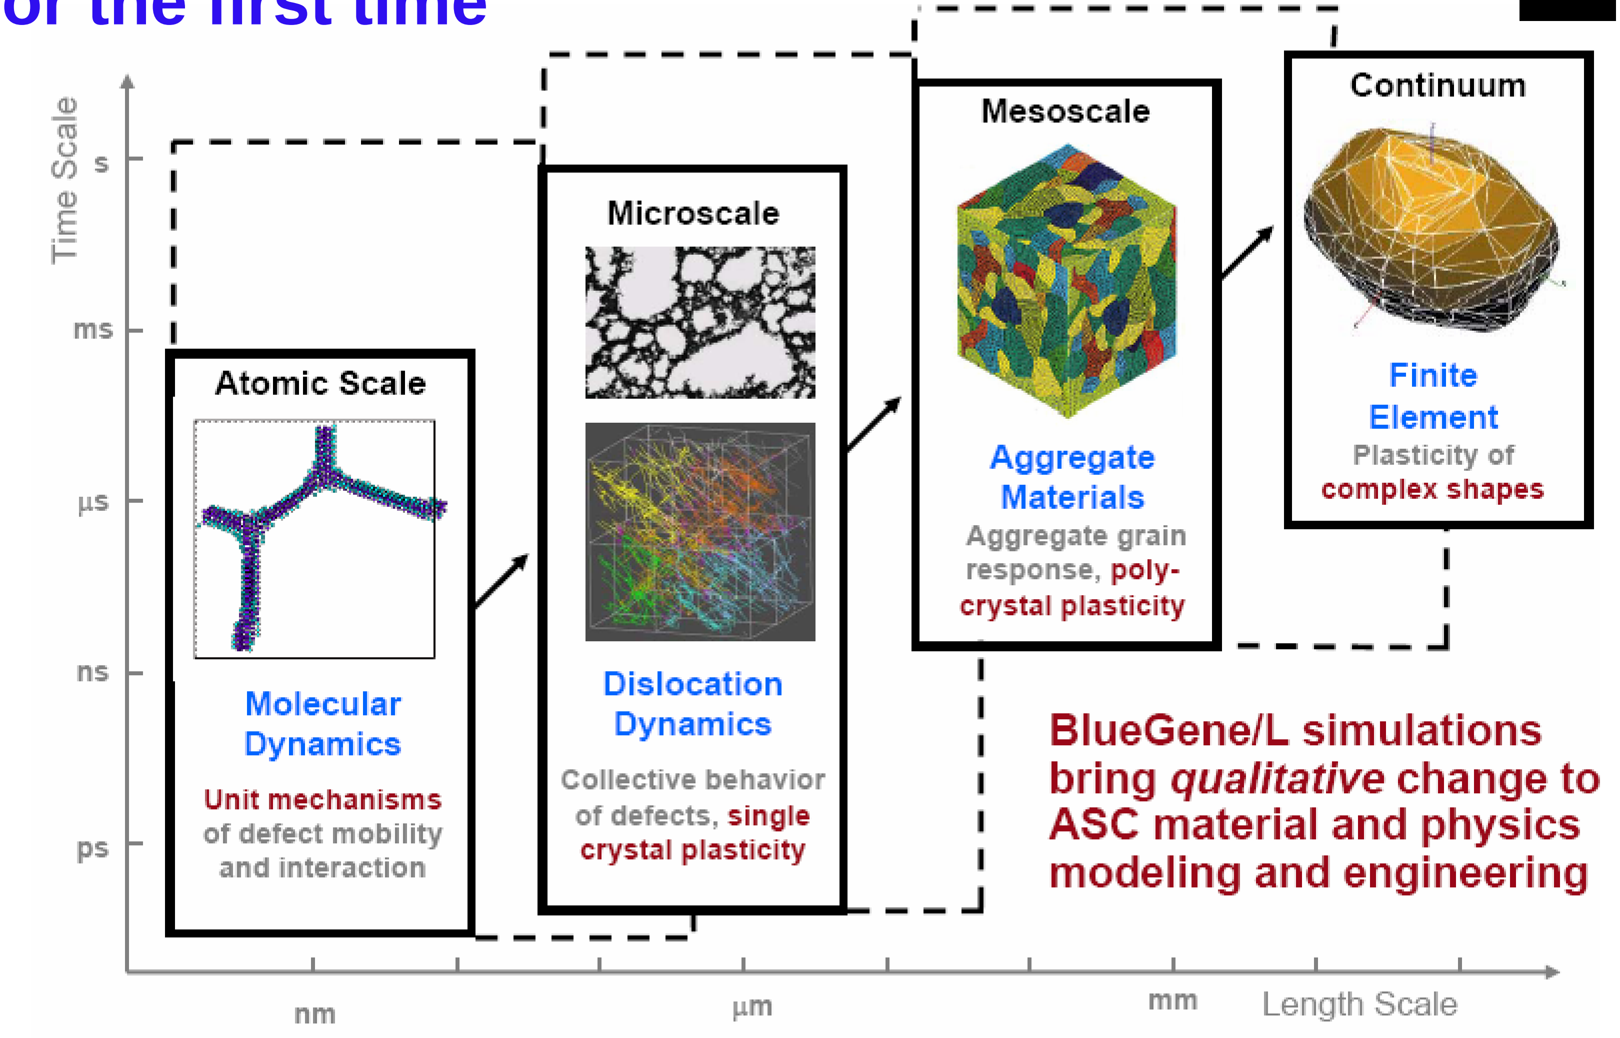

BlueGene/L simulations bring *qualitative* change to ASC material and physics modeling and engineering

Courtesy of Steve Louis @ LLNL

Some Grand Challenges for Healthcare and Life Sciences

- **Whole cell modeling - e.coli** Genetically engineer e-coli strains to increase drug production by fermentation

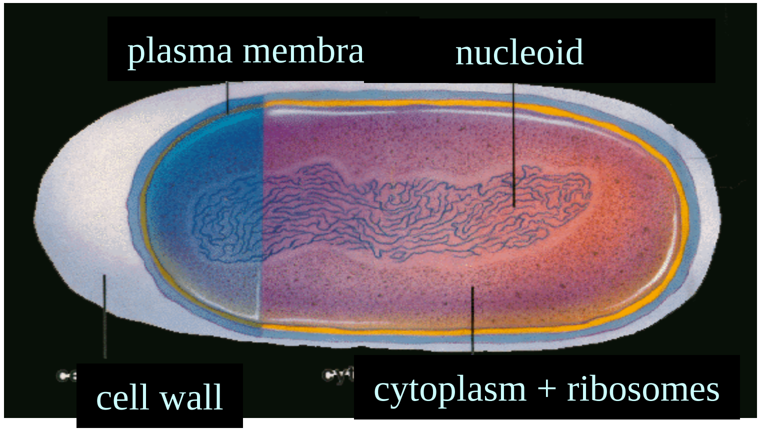

- ▪ E. coli is the most popular target
  - It's simple: "only" ~4000 genes, no nucleus, unitary genes, no organelles
  - It's well studied
- ▪ How might Blue Gene impact
  - Atomistic level
  - Chemical kinetics
  - Continuous models

- **Drug delivery modeling**

Develop manufacturing processes that insure the right dose of a drug is effectively delivered

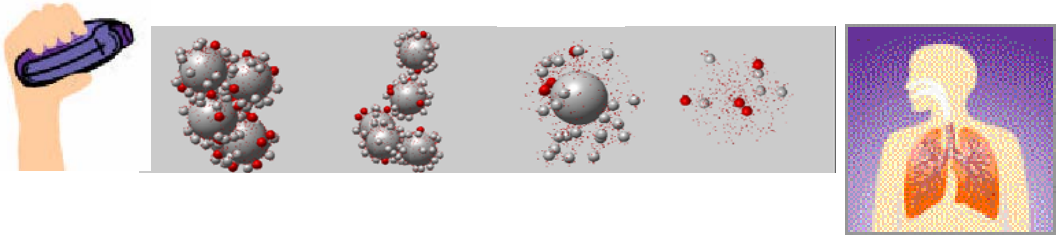

- ▪ Inhaler delivery of drugs
  - Multi-scale – atomistic, chemical, fluid flow
  - Manufacturing complex model
- ▪ How might Blue Gene impact
  - Atomistic level
  - Chemical kinetics
  - Continuous models

Increasing Importance of Engineering, Mathematical and Computational Sciences in Human Disease Research – Computational Medicine

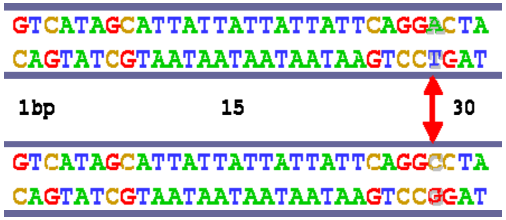

Genetic Variability

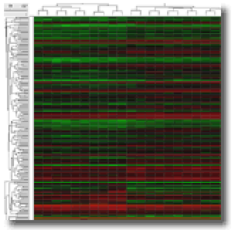

Gene Expression Profiling

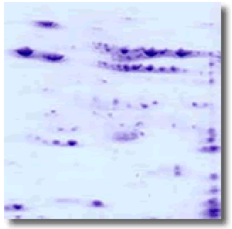

Protein Expression Profiling

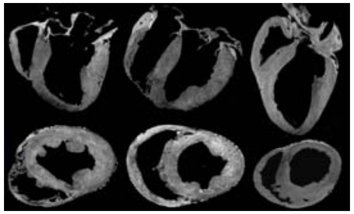

Multi-Modal Imaging

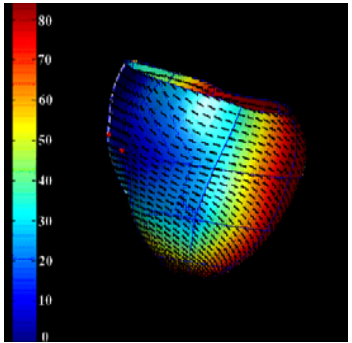

Data Analysis And Modeling

Goals – How Best To:

- analyze these data sets to gain novel insights regarding disease mechanisms and to perform risk prediction targeted to the individual (*statistical inference*, *pattern discovery/classification*, *computational anatomy*)
- synthesize computational models of biological systems and disease processes that provide insights into disease mechanisms and novel therapies (*dynamical systems theory*, *probability theory*, *stochastic processes*)
- Distribute multi-scale data, data analysis methods and computational models to basic and clinical researchers through computational grids ("bio-grids") – Blue Gene potential compute engine

Courtesy: Rai Winslow, Johns Hopkins University

Multi-scale in Physiome Project

Hunter's Group – converting CMISS (**C**ontinuum **M**echanics, **I**mage analysis, **S**ignal processing and **S**ystem Identification) to potentially use Blue Gene

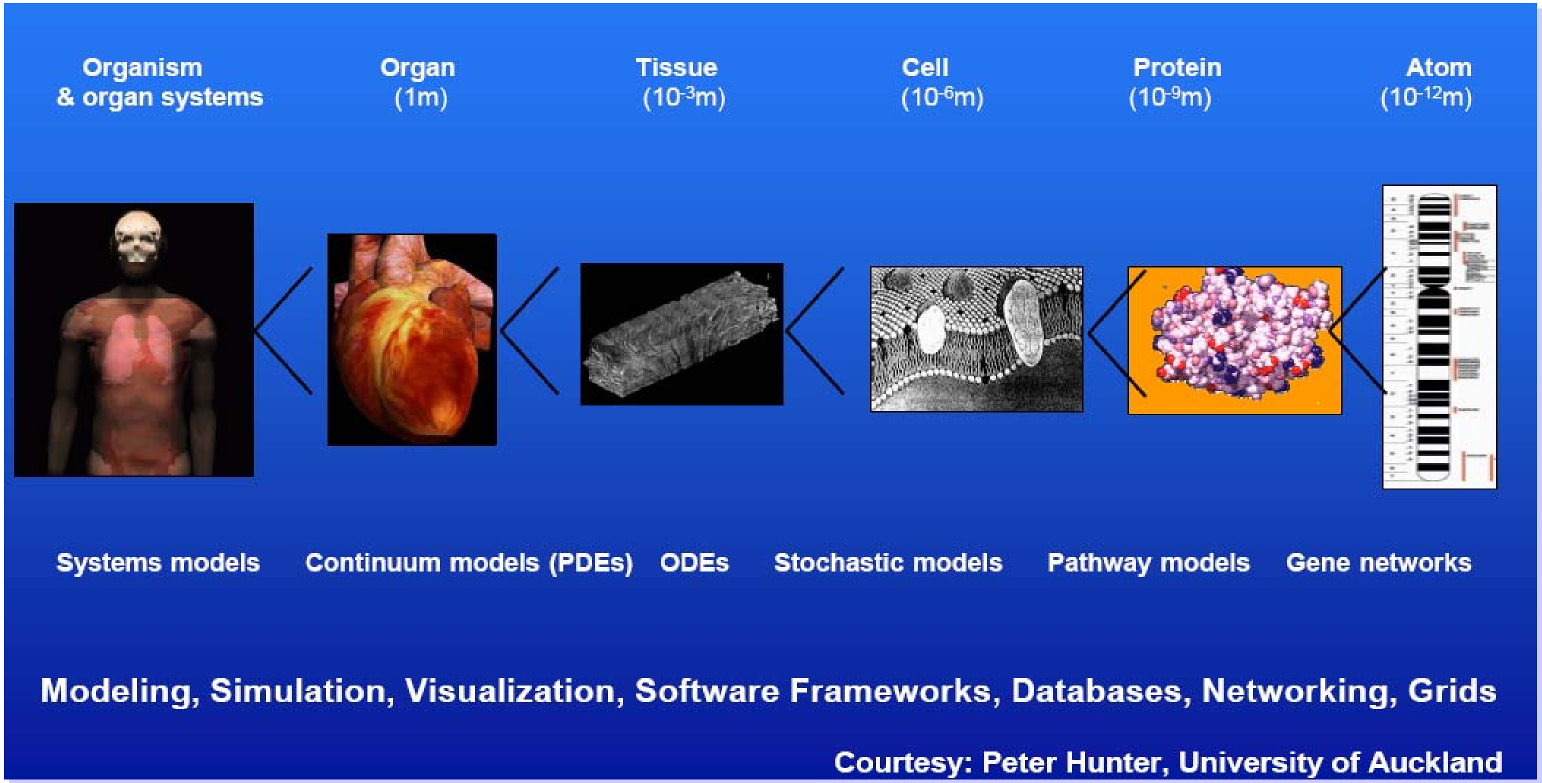

Courtesy: Peter Hunter, University of Auckland

- Blue Gene Impact
  - Atomistic to System Model scales – tightly coupled

Biomedical and Molecular Imaging – Exploit Parallelism – Use Blue Gene JHU

Dean Wong

High Resolution Research Tomograph (HRRT)

- **Resolution ~ 2.5 mm**
- **Sensitivity ~ 6%**
- **Number of detectors 119,800**
- **Number of Lines of response (LOR) 4.5 Billion !!**

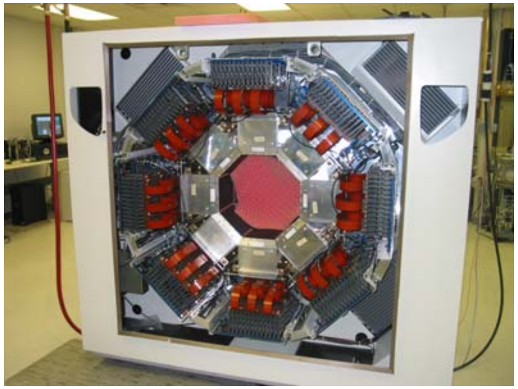

Dynamic (4D) PET Imaging

Quantity measured in PET: *in vivo* regional concentration of the radiotracer

- Use multiple time frames to “measure” the physiologic or metabolic process
- Can extract how various compartments interact

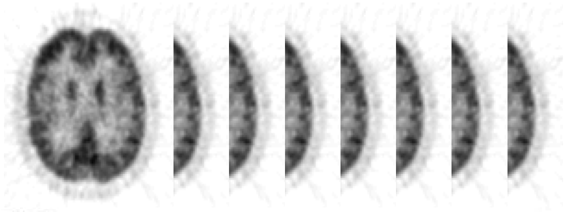

Radiology - Storage/Computational Issues

- ~30GB of raw/list-mode data per study
- Each study divided into ~30 frames and reconstructed
- Currently, computation takes:
  - 15 hours (span3), 7 hours (span9:lower-res) per study!
  - 8 nodes/frame, 4 frames processed at a time
  - Each reconstructed image is 50MB (i.e. 1.5GB /study)

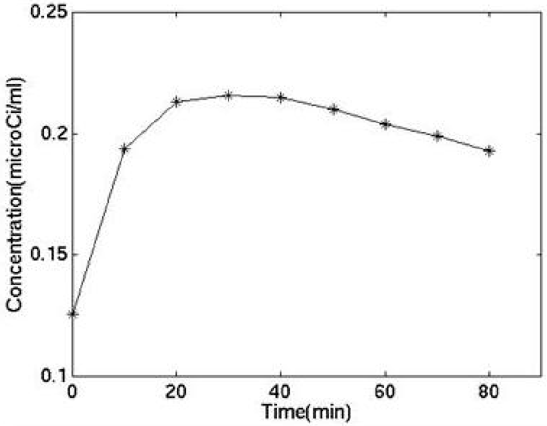

Goal: 20 studies/week: clearly not achieved in span3

BG/L Consortium – Need to Identify and Enhance Bio Applications

LABORATORIES:

- Ames National Lab/Iowa State University
- Argonne National Laboratory
- Brookhaven National Laboratory
- Fermi National Laboratory
- Jefferson Laboratory
- Lawrence Berkeley National Laboratory
- Lawrence Livermore National Lab
- Oak Ridge National Laboratory
- Pacific Northwest National Laboratory
- Princeton Plasma Physics Laboratory

UNIVERSITIES:

- Boston University
- California Institute of Technology
- Columbia University
- Cornell University
- DePaul University
- Harvard University
- Illinois Institute of Technology
- Indiana University
- Iowa State University

UNIVERSITIES (cont'd):

- Louisiana State University
- Massachusetts Institute of Technology
- National Cntr for Atmospheric Research
- New York University - Courant Institute
- Northern Illinois University
- Northwestern University
- Ohio State University
- Pennsylvania State University
- Pittsburgh Super Computing Center
- Princeton University
- Purdue University
- Rutgers University
- Stony Brook University
- Texas A&M University
- University of California - Irvine
- University of California - San Francisco
- University of California - SDSC
- University of Chicago
- University of Colorado - JILA
- University of Delaware

UNIVERSITIES (cont'd):

- University of Illinois Urbana-Champaign
- University of Minnesota
- University of North Carolina
- University of S California - ISI
- University of Texas at Austin - TACC
- University of Utah
- University of Wisconsin

INDUSTRY:

- Engineered Intelligence Corporation
- IBM

INTERNATIONAL:

- ASTRON/LOFAR
- CERN
- EPFL
- NIC Juelich
- AIST
- National University of Ireland/ICHEC
- Trinity College, Trinity Centre for HPCC
- John von Neumann Inst. for Computing
- NIWS Co., Ltd.
- University of Edinburgh, EPCC

Remarks

- Is Blue Gene a systems for computational science & engineering?
  - Starting to see effective use of lots of processors in this domain
- Need to re-think how to tackle problems
- Think of problems might tackle that would not have dreamed of doing without Blue Gene
- Systems of complex systems will need multi-disciplinary teams, open access and enthusiasm in Universities, Government Labs and Industry (Hospitals, Pharmas, Software Vendors)

BECAUSE *SPEED* MATTERS.....

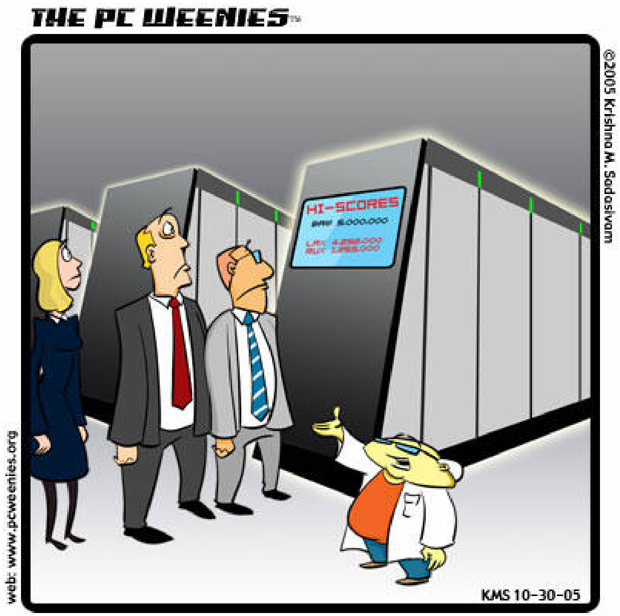

"OUR SUPERCOMPUTER IS CAPABLE OF PERFORMING TRILLIONS OF COMPUTATIONS PER SECOND, BUT WE PRIMARILY USE IT FOR BIT TORRENT AND GAMING."

National Geographic Web site

Familytree DNA

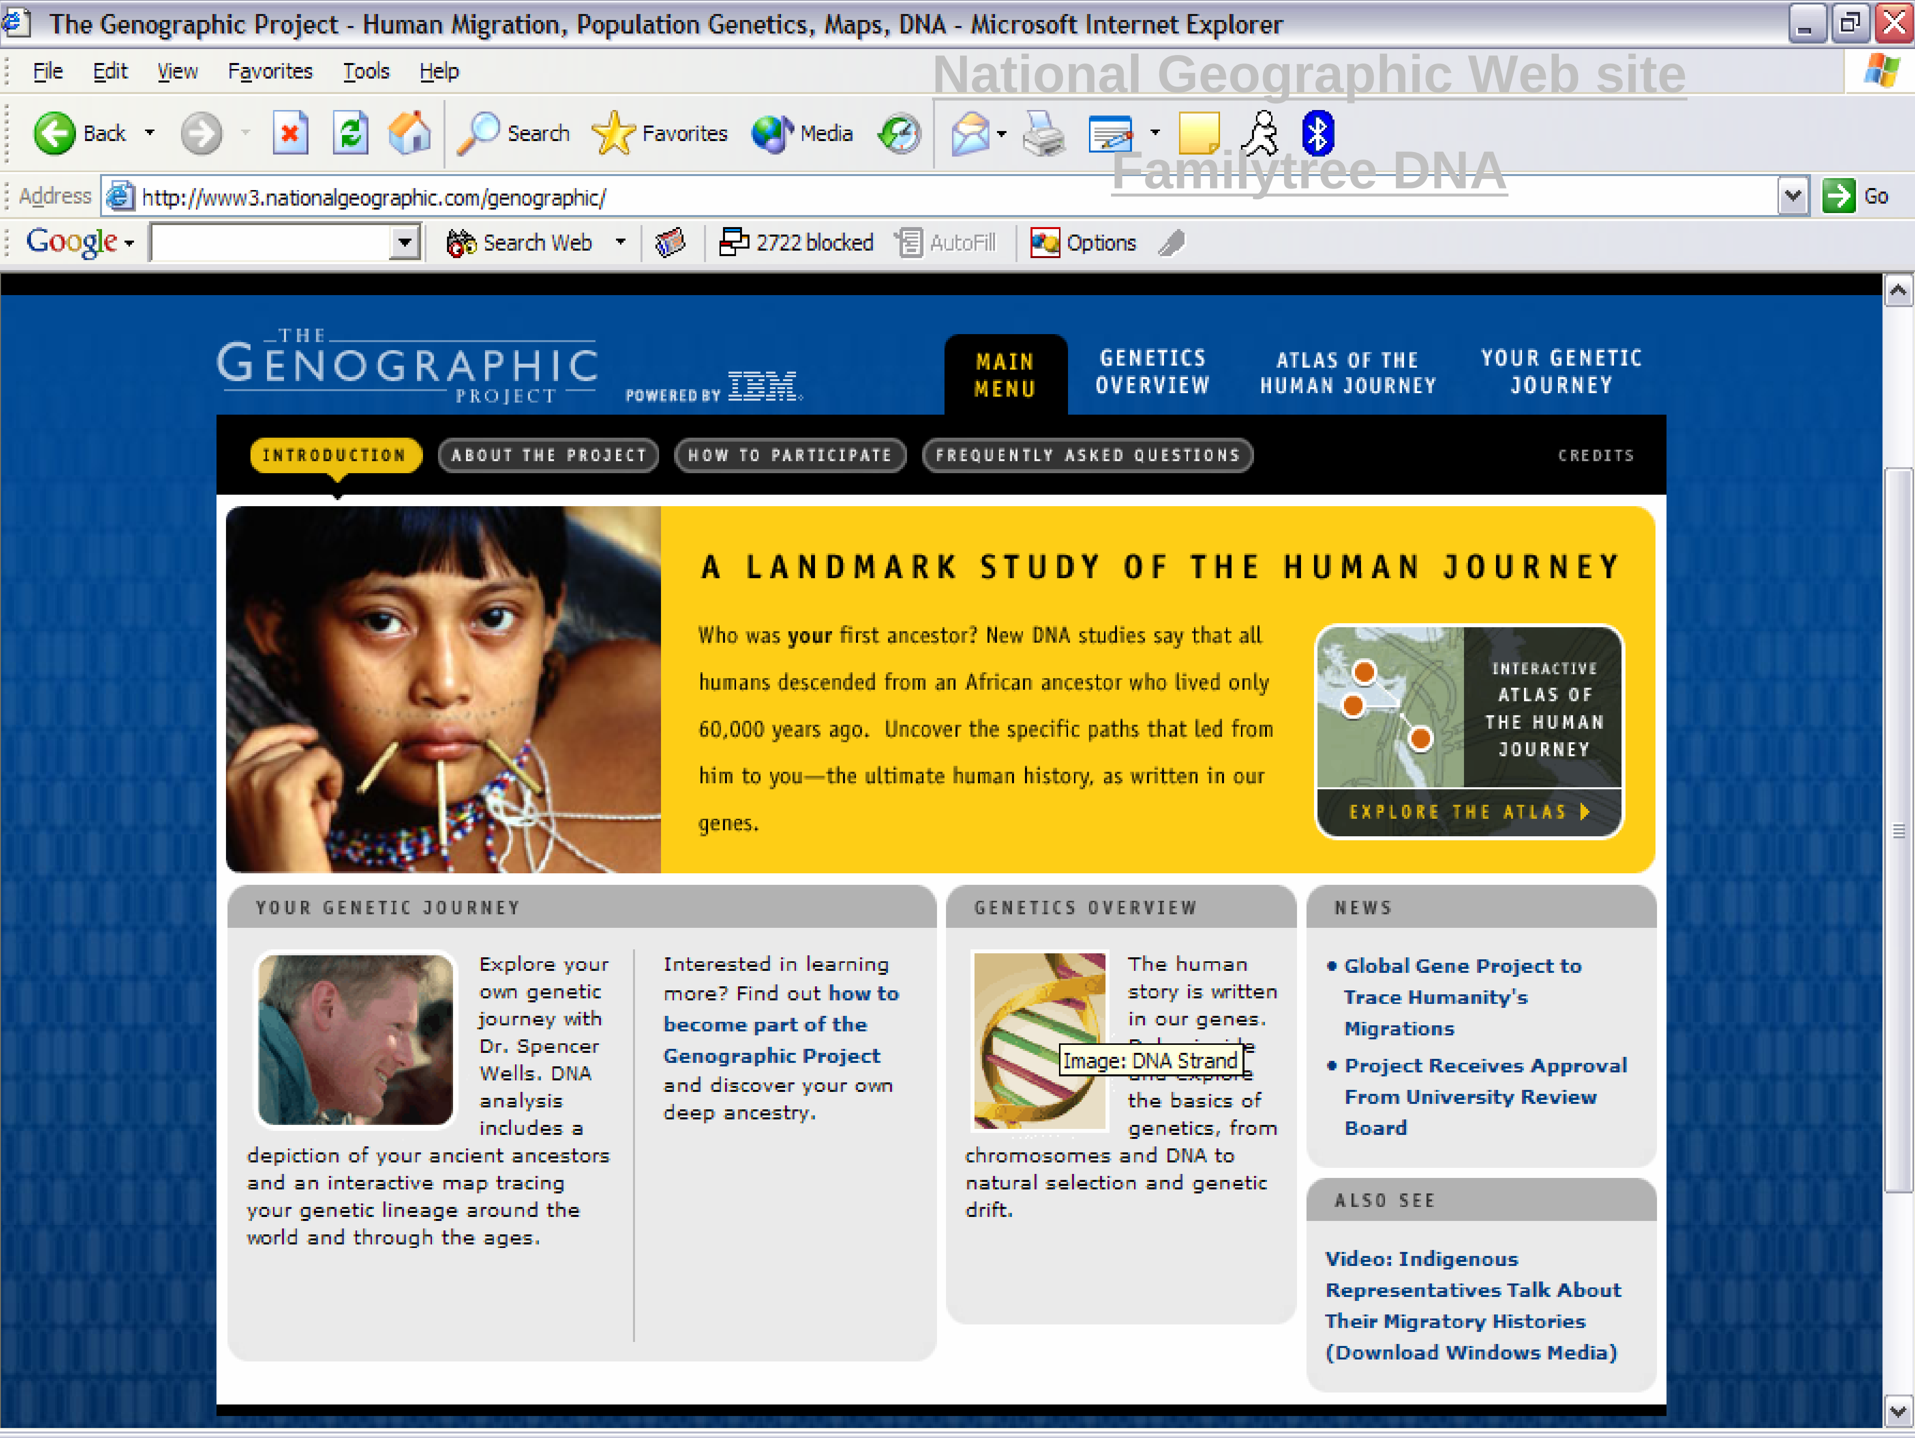

A LANDMARK STUDY OF THE HUMAN JOURNEY

Who was **your** first ancestor? New DNA studies say that all humans descended from an African ancestor who lived only 60,000 years ago. Uncover the specific paths that led from him to you—the ultimate human history, as written in our genes.

YOUR GENETIC JOURNEY

Explore your own genetic journey with Dr. Spencer Wells. DNA analysis includes a depiction of your ancient ancestors and an interactive map tracing your genetic lineage around the world and through the ages.

Interested in learning more? Find out **how to become part of the Genographic Project** and discover your own deep ancestry.

GENETICS OVERVIEW

The human story is written in our genes. [illegible] and explore the basics of genetics, from chromosomes and DNA to natural selection and genetic drift.

NEWS

- **Global Gene Project to Trace Humanity's Migrations**
- **Project Receives Approval From University Review Board**

ALSO SEE

Video: Indigenous Representatives Talk About Their Migratory Histories (Download Windows Media)

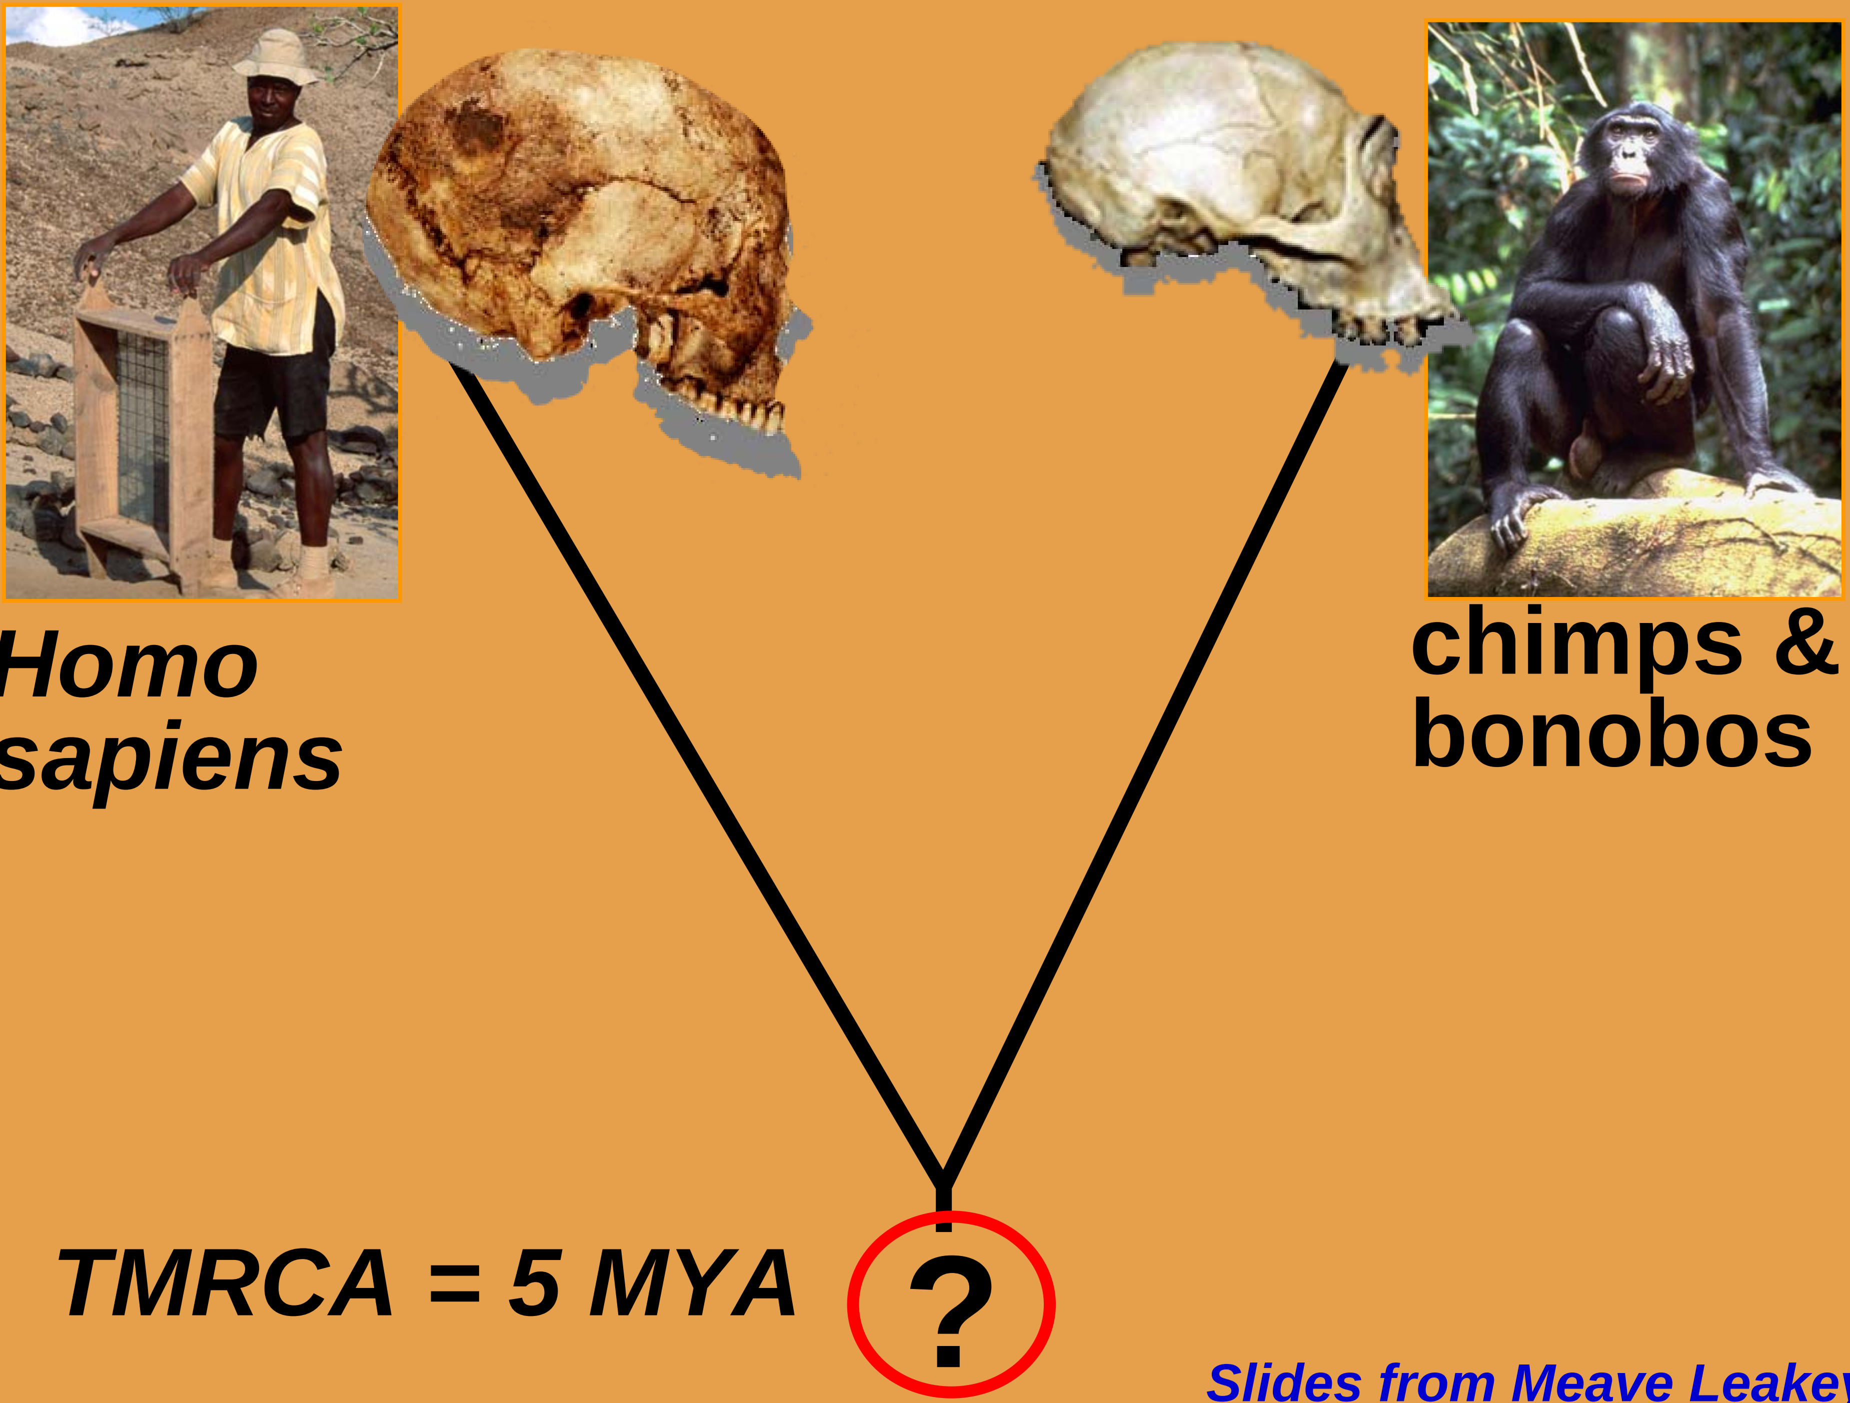

Homo sapiens

chimps & bonobos

TMRCA = 5 MYA

?

Slides from Meave Leakey

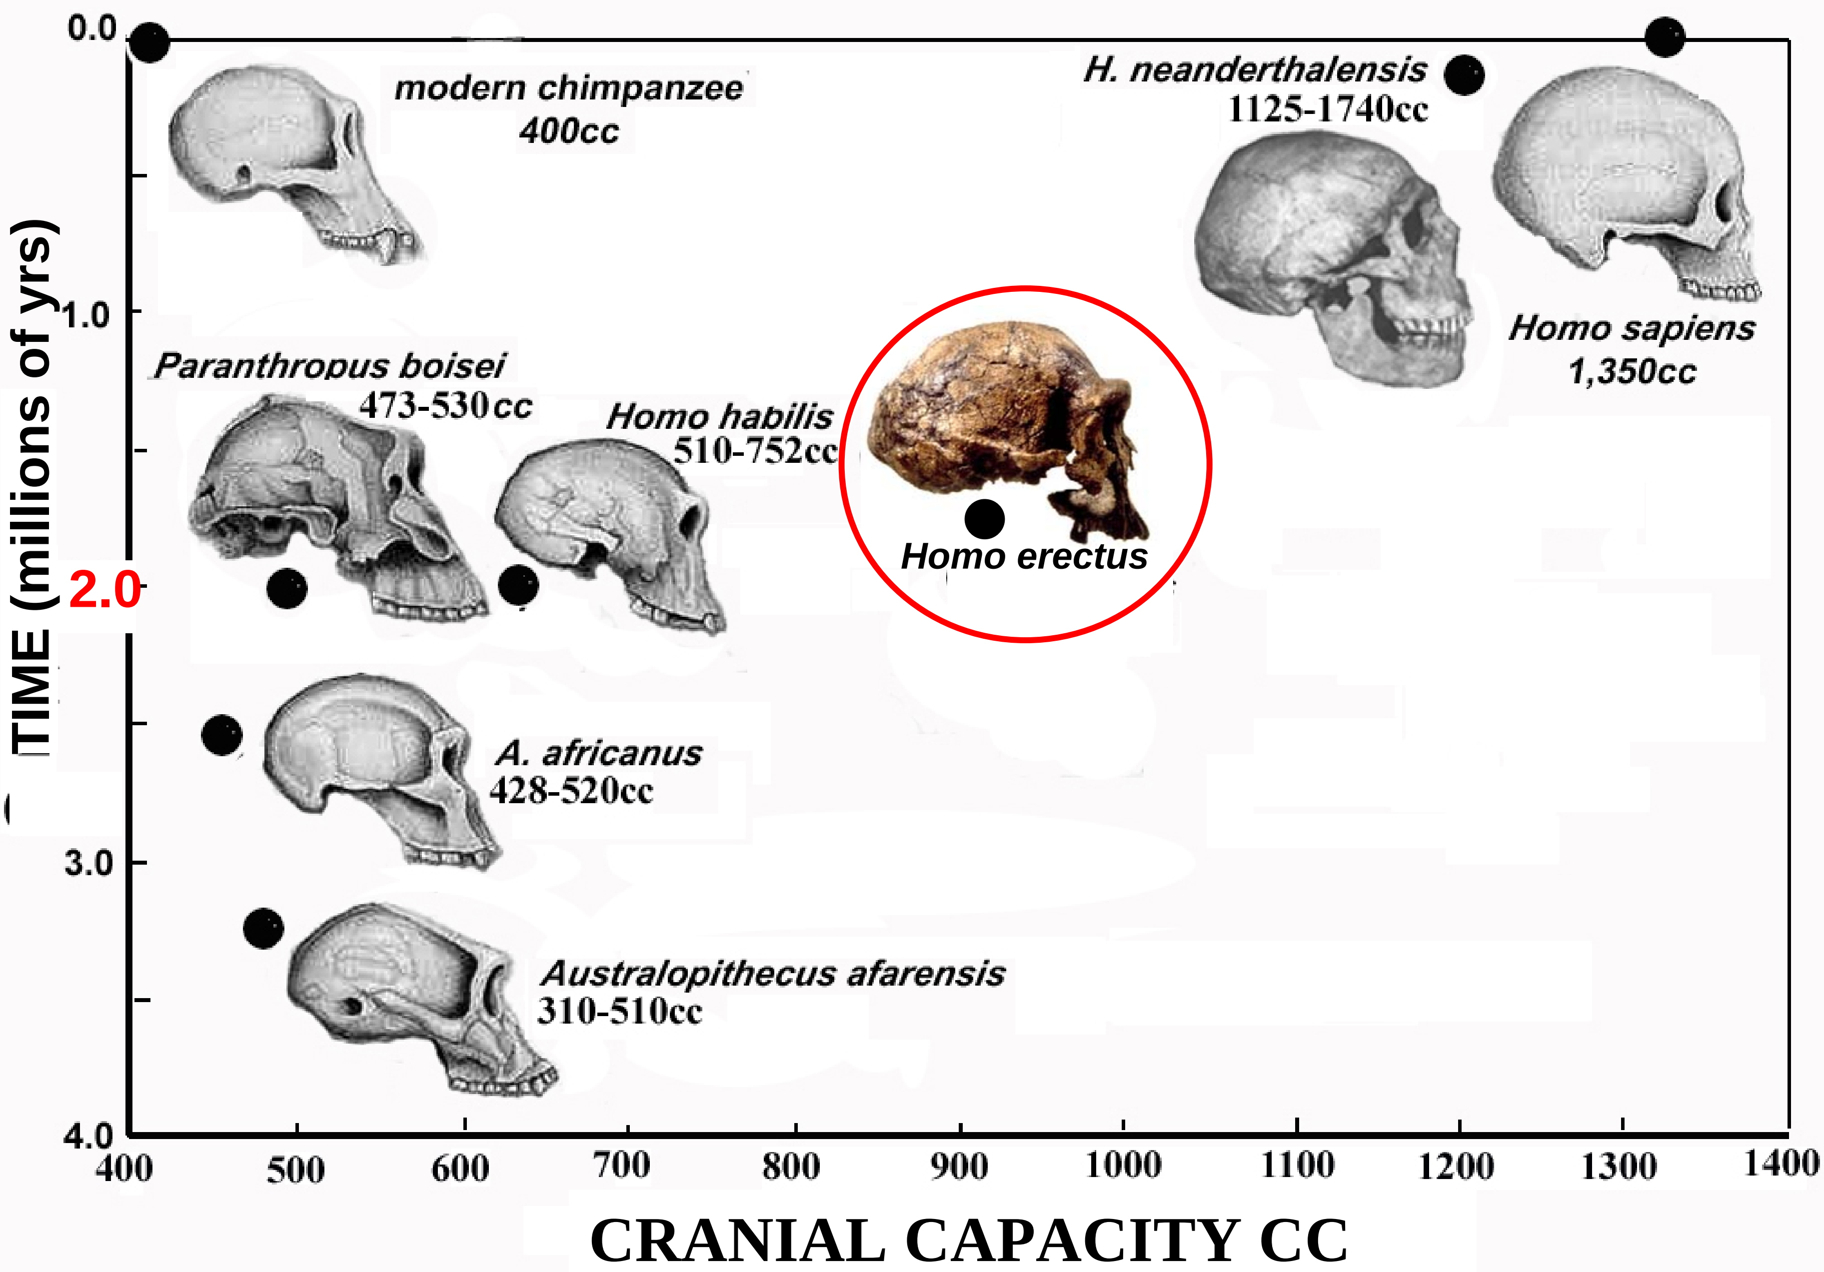
0.0
1.0
2.0
3.0
4.0
TIME (millions of yrs)
400
500
600
700
800
900
1000
1100
1200
1300
1400
CRANIAL CAPACITY CC
modern chimpanzee
400cc
H. neanderthalensis
1125-1740cc
Homo sapiens
1,350cc
Paranthropus boisei
473-530cc
Homo habilis
510-752cc
Homo erectus
A. africanus
428-520cc
Australopithecus afarensis
310-510cc

Every creature alive today had ancestors

Not all fossils had descendents

Mitochondrial DNA (mtDNA)

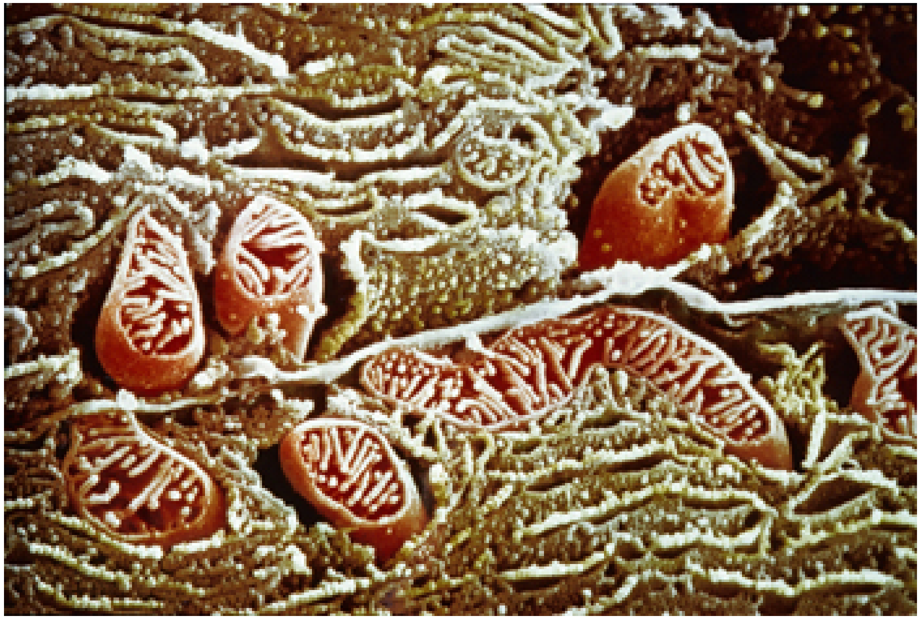

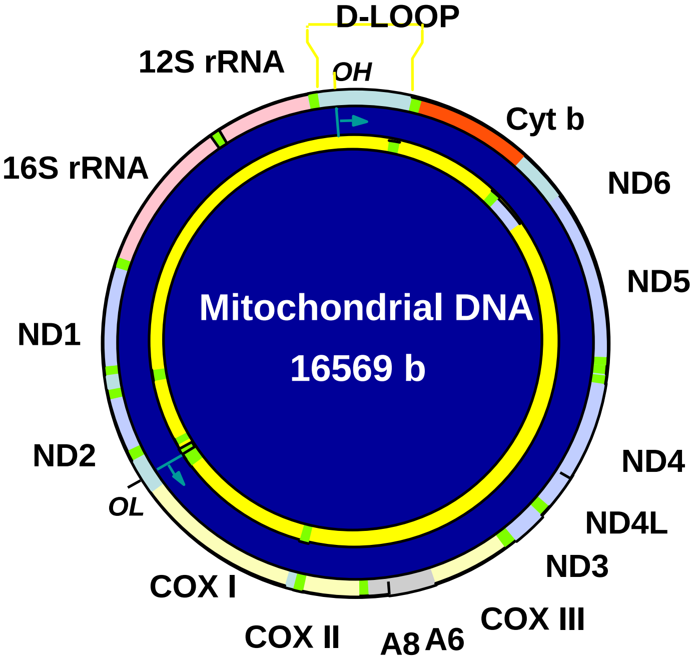

- ~ bacterial origin. Symbiosis 2 billion years ago
- Code for many critical genes (eg. tRNA, rRNA)
- 1981: Human mitochondrial genome was sequenced
- Strict maternal inheritance, no recombination
- High mutation rate, μ = 0.0025 mutations/generation

Properties of the Y chromosome

- Paternal inheritance
- No recombination
  - (except the ‘pseudoautosomal regions’ at the ends of the chromosome)
- Consequences:
  - Long stretch of DNA, highly informative
  - Patterns of variation are simple to interpret

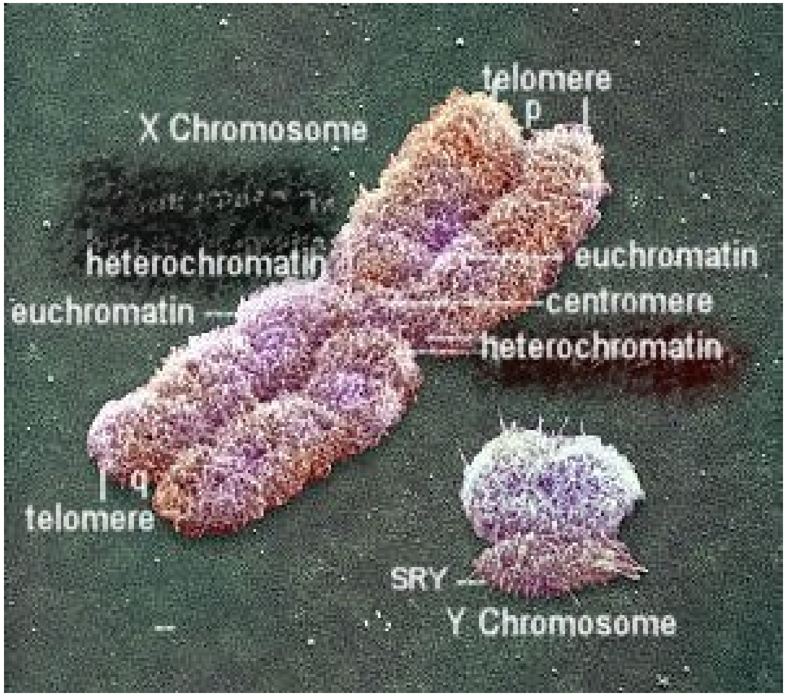

Uniparental Inheritance of Y & mtDNA

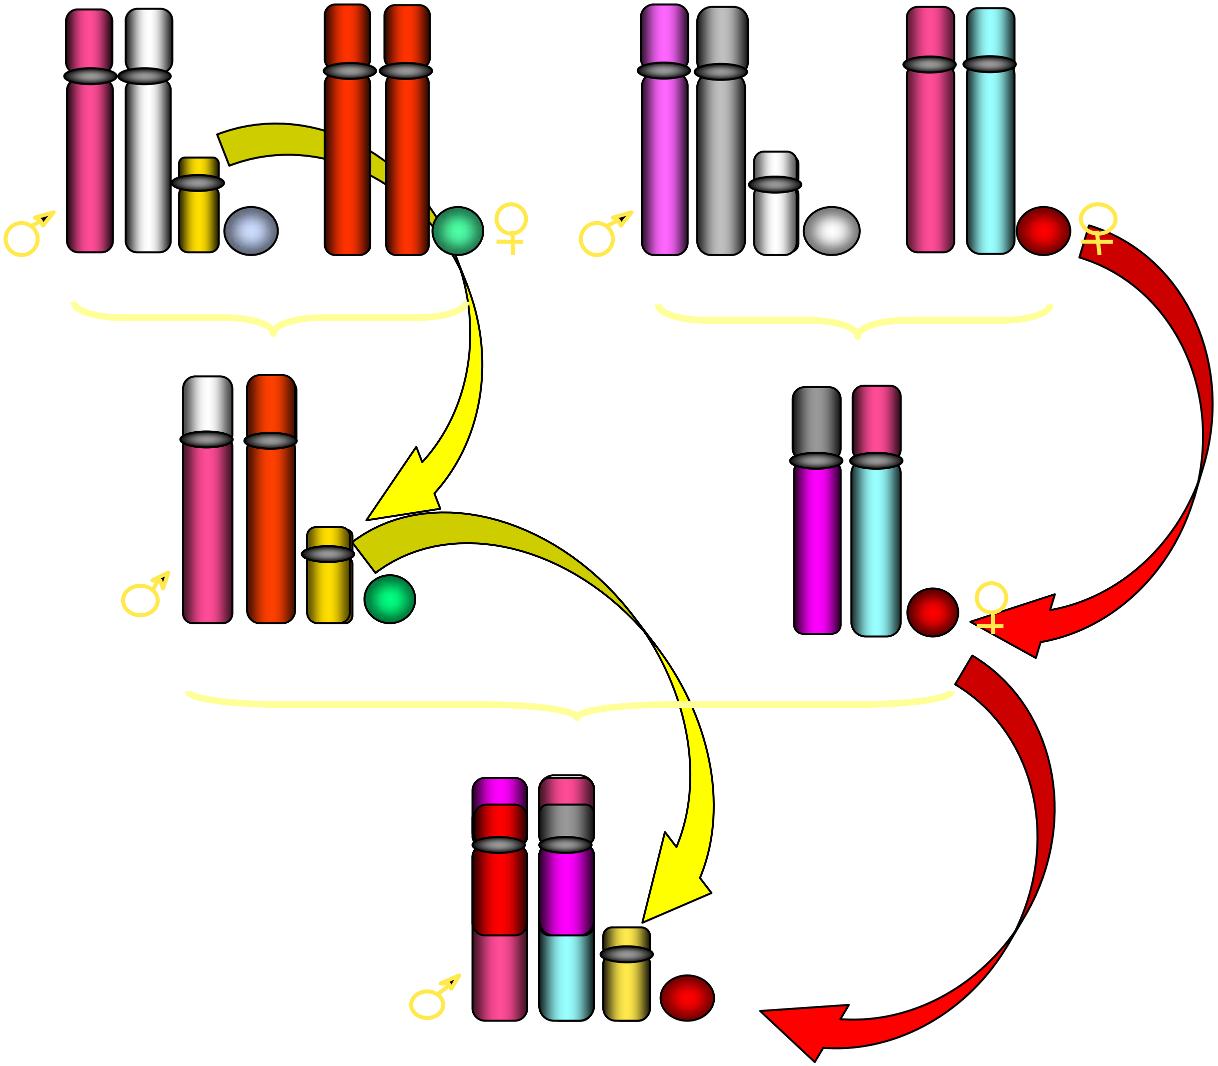

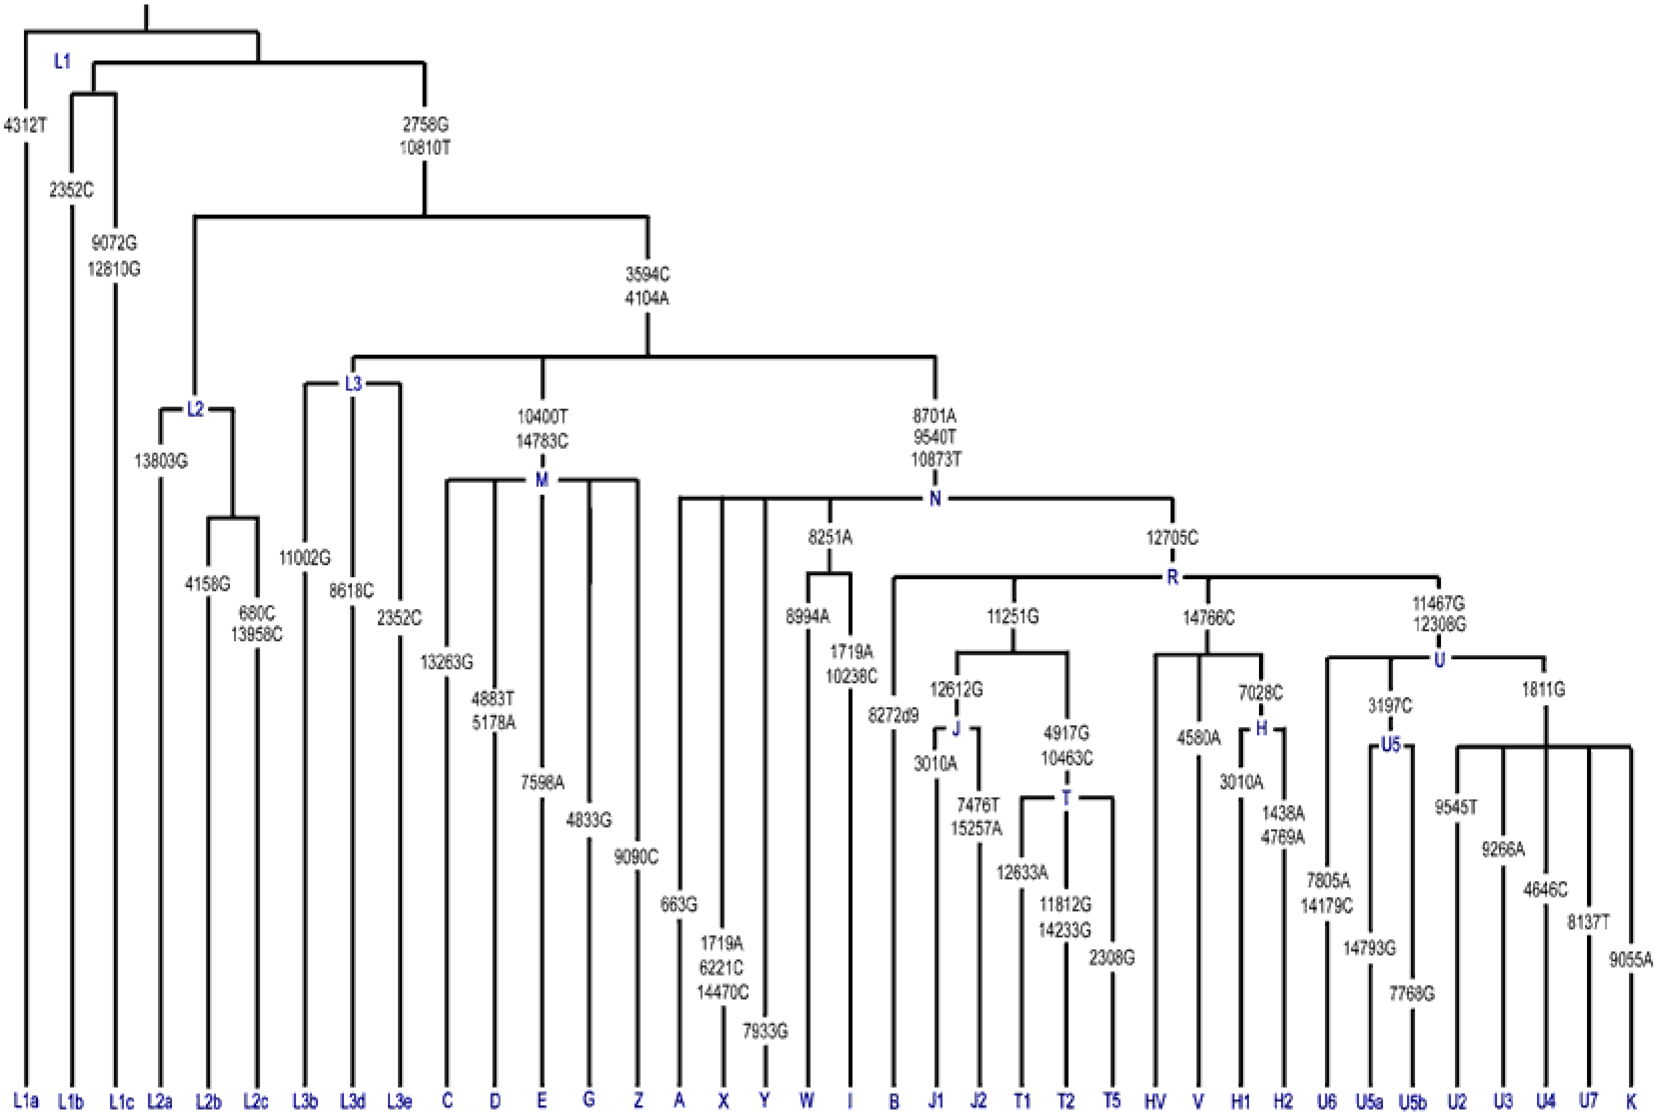
L1
4312T
2758G
10810T
2352C
9072G
12810G
3594C
4104A
L3
L2
10400T
14783C
8701A
9540T
10873T
13803G
M
N
8251A
12705C
11002G
4158G
8618C
680C
13958C
2352C
R
8994A
11467G
12308G
11251G
14766C
13263G
1719A
10238C
U
12612G
7028C
4883T
5178A
8272d9
3197C
1811G
J
H
4917G
10463C
U5
4580A
3010A
3010A
7598A
7476T
15257A
T
1438A
4769A
4833G
9545T
9090C
9266A
12633A
7805A
14179C
663G
11812G
14233G
4646C
1719A
6221C
14470C
8137T
14793G
2308G
9055A
7768G
7933G
L1a
L1b
L1c
L2a
L2b
L2c
L3b
L3d
L3e
C
D
E
G
Z
A
X
Y
W
I
B
J1
J2
T1
T2
T5
HV
V
H1
H2
U6
U5a
U5b
U2
U3
U4
U7
K

Human mtDNA Migrations

http://www.mitomap.org/mitomap/WorldMigrations.pdf

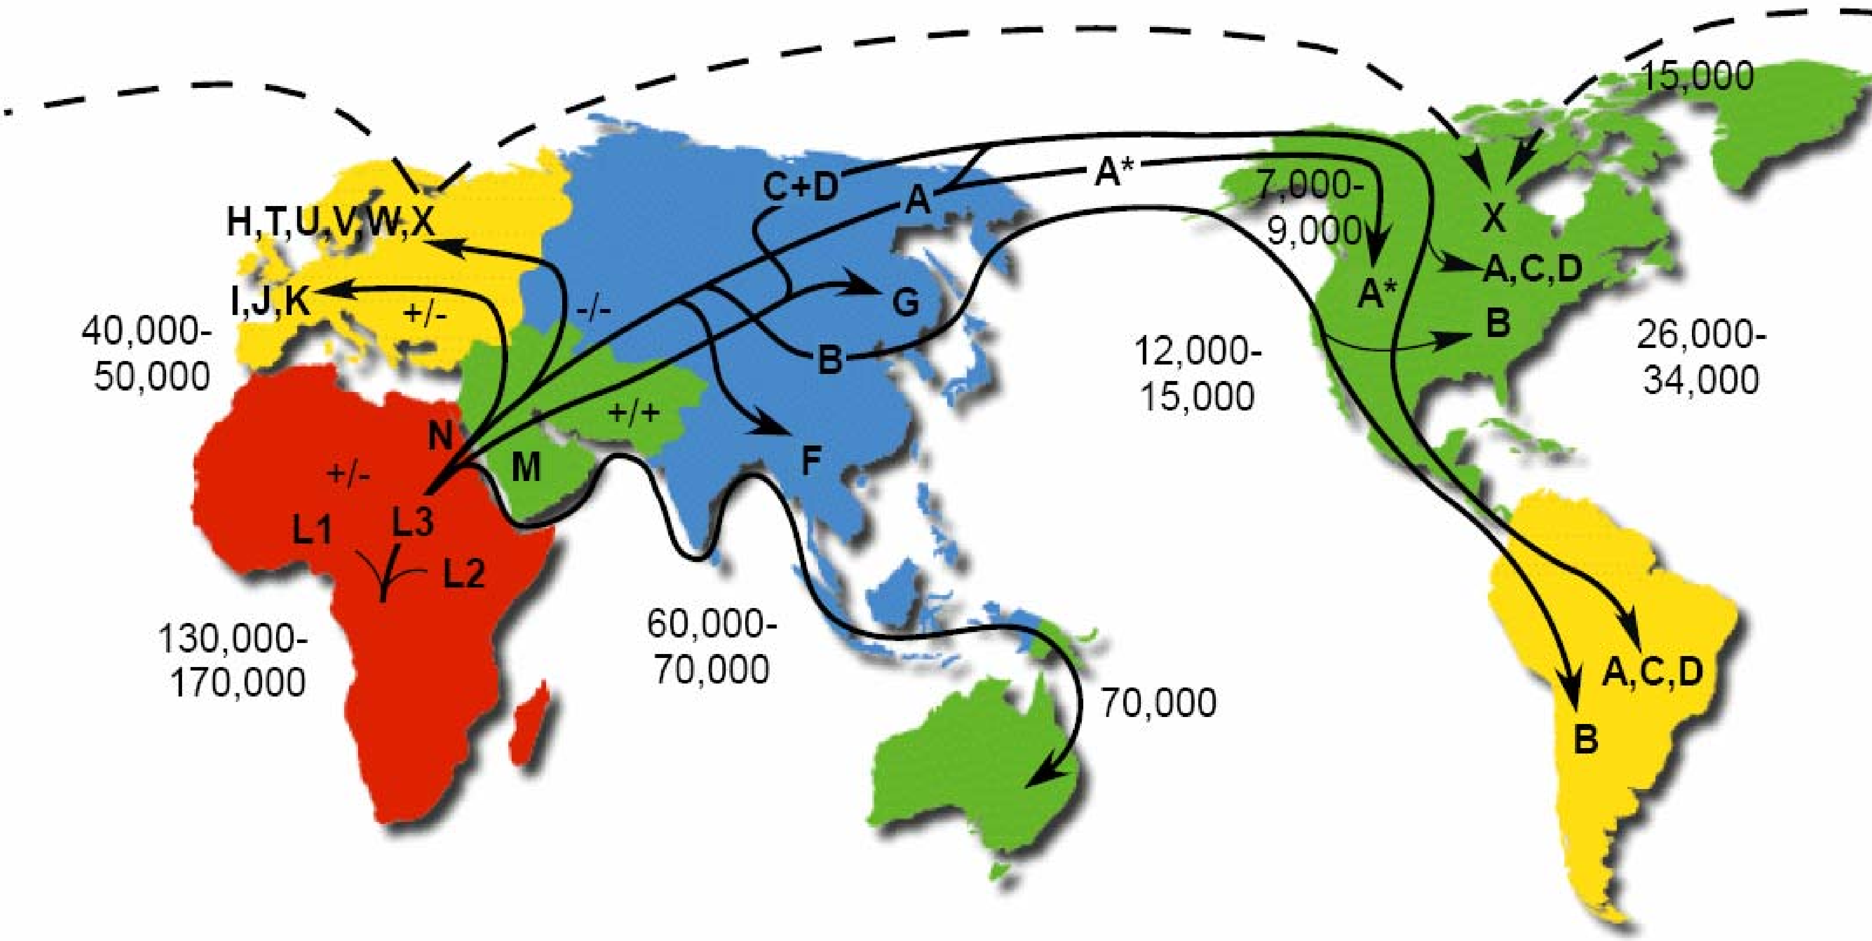

+/-, +/+, or -/- = Dde I 10394 / Alu I 10397

* = Rsa I 16329

Mutation rate = 2.2 - 2.9 % / MYR

Time estimates are YBP

N clade

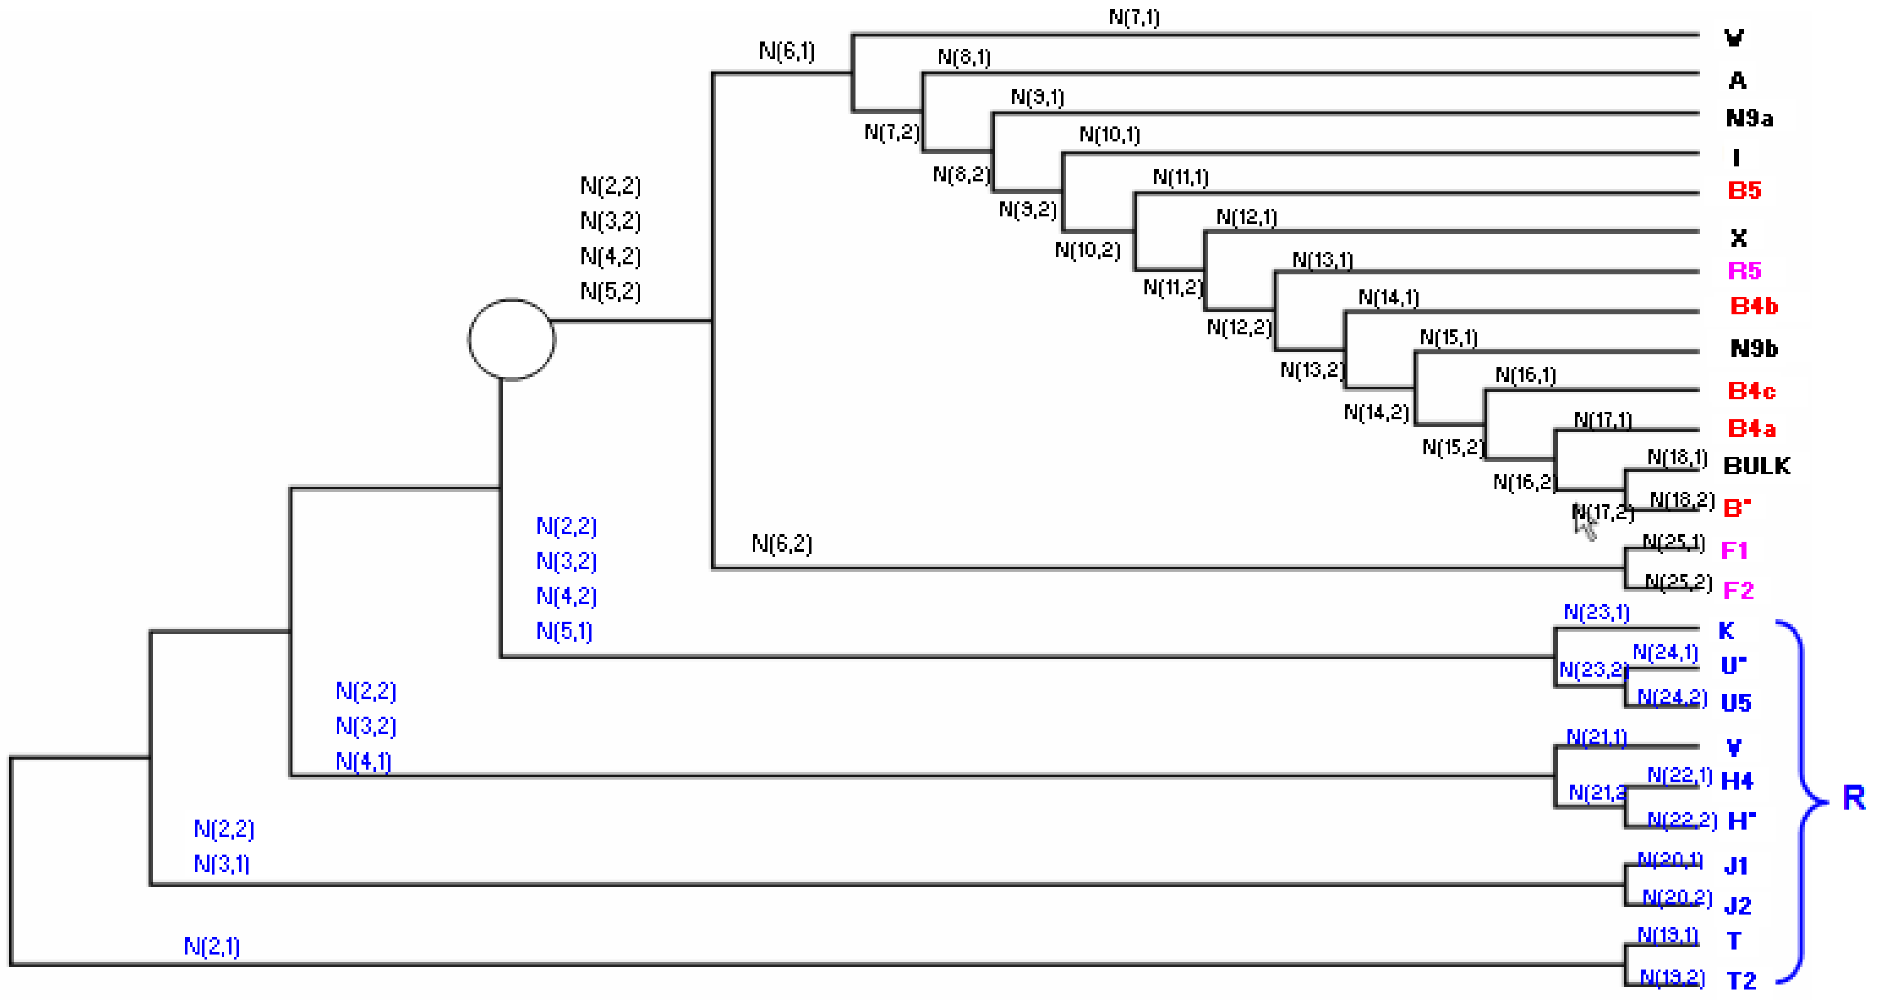
N(7,1)
N(6,1)
N(8,1)
N(7,2)
N(9,1)
N(8,2)
N(10,1)
N(9,2)
N(11,1)
N(10,2)
N(12,1)
N(11,2)
N(13,1)
N(12,2)
N(14,1)
N(13,2)
N(15,1)
N(14,2)
N(16,1)
N(15,2)
N(17,1)
N(16,2)
N(18,1)
N(18,2)
N(17,2)
N(25,1)
N(25,2)
N(6,2)
N(2,2)
N(3,2)
N(4,2)
N(5,2)
N(2,2)
N(3,2)
N(4,2)
N(5,1)
N(23,1)
N(24,1)
N(23,2)
N(24,2)
N(2,2)
N(3,2)
N(4,1)
N(21,1)
N(22,1)
N(21,2
N(22,2)
N(2,2)
N(3,1)
N(20,1)
N(20,2)
N(2,1)
N(19,1)
N(19,2)
W
A
N9a
I
B5
X
R5
B4b
N9b
B4c
B4a
BULK
B*
F1
F2
K
U*
U5
V
H4
H*
J1
J2
T
T2
R

Successes so Far

- ~ 130,000 samples collected from public
- Analysis of 1737 publically available sequences
    - – New, accurate tree of human Migration
    - – robust protocols for haplogroup assignment
    - – method to find ancient mutations and ages of clades
    - – direct evidence for 2 migrations giving rise to N and M clades
    - – Inferring movement of peoples (eg. migration history into China, Japan, India, Americas)

Physics is what physicists do at night.

R. Feynman

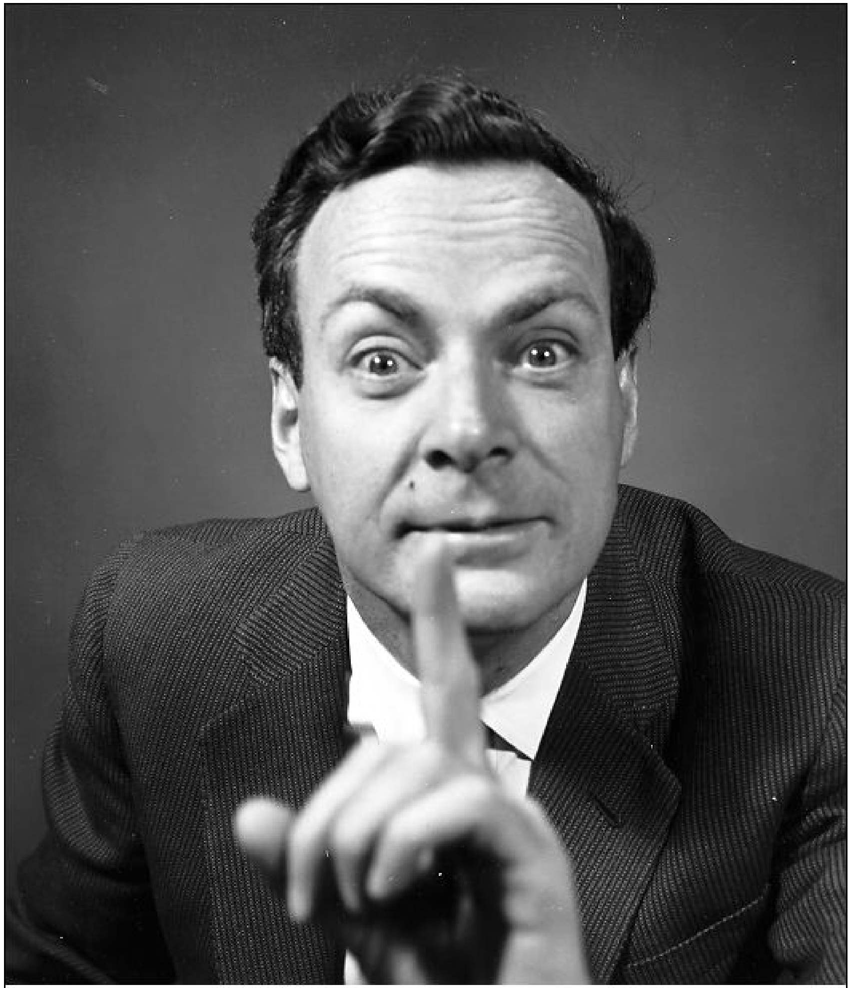

Biology some
~~Physics~~ is what ^ physicists do at night.

E. Schrödinger

WHAT IS LIFE?

DOMO ARIGATO