From the Web to the Grid

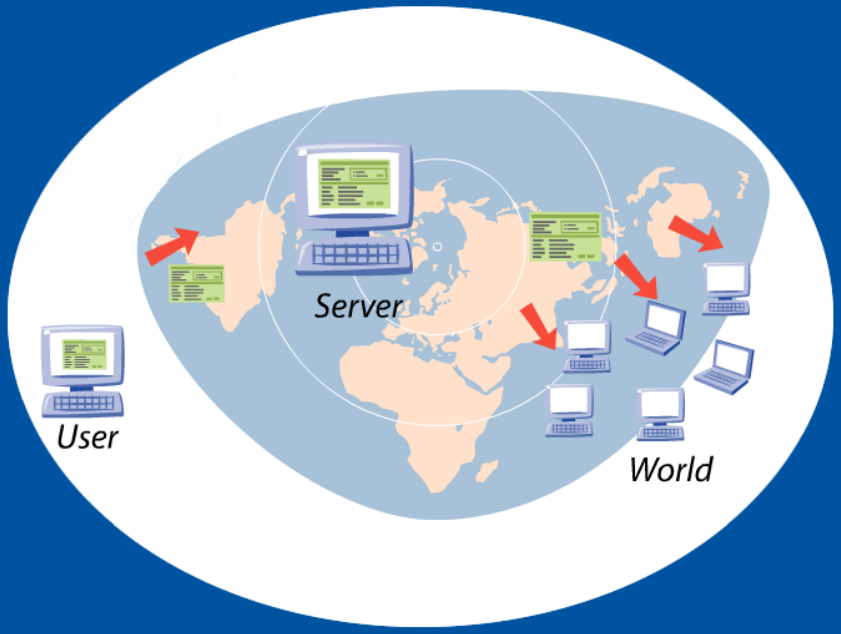

CERN

How did the Grid start?

- Name “Grid” chosen by analogy with electric power grid (Foster and Kesselman 1997)
- Vision: plug-in computer for processing power just like plugging in toaster for electricity.
- Concept has been around for decades (distributed computing, metacomputing)
- Key difference with the Grid is to realise the vision on a global scale.

Different Grids for different needs

There is as yet no unified Grid, like there is a single web.

Rather **there are many Grids** for many applications:

- **Enterprise Grids** link together PCs within one company.
- **Volunteer computing** links together public computers.
- **Scientific Grids** link together major computing centres.
- Latest trend **federates national Grids** into global Grid infrastructure.
- High Energy Physics is a driving force for this.

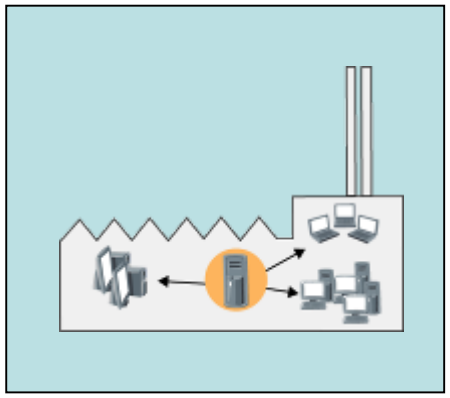

The LHC data challenge

- 40 million bunch collisions per second
- After filtering, ~100 collisions of interest per second per detector
- > 1 Megabyte of data per collision recording rate > 1 Gigabyte/sec
- 10^{10} collisions recorded each year stored data ~15 Petabytes/year

  …for more than 10 years

1 Megabyte (1MB)
A digital photo

1 Gigabyte (1GB)
= 1000MB
5GB = A DVD movie

1 Terabyte (1TB)
= 1000GB
World annual book production

1 Petabyte (1PB)
= 1000TB
Annual production of one LHC experiment

1 Exabyte (1EB)
= 1000 PB
3EB = World annual information production

Data Storage for the LHC

- LHC data correspond to about 20 million CDs each year!

Where will the experiments store all of these data?

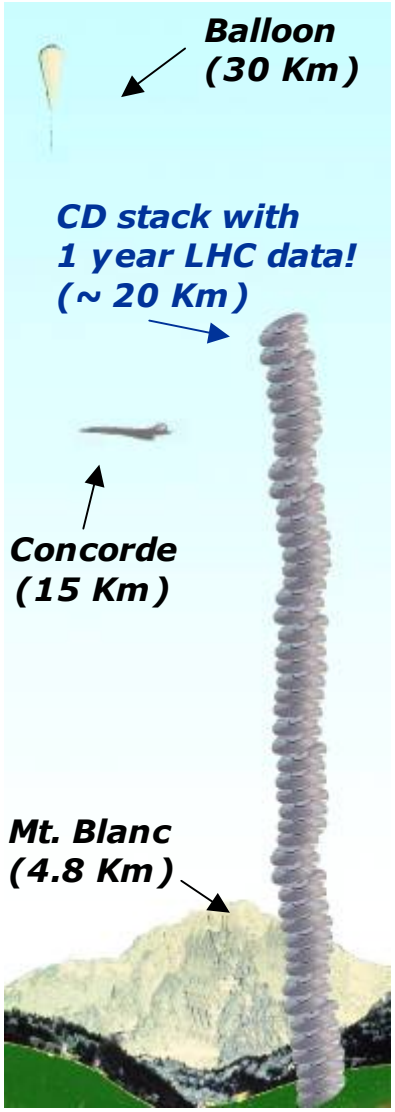

Data Processing for the LHC

- LHC data analysis requires a computing power equivalent to ~ 100,000 of today's PC processors!

Where will the experiments find such a computing power?

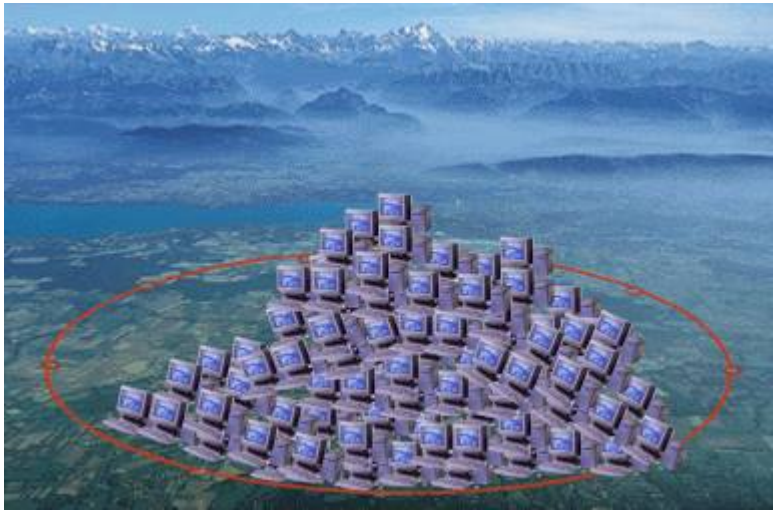

Computing for LHC

- Problem: even with Computer Centre upgrade, CERN can provide only a fraction of the necessary resources.

- Solution: CERN has over 250 partner institutes in Europe, over 200 in rest of the world. Most have significant computing resources. Build a Grid that **unites these computing resources.**

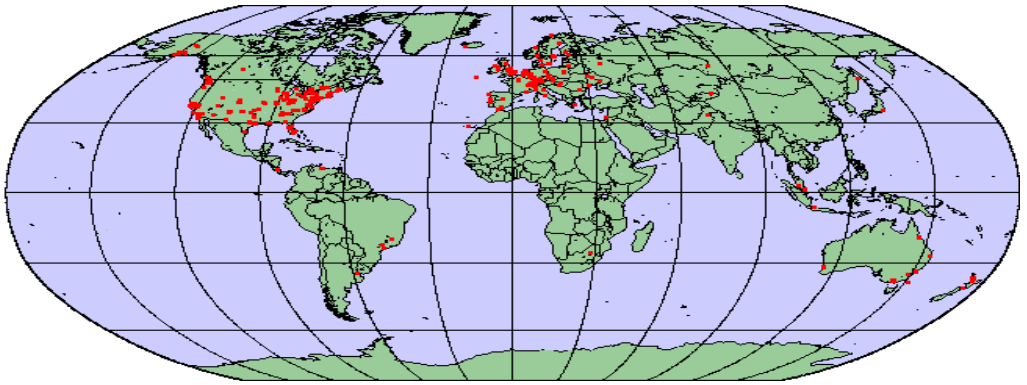

How does the Grid work?

- It relies on advanced software, called **middleware**.
- Middleware automatically finds the **data** the scientist needs, and the **computing power** to analyse it.
- Middleware balances the load on different resources. It also handles **security, accounting, monitoring** and much more.

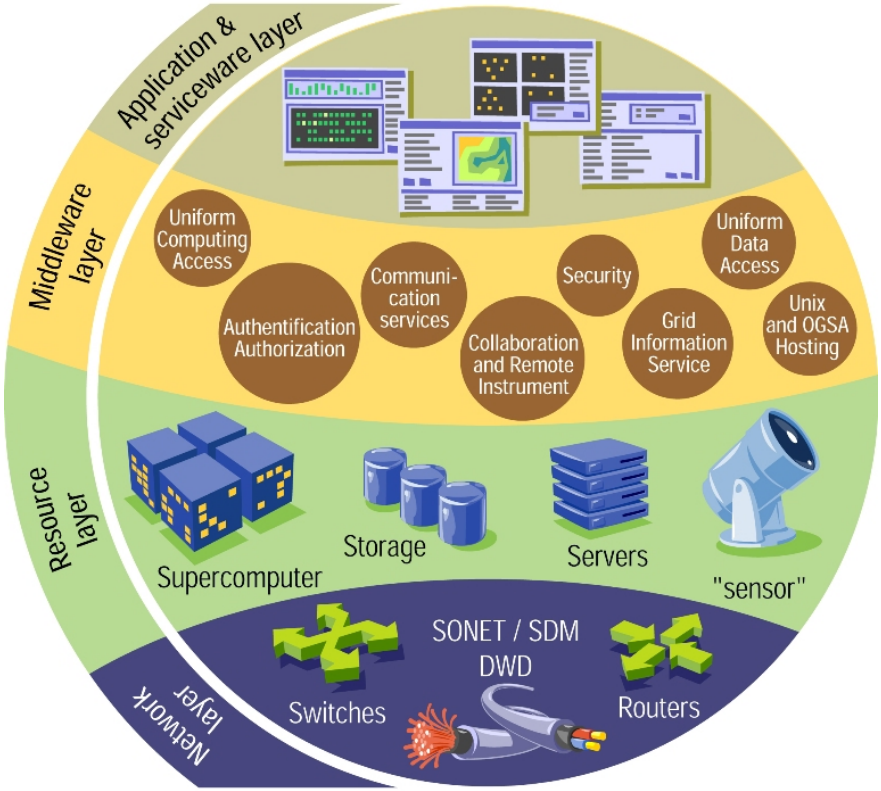

Grid Applications for Science

- Medical/Healthcare *imaging, diagnosis and treatment*
- Bioinformatics *study of the human genome and proteome*
- Nanotechnology *design of new materials from the molecular scale*
- Engineering *design optimization, simulation, failure analysis*
- Natural Resources and the Environment *climate modelling, earth observation*

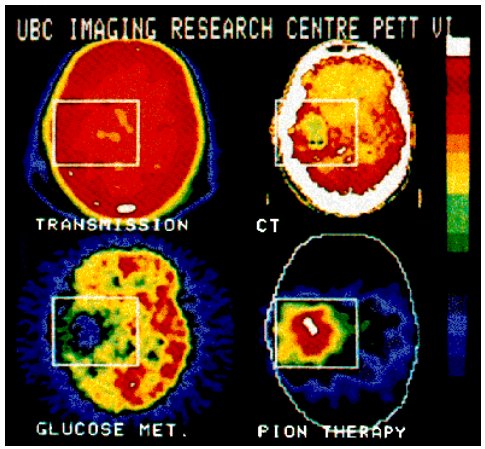

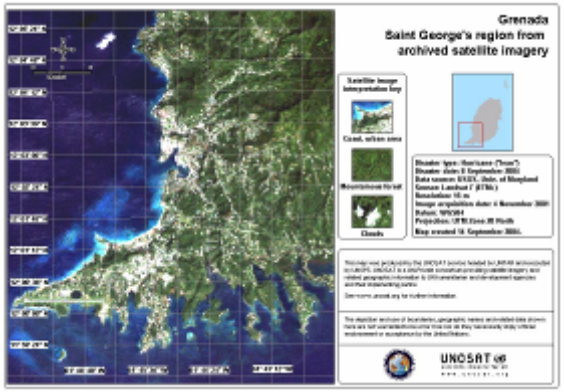

Grid @ CERN

- CERN projects:
  - LHC Computing Grid (LCG)

- EU-funded projects led by CERN:
  - Enabling Grids for E-SciencE (EGEE)

- Industry funded projects:
  - CERN openlab for DataGrid applications

LHC Computing Grid Project (LCG)

- Timeline:
  - 2002: start project
  - 2003: service opened (LCG-1 started in September with 12 sites)
  - 2005 >100 sites contributing, >10k CPUs,
  - 2006: 1GB/s sustained to Tier-1 centres
  - 2007: full operation of LHC computing service.

Worldwide LHC Computing Grid (WLCG)

- More than 100 computing centres
- 12 large centres for primary data management: CERN (Tier-0) and eleven Tier-1s
- 38 federations of smaller Tier-2 centres
- 32 countries involved

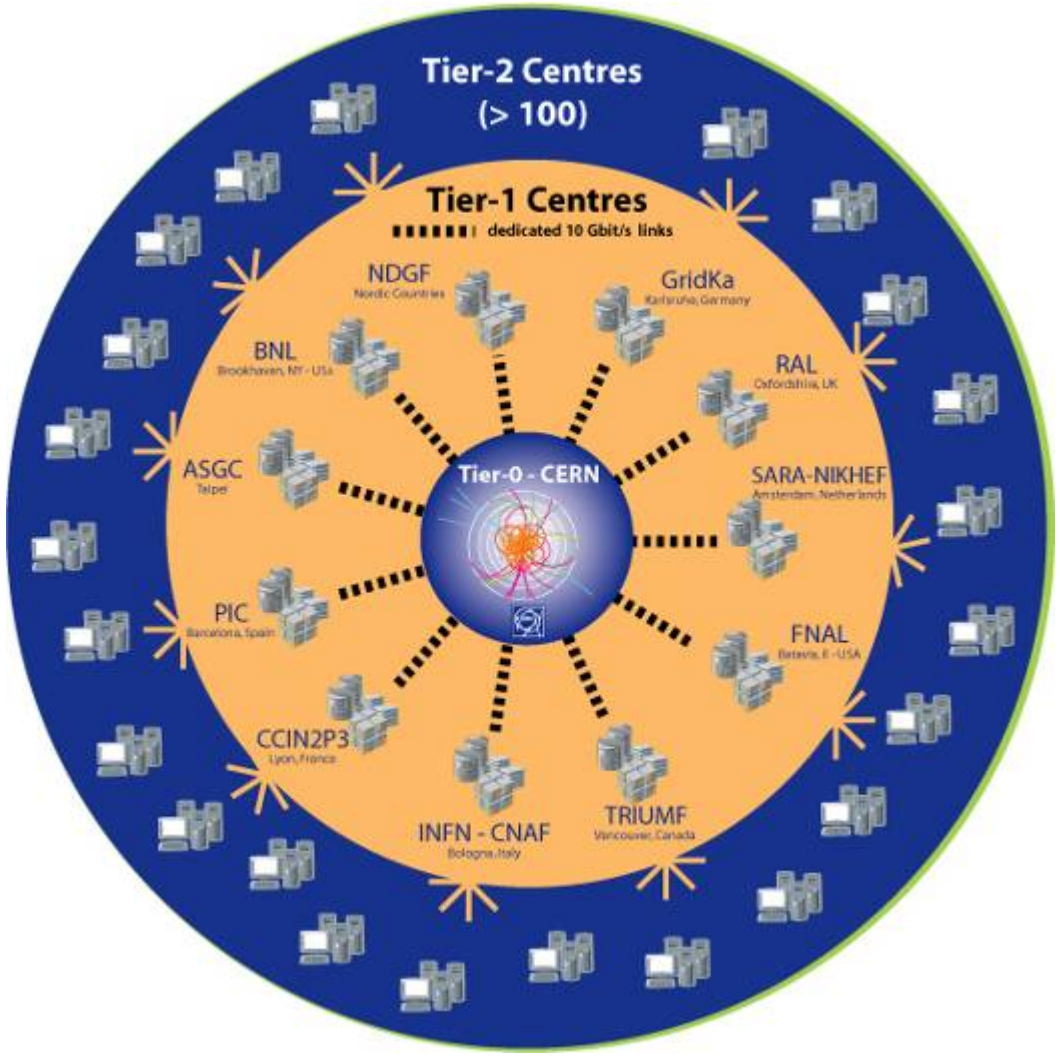

CERN openlab

- Platform competence centre
- Grid interoperability centre
- Security activities
- Networking activities
- Student programme
- Joint events

www.cern.ch/openlab

PARTNERS

CONTRIBUTORS

STONESOFT

SETI@home: >500,000 CPUs

Einstein@home: searching for pulsar gravitational waves with LIGO and GEO

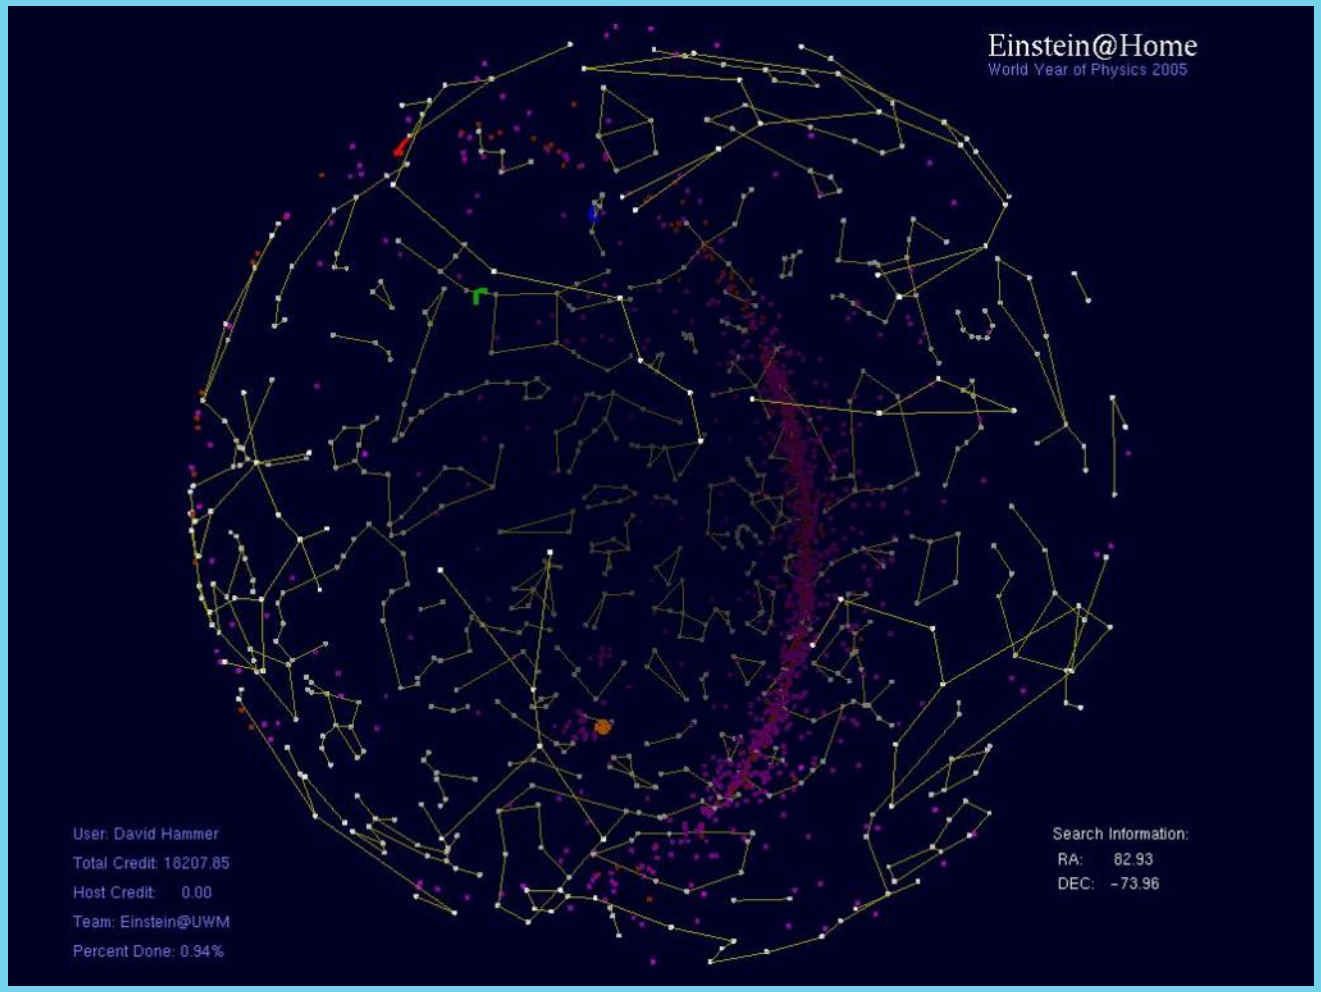

ClimatePrediction.net: modelling future of Earth's climate

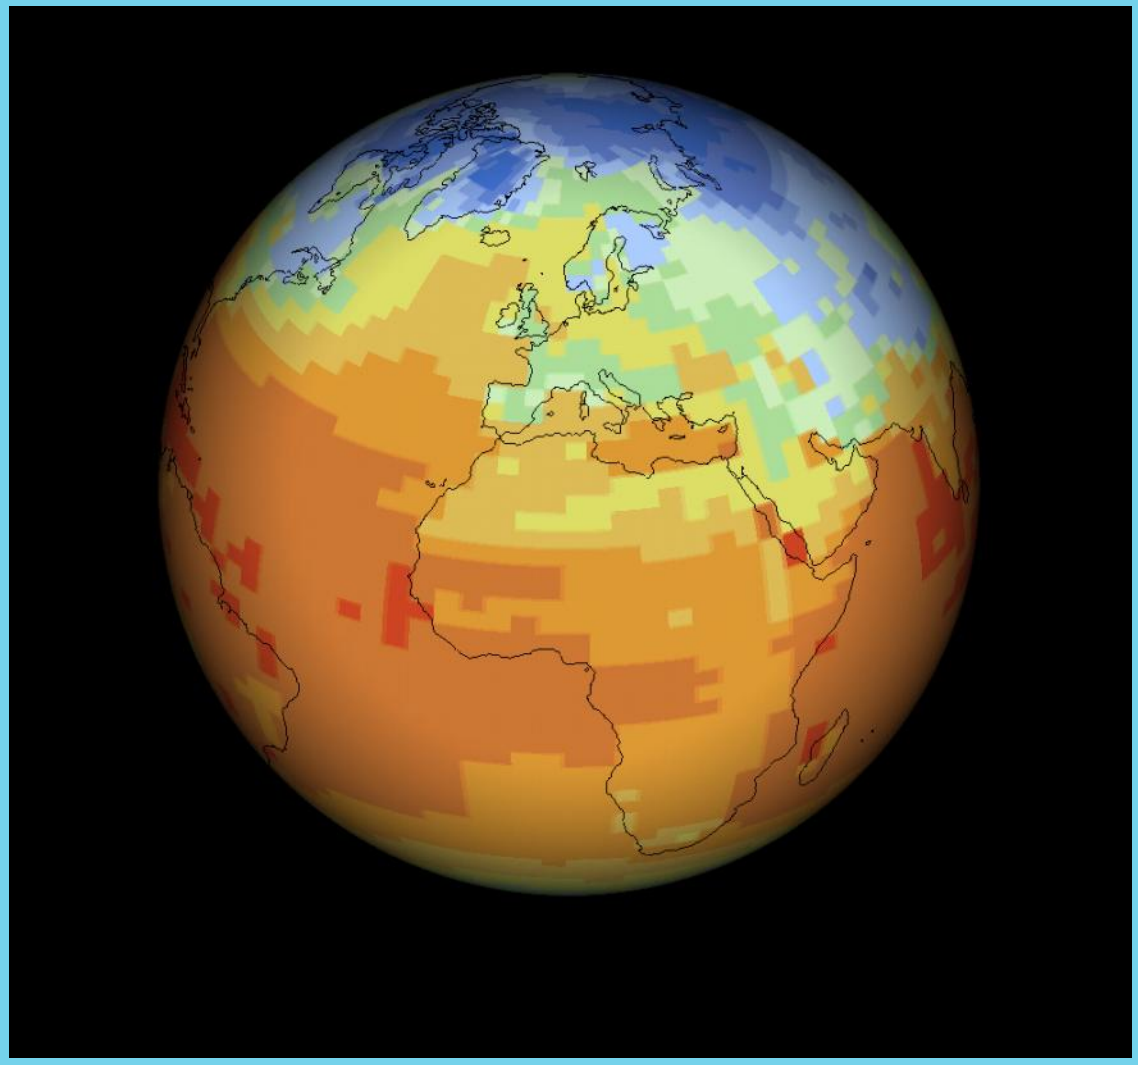

Folding@home: >1 petaflop using 50k Playstation-3s

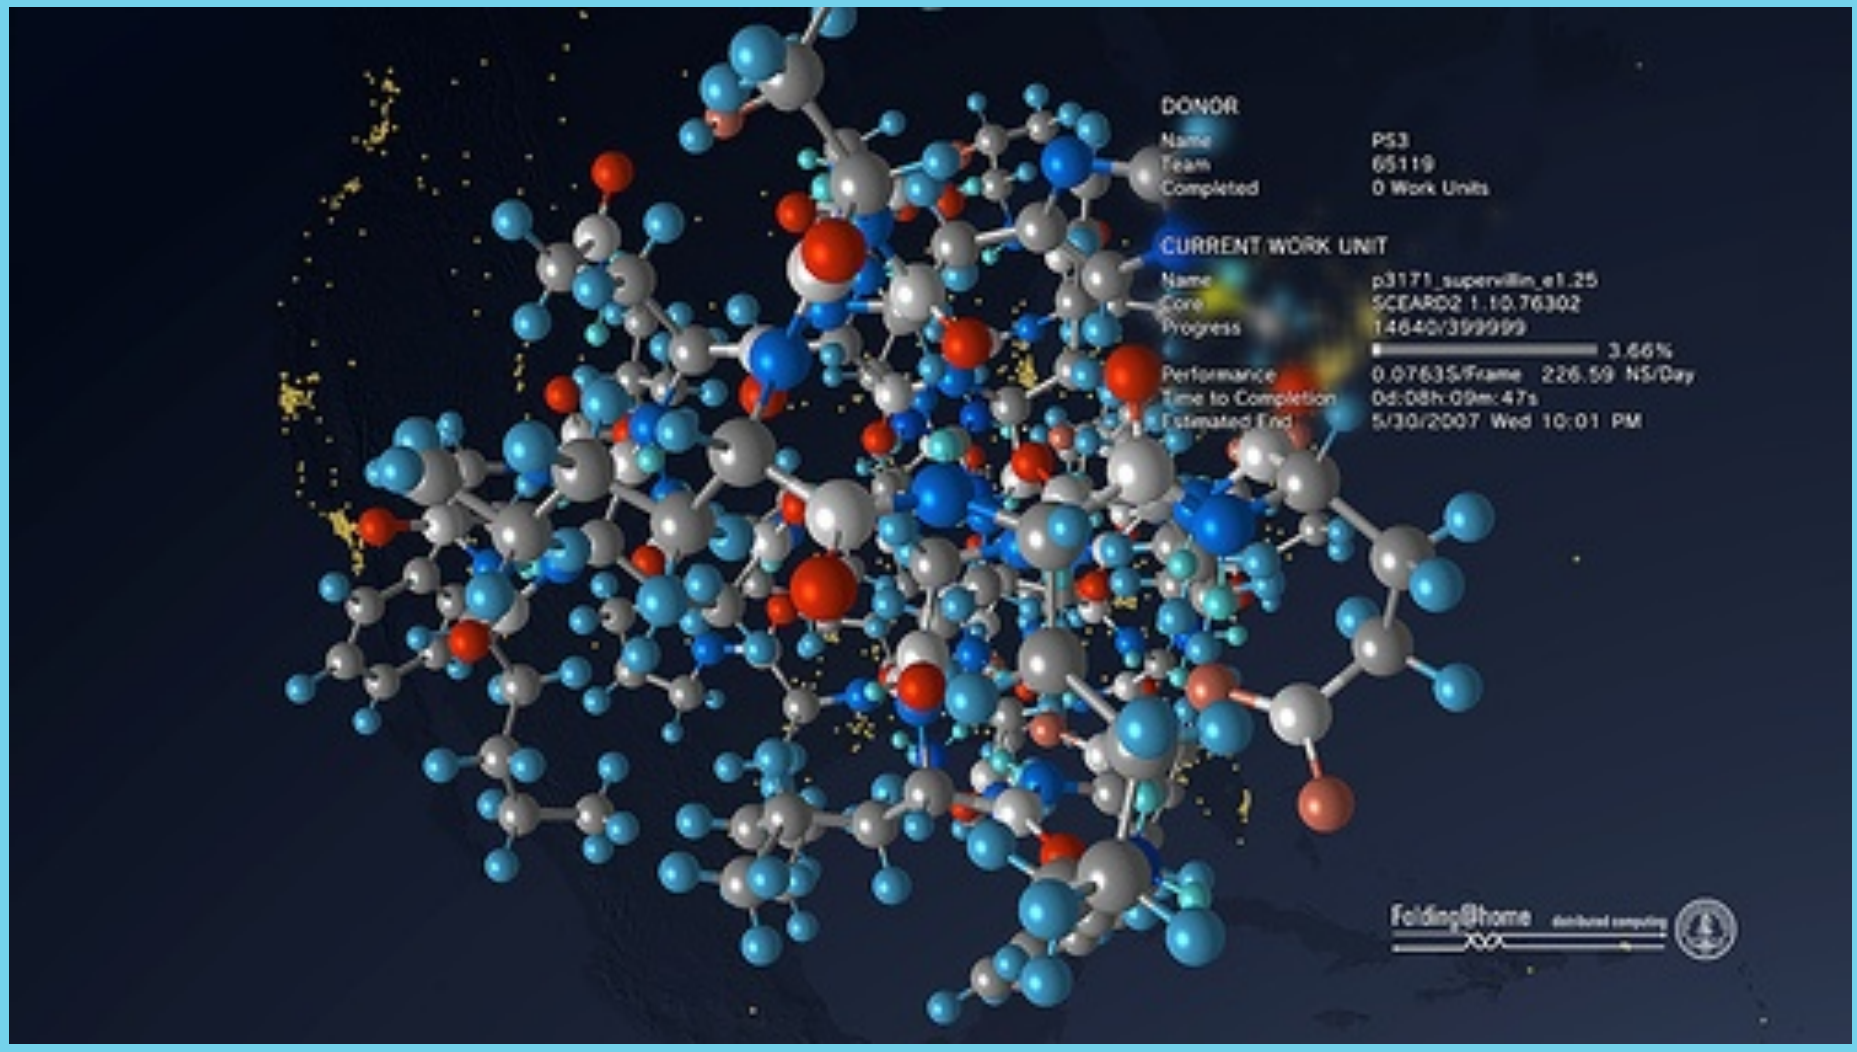

Spot the ~~computers~~ scientists!

GalaxyZoo: classify galaxies

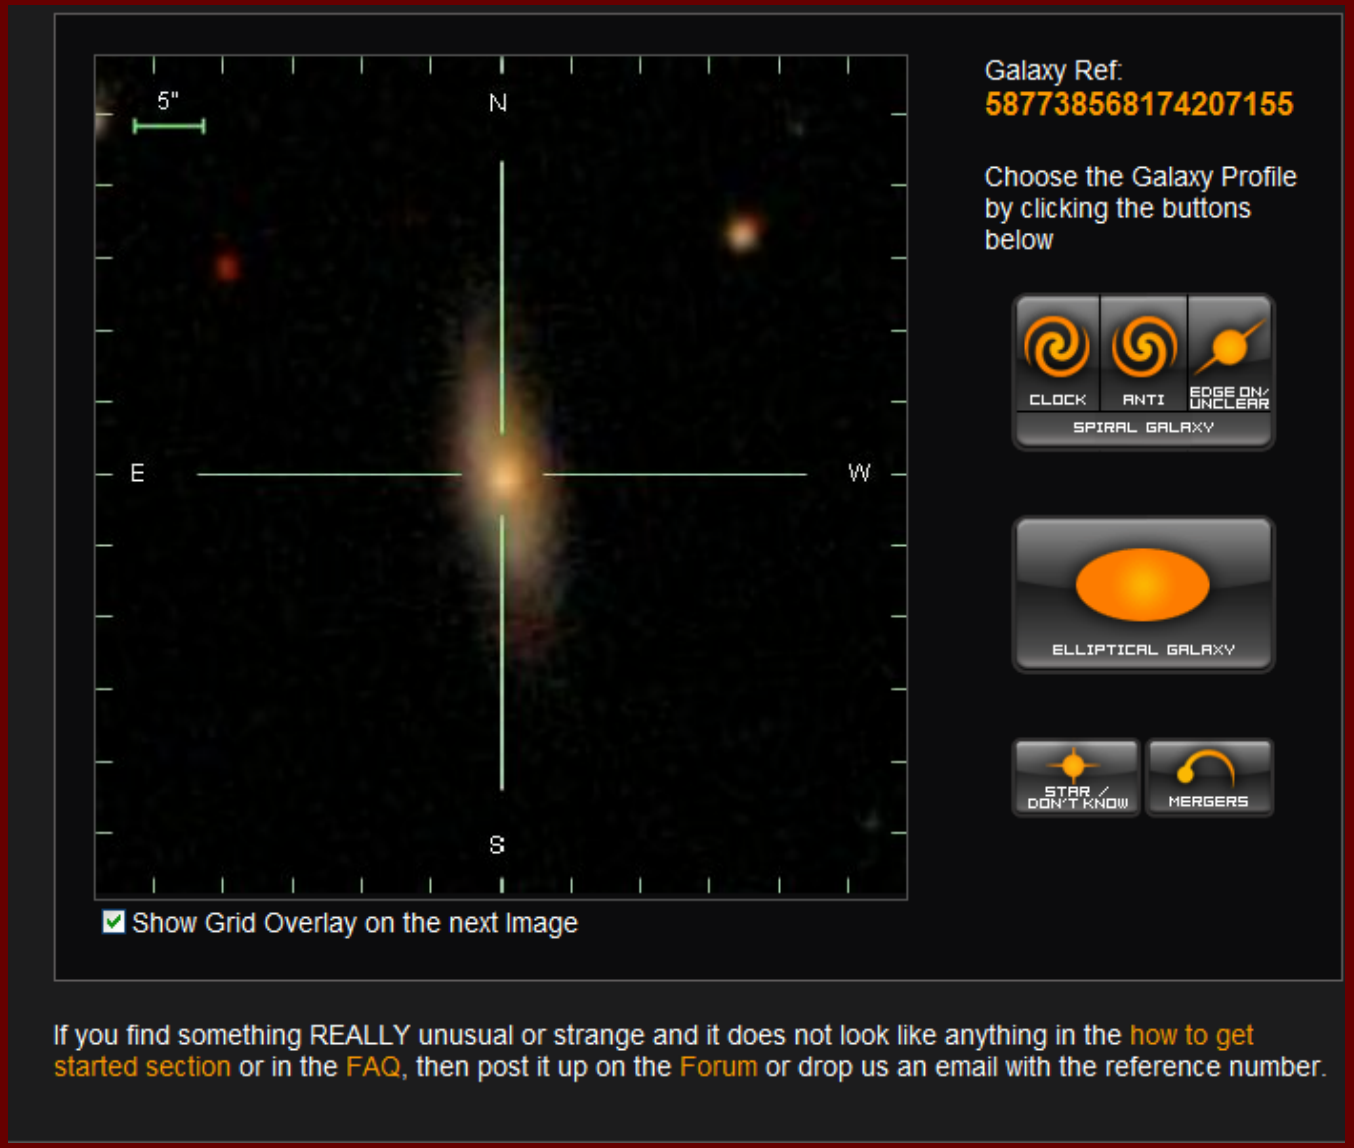

Stardust@home: find cosmic dust tracks with a virtual microscope

Herbaria@home:
digitize 19th century plant archives

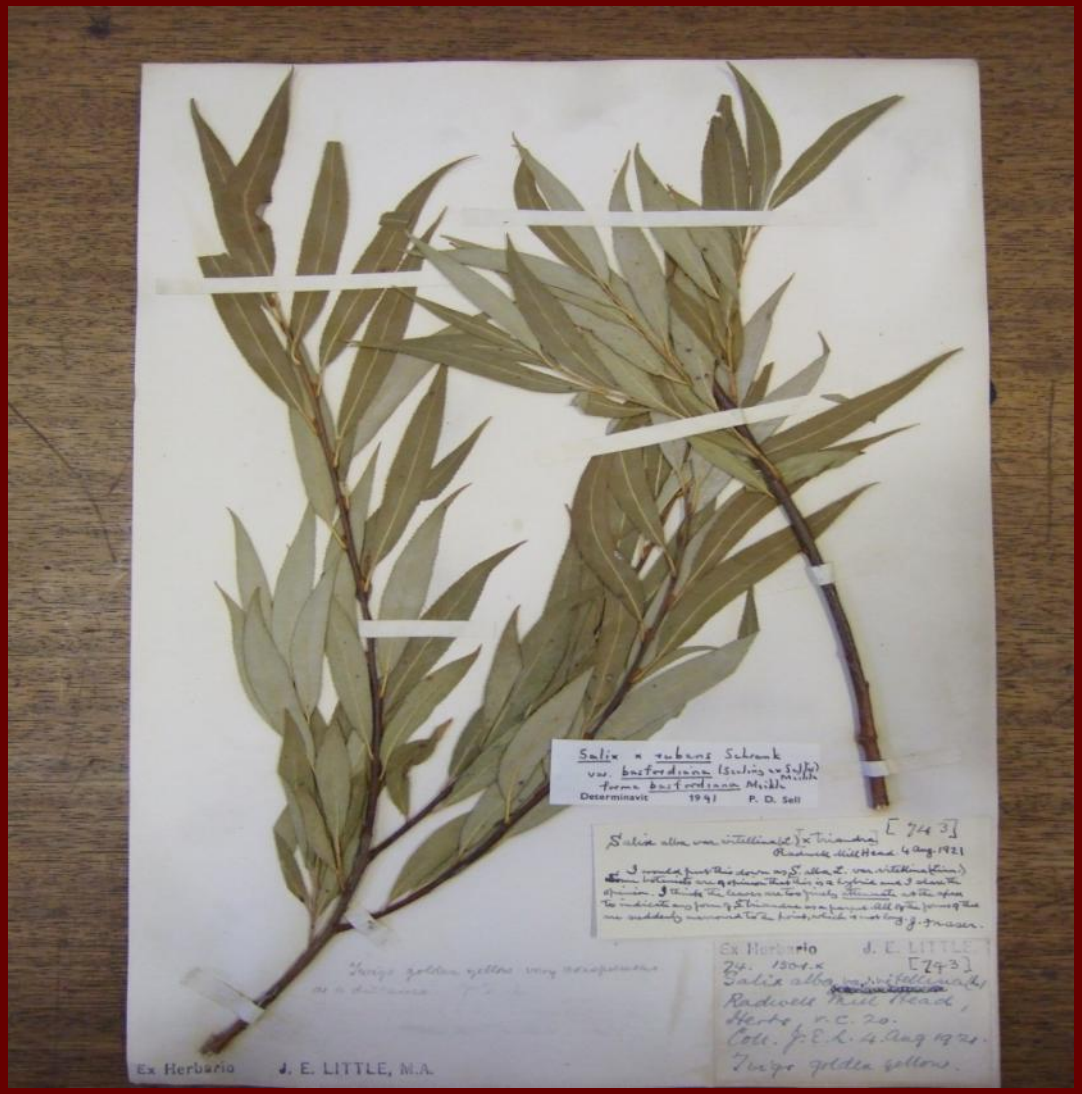

What is BOINC?

- "**B**erkeley **O**pen **I**nfrastructure for **N**etwork **C**omputing"
- **Software platform for distributed computing using volunteered computer resources**
- **http://boinc.berkeley.edu**
- **Uses a volunteer PC's unused CPU cycles to analyse scientific data**
- **Client-server architecture**
- **Free and Open-source**
- **Also handles DESKTOP GRIDS**

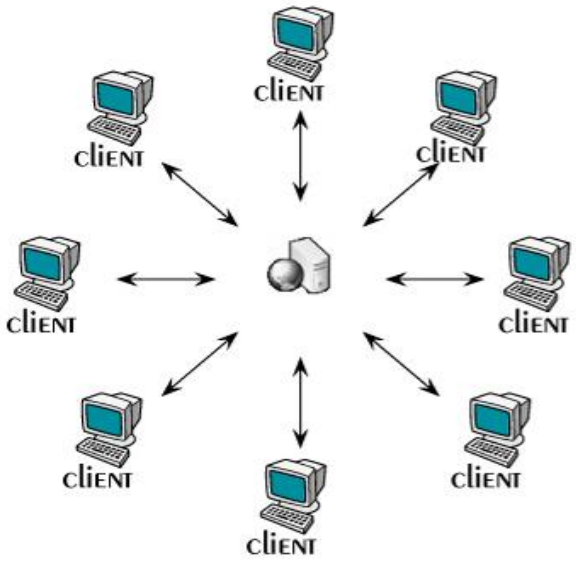

Basic structure of BOINC

- Interaction between **client and server**

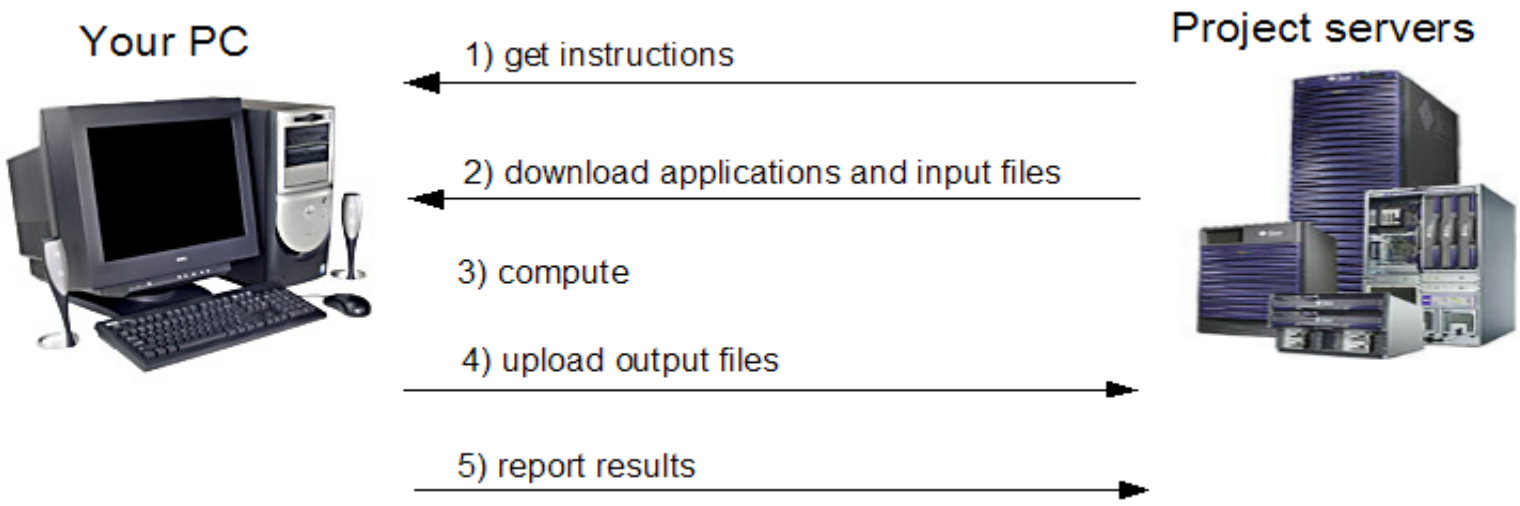

Some volunteer computing projects

SCIENCE
SETI@home (BOINC)
evolution@home
eOn
climateprediction.net (BOINC)
Muon1
LHC@home (BOINC)
Einstein@Home(BOINC)
BBC Climate Change Experiment (BOINC)
Leiden Classical (BOINC)
QMC@home (BOINC)
NanoHive@Home (BOINC)
µFluids@Home (BOINC)
Spinhenge@home (BOINC)
Cosmology@Home (BOINC)
PS3GRID (BOINC)
Mars Clickworkers

LIFE SCIENCES
Parabon Computation
Folding@home
FightAIDS@home
Übero
Drug Design Optimization Lab (D2OL)
The Virtual Laboratory Project
Community TSC
Predictor@home (BOINC)
XGrid@Stanford
Human Proteome Folding (WCG)
CHRONOS (BOINC)
Rosetta@home (BOINC)
RALPH@home (BOINC)
SIMAP (BOINC)
malariacontrol.net (BOINC)
Help Defeat Cancer (WCG)
TANPAKU (BOINC)
Genome Comparison (WCG)
Docking@Home (BOINC)
proteins@home (BOINC)
Help Cure Muscular Dystrophy (WCG)

MATHEMATICS & CRYPTOGRAPHY
Great Internet Mersenne Prime Search
Proth Prime Search
ECMNET
Minimal Equal Sums of Like Powers
MM61 Project
3x + 1 Problem
Distributed Search for Fermat Number Divisors
PCP@Home
Generalized Fermat Prime Search
PSearch
Seventeen or Bust
Factorizations of Cyclotomic Numbers
Goldbach Conjecture Verification
The Riesel Problem
The 3*2^n-1 Search
NFSNET
Search for Multifactorial Primes
15k Prime Search
ElevenSmooth
Riesel Sieve
The Prime Sierpinski Project
P.I.E.S. - Prime Internet Eisenstein Search
Factors of k*2^n±1
XYYXF
12121 Search
2721 Search
Operation Billion Digits
SIGPS
Primesearch

INTERNET PERFORMANCE
Gómez Performance ($)
Network Peer
NETI@home
dCrawl
DIMES
Red Library DLV
Majestic-12
Boitho
PeerFactor
DepSpid
Pingdom GIGRIB
Project Neuron(BOINC)

ECONOMICS
MoneyBee
Gstock

GAMES
ChessBrain
Chess960@home (BOINC)

ART
Electric sheep
Internet Movie Project
RenderFarm@home (BOINC)

Source: http://distributedcomputing.info/

The BOINC community

- Competition between individuals and teams for “credit”.
- Websites and regular updates on status of project by scientists.
- Forums for users to discuss the science behind the project.
- E.g. for LHC@home, the volunteers show great interest in CERN and the LHC.
- Supply each other with scientific information and even help debug the project.

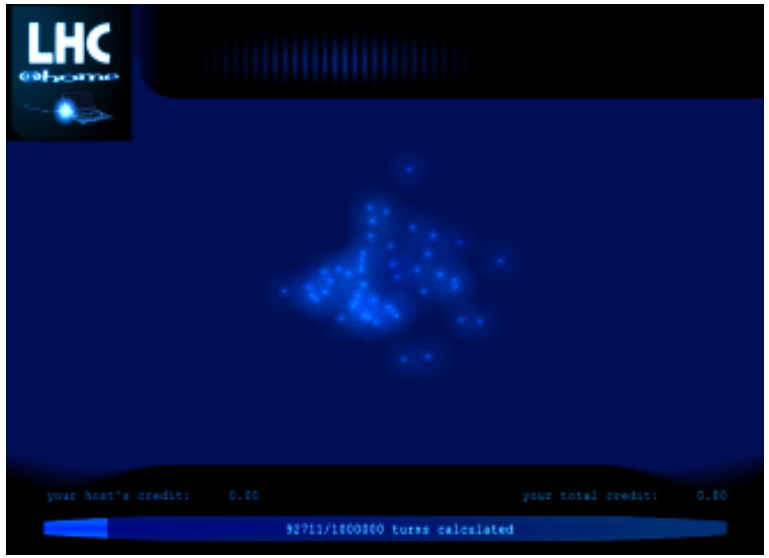

LHC@home screensaver

LHC@home

- Calculates stability of proton orbits in CERN's new LHC accelerator
- System is nonlinear and unstable so numerically very sensitive. Hard to get identical results on all platforms
- About 40 000 users, 70 000 PC's… over 1500 CPU years of processing
- Objectives: extra CPU power and raising public awareness of CERN and the LHC - both successfully achieved.
- Started as an outreach project for CERN 50th Anniversary 2004; used for Year of Physics (Einstein Year) 2005

SixTrack program

SixTrack is a Fortran program by F. Schmidt, based on DESY program
SixTrack simulates 60 particles for 100k-1M LHC orbits
Can include measured magnet parameters, beam-beam interactions
LHC@home revealed reproducibility issues, solved by E. McIntosh

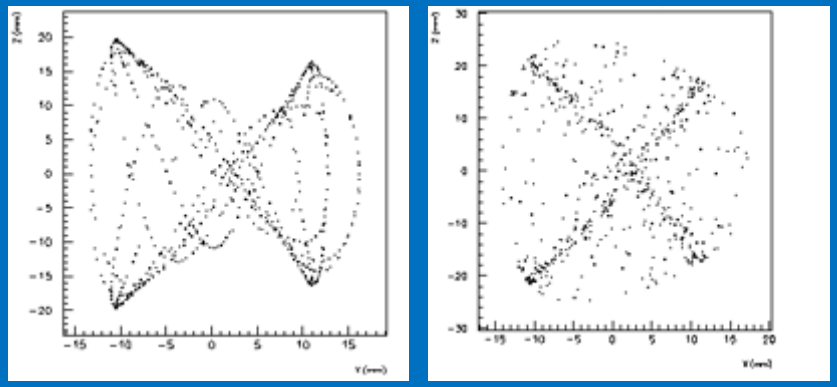

Phase space images of a particle for a stable orbit (left)
and unstable chaotic orbit (right).

>3000 CPU-years
>60k volunteers

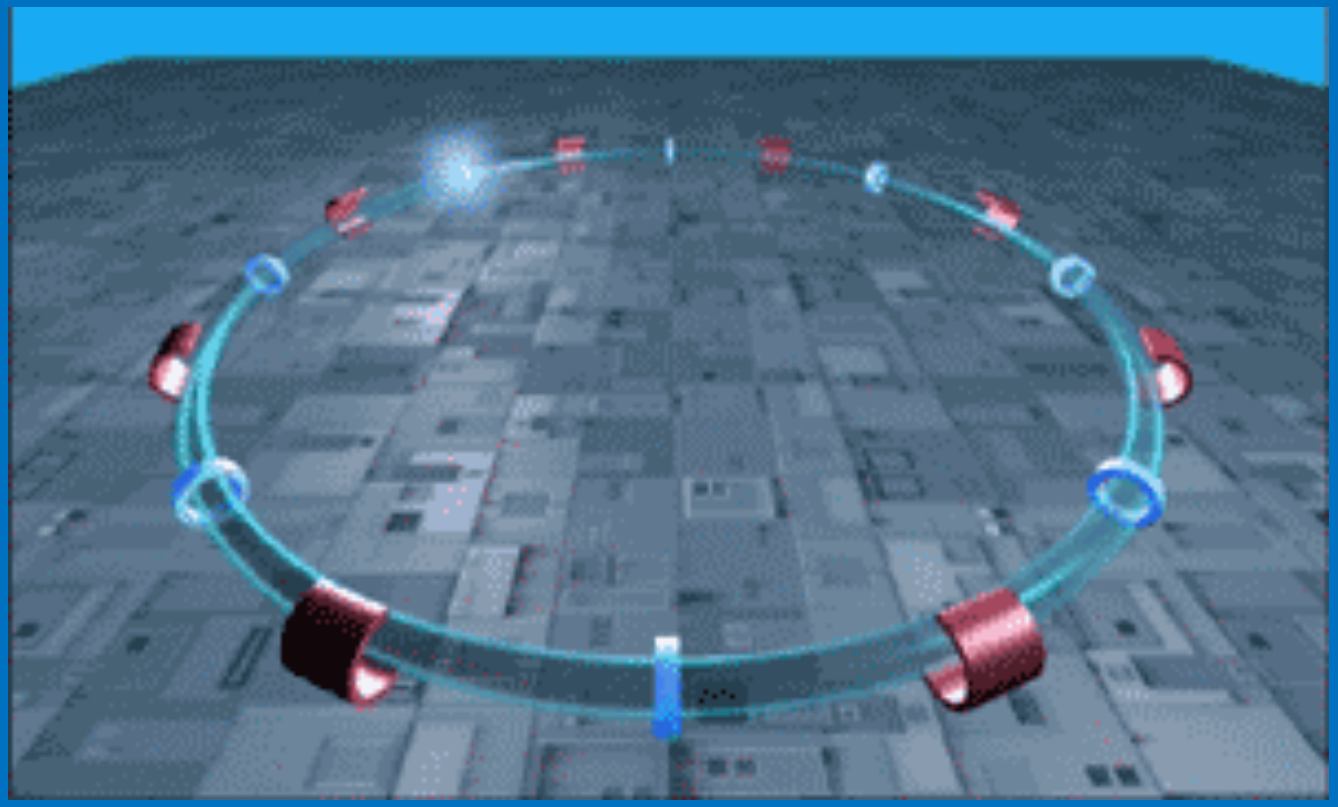

BOINC for real LHC physics: *LHC@home 2.0*

Challenge was issued in 2006:

"Why don't you run real LHC physics on BOINC?"

- **After 4 years, with students and volunteers, we did it !!**

- **Allow "any" PC to run a full LHC physics application**
- **Let all those PC's look like a "standard Data Centre"**

Monte Carlo event generation

A lab for testing theory models and detector performance with high statistics

The MC description of LHC events is tremendously complex

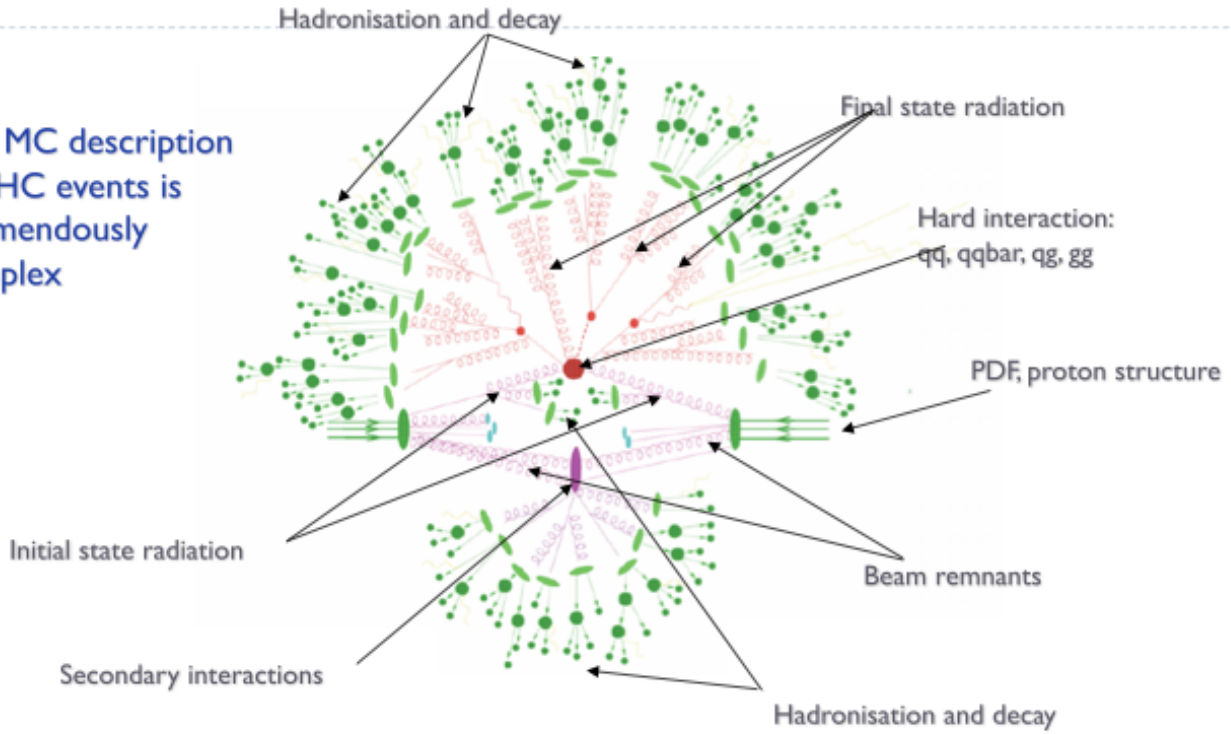

This is a schematization to be able to cut down the problem in pieces and model them in a different way. The "pieces" are correlated !

BOINC for real LHC physics

Technical Challenges

- **Using normal BOINC: each application must be ported to every volunteer's PC (mostly Windows). But the LHC experiments run Linux. Porting to Windows is impractical.**
- **Experiment code changes often: all executing PC's must be updated. And the code size is VERY BIG (10 GBytes).**
- **Experiments have their own job management systems and do not want to use BOINC job scheduling.**
- **Volunteer PC's are an "untrusted" resource.**

CernVM + Co-Pilot

- Using **Virtualization** and **CernVM** it is possible to solve the porting problems for application deployment on volunteers' machines

- With CernVM's **Co-Pilot** it is possible to connect existing Grid infrastructures of LHC experiments with the BOINC volunteer resources, removing the need for physicists to change their procedures

CernVM
SoFTware Appliance

BOINC – CernVM Architecture

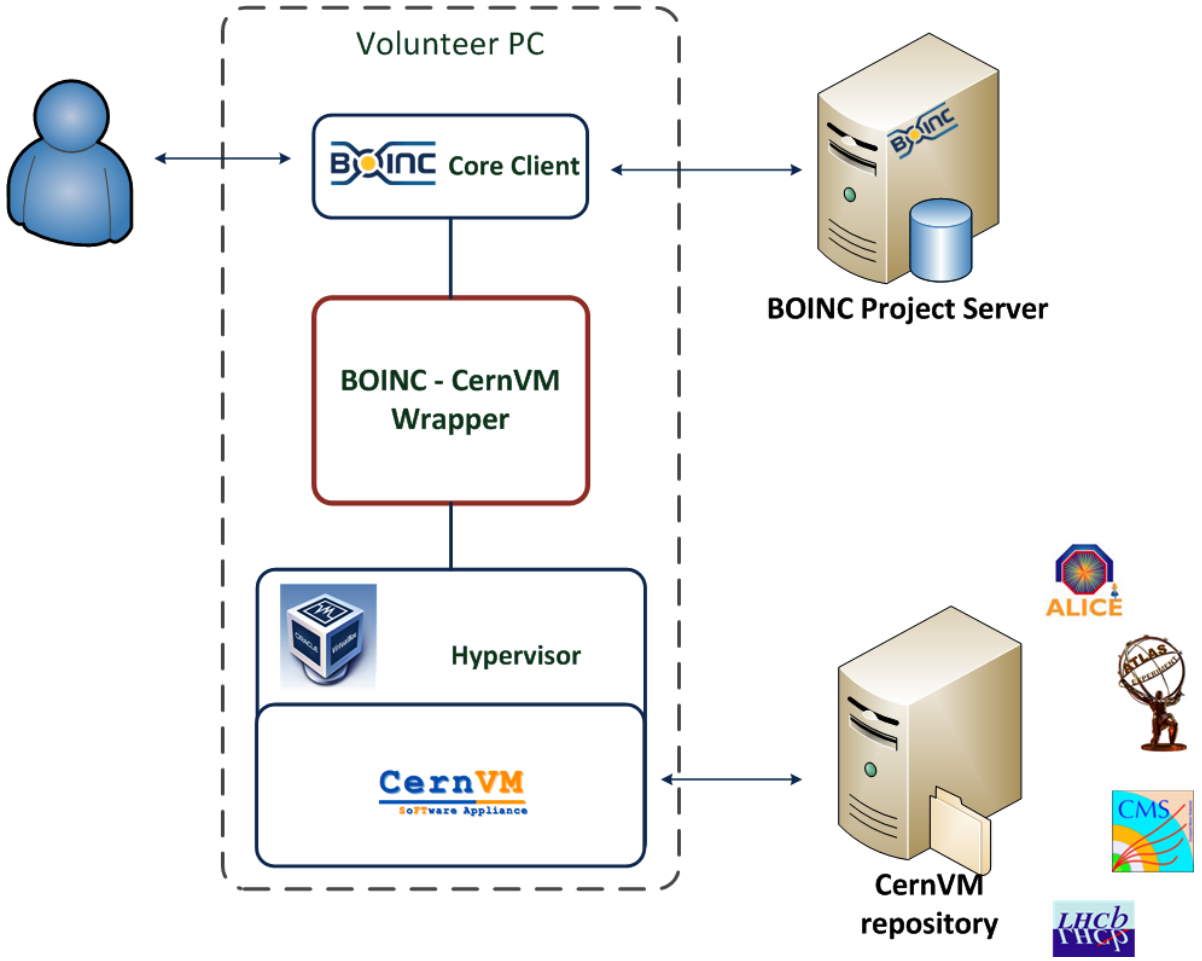

BOINC + CernVM + Co-Pilot => Grid (Volunteer Cloud)

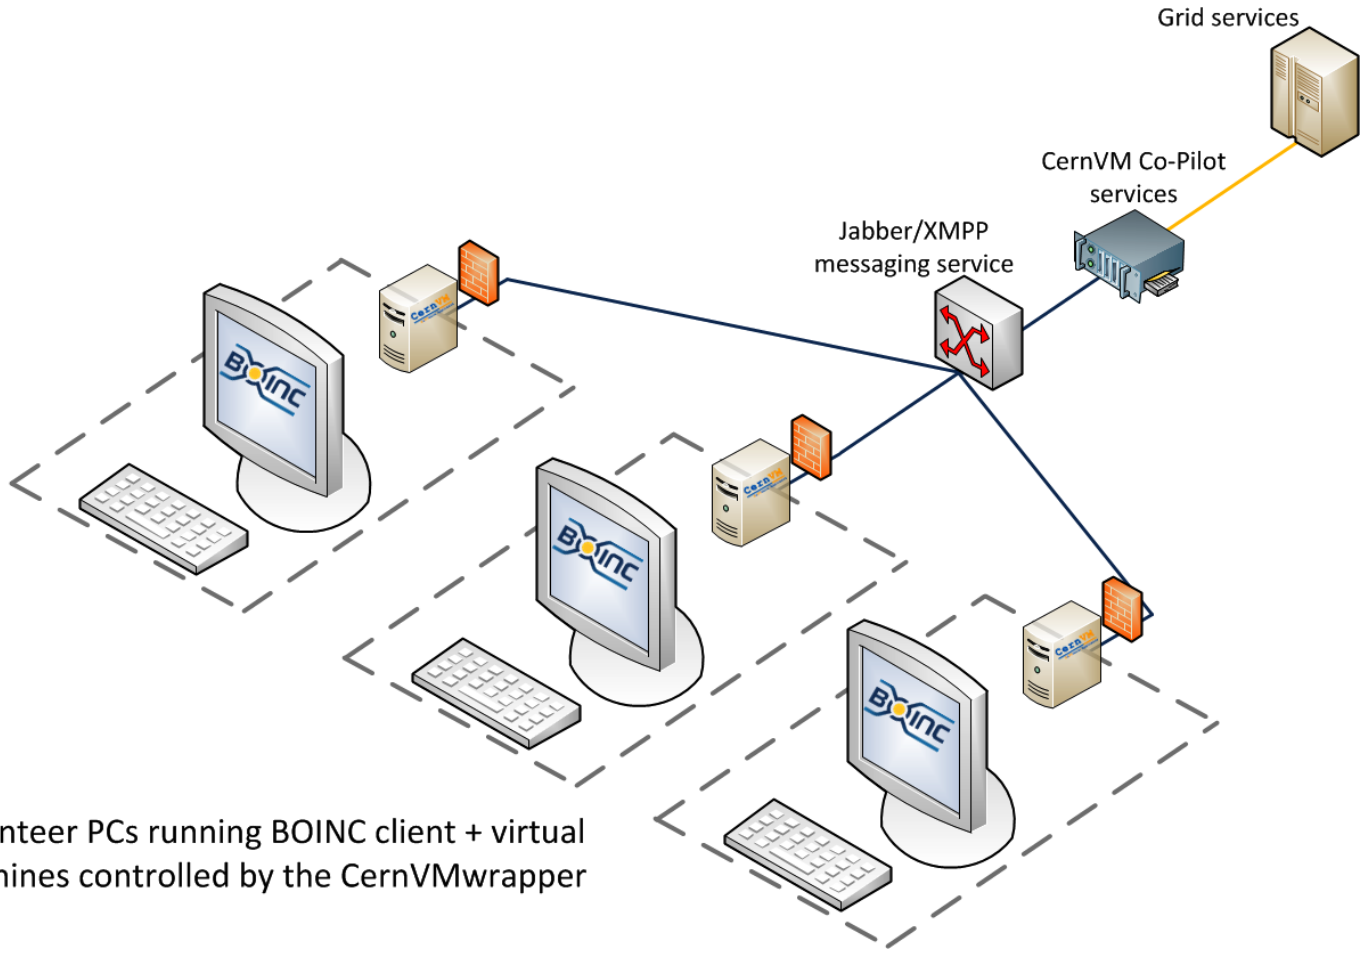

Project Contributors (since 2006)

The CernVM project team / CERN-PH:

Predrag Buncic, Jakob Blomer, Carlos Aguado, Artem Harutyunyan

STUDENTS (usually 2 months) - supervisor Ben Segal:

Helen McGlone (UK - 2006)
Daniel Alvarez (Spain - 2006)
Daniel Lombrana (Spain - 2006)
David Weir (UK - 2007)
Ilya Gorbunov (Russia - 2008)
Jarno Rantala (Finland - 2009)
Arturo Sanchez (Venezuela - 2009)
Rohit Yadav (India - 2010)
Wenjing Wu (China - 2010)
Jie Wu (China - 2010)

VOLUNTEER

David Garcia / CERN-BE (2009-2010)

The LHC Volunteer Cloud

Final Summary:

- Solved porting problem to all platforms: Windows, Linux, Mac
- Solved image size and image updating problems
- Solved job production interface problem
- Solved problem of untrusted clients

- **All done without changing the existing BOINC infrastructure or any LHC physicists' procedures**

- ***We have built an LHC "Volunteer Cloud" …***

Future Plans for LHC@home 2.0

- Go on from current "alpha" testing phase (using about 150 volunteers) to "beta" testing (about 1000 volunteers), still running the Theory Monte Carlo application
- Start running an LHC experiment application, e.g. for ATLAS or CMS
- Add features to BOINC client to show more application details while physics jobs are running
- Develop a "Volunteer Thinking" version of the Theory application (currently in a very early prototype form)

Africa@home: Malaria Control

- 3 month student project (in 2005)
- Malaria epidemiology project (STI)
- 6000 CPUs joined in 2 weeks
- >100 CPU-years in 2 months
- Demoed at WSIS Tunis 2005
- Went public July 2006
- Workshop in S. Africa in July 2007

MalariaControl.net:
modeling the spread of the disease

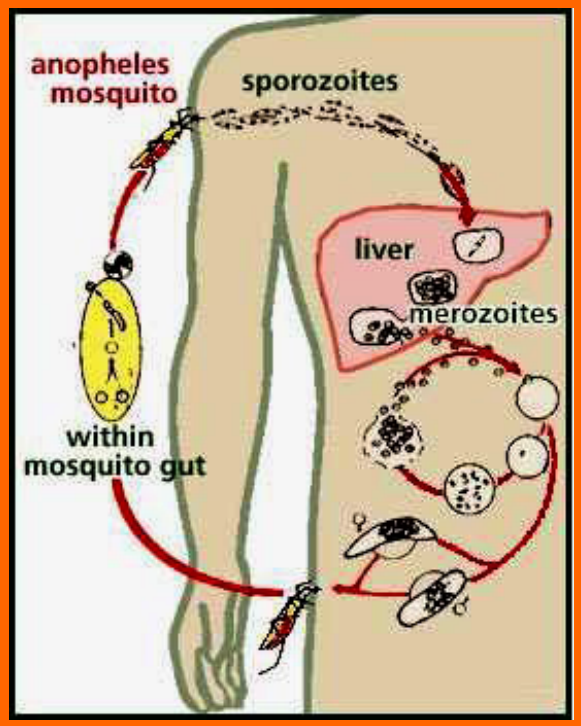

STI population model:

Individuals (humans, mosquitoes) with properties (age, immunity…) and interactions (infect…)

Evolve model, observe results, adjust parameters, repeat to fit field data (deaths, $ spent…)

MalariaControl.net: scientific results

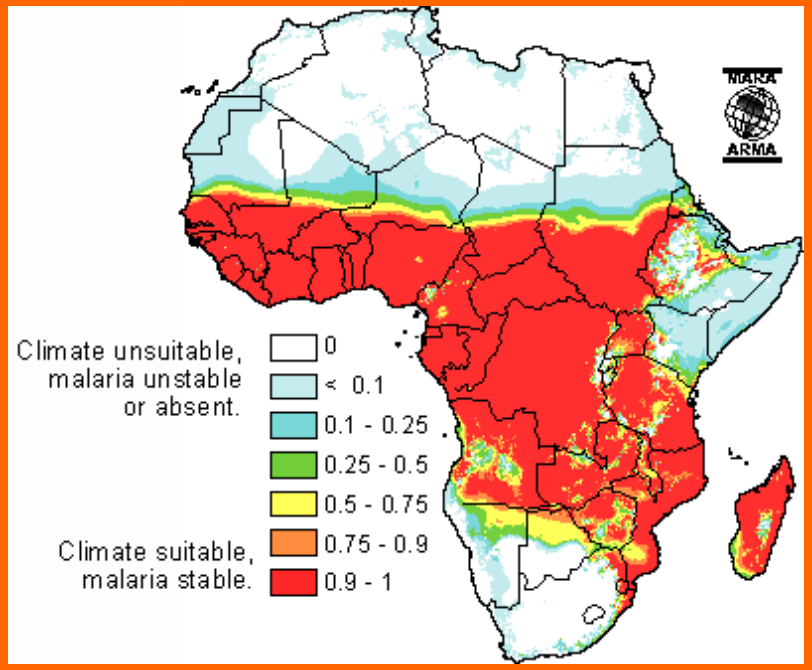

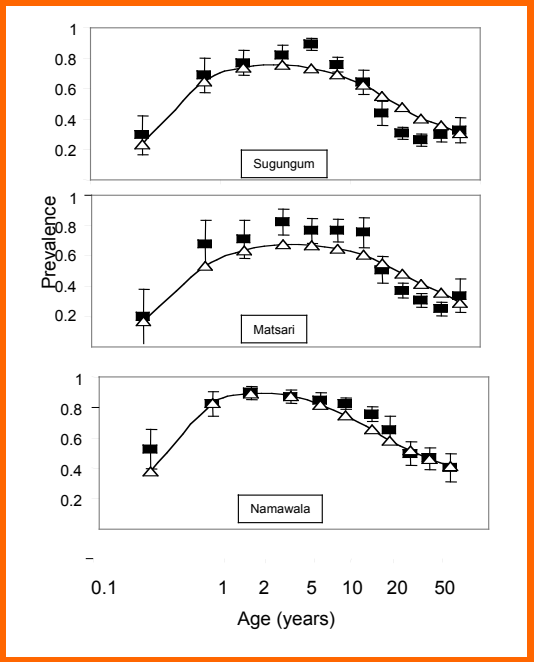

14 articles on STI model published in Am. J. Trop. Med. Hyg., 75 (supplement), 2006.

Volunteer computing enables detailed models, more parameters, projecting future scenarios.

MalariaControl.net: health impact

STI model predictions of cost-effectiveness: Vaccine **$1 - $10** per dose with 52% efficacy
= $4.73 - $34.43 per fully-immunized child **= $450 - $3,500** per death averted
= $12 - $96 per disability adjusted life year **= $2.7M - $19.8M** per year for Mozambique

- STI data reviewed by PATH Malaria Vaccine Initiative with stakeholders in Mozambique.
- Instrumental in Mozambique now planning for possible future use of a malaria vaccine.

Africa@home:
empower African scientists

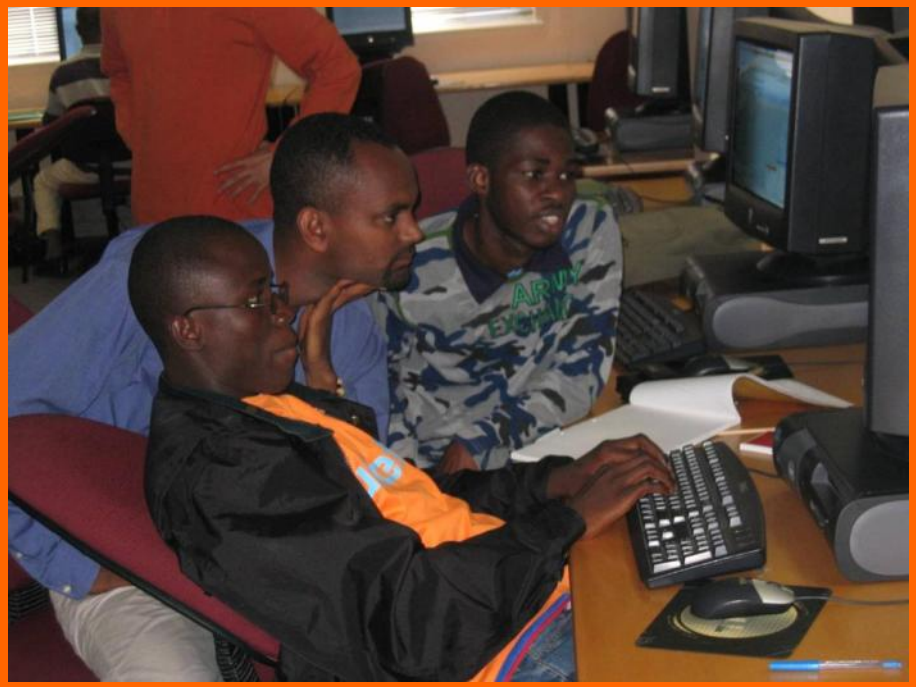

Partnership: CERN, Swiss Tropical Institute, Uni. Geneva, World Health Org, 2 NGOs
Africa@home workshops: >50 African scientists from 20 countries (South Africa, Mali)
Africa@home projects: MalariaControl.net, HIVMM, AfricaMap, Autodock (w. HealthGrid)
Africa@home servers: Uni. Cape Town, Uni. Geneva