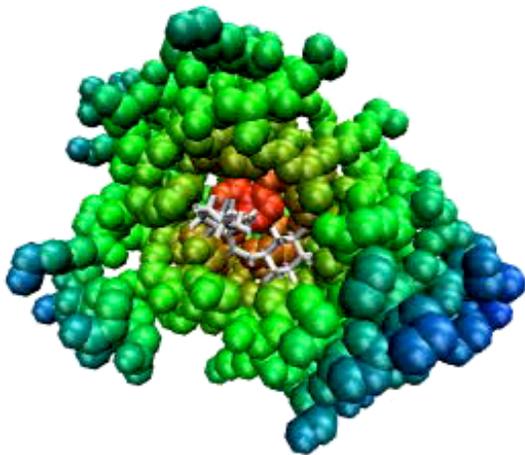


Distributed
computing
map
←

Examining Biomolecular Phenomena with Massively Parallel Computation



Guha Jayachandran, Ph.D.
GuhaJ@cs.stanford.edu
June 20, 2007

Outline

- Some history, the goal and the problem
- Folding@Home: infrastructure, technologies, and participation
- Example parallel techniques, with results
- The future





History, Goal, and Problem

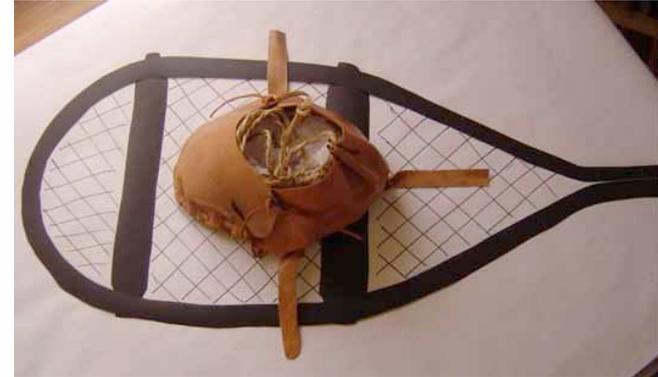


A Corpse in the Alps



Why interesting?

His Possessions



Biomedicine Not New

- Traditional Chinese medicine and Ayurveda both several thousand years old
 - Many compounds now being studied

- Aspirin's chemical forefather known to Hippocrates



- Even inoculation at least 2000 years old



- Of course, many useless drugs too

Revolutions

Computational



Biomedical

The Molecular Cast

“...all things are made of atoms, and that everything that living things do can be understood in terms of **the jiggings and wiggings of atoms.**”

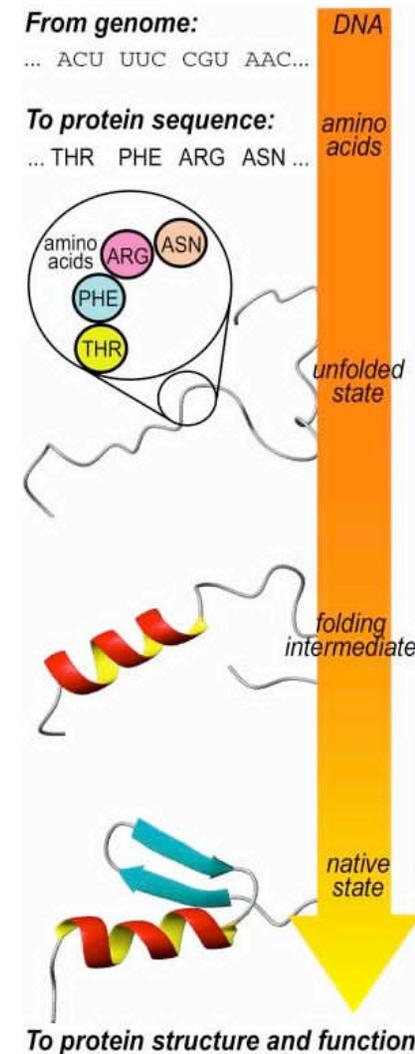
- Richard Feynman



The Molecular Cast

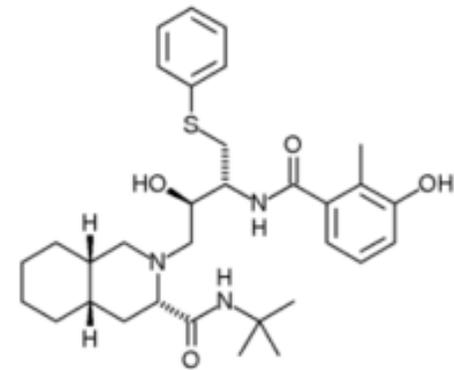
■ Protein

- Carries out the functions of life at the cellular level
- Paradigm for nanotechnology
- Folds into a given 3D structure
 - Sequence → structure → function
 - Misfolding can cause disease



The Molecular Cast

- Small molecule ligand
 - Low molecular weight organic molecule
 - Can modulate proteins' activity by binding to them
 - Many drugs are small molecules that work by binding to a given protein
 - Probability of a ligand and protein being bound is a function of their binding free energy



Nelfinavir: Binds to HIV-1 and HIV-2 proteases, inhibiting them from cleaving viral protein. Used in treatment of AIDS.

Grand Challenges in Biophysics

Biophysics	Biocomputation
<ul style="list-style-type: none">■ Protein structure and dynamics<ul style="list-style-type: none">■ Explain structure	<p>Predict folded structure given sequence</p>
<ul style="list-style-type: none">■ Elucidate folding mechanisms	<p>Simulate folding from unfolded state... in detailed model, to long time</p>
<ul style="list-style-type: none">■ Protein-ligand binding<ul style="list-style-type: none">■ Say how well a given ligand will bind to a protein	<p>Accurately compute solvation and binding free energies</p>

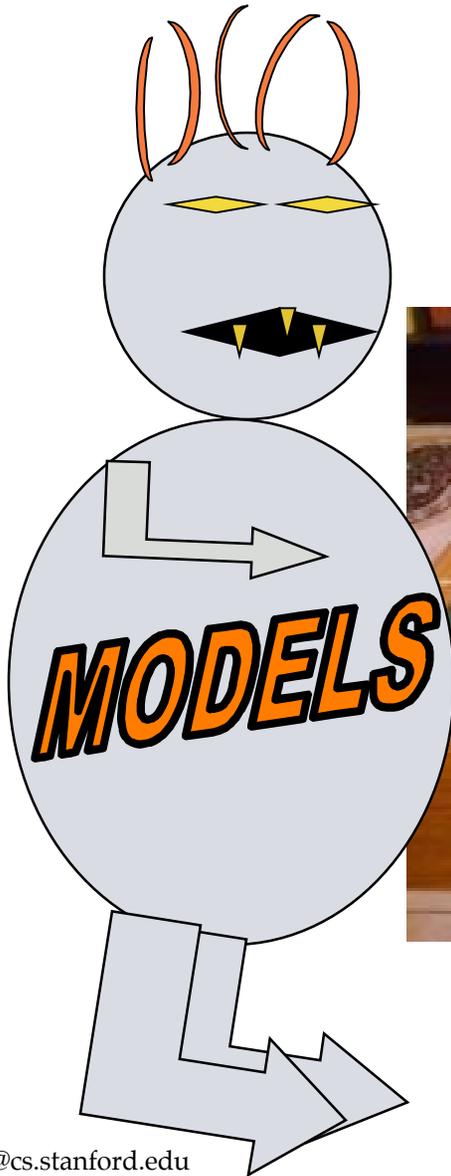
Virtual Airplanes

Why compute? To make possible the otherwise impossible.

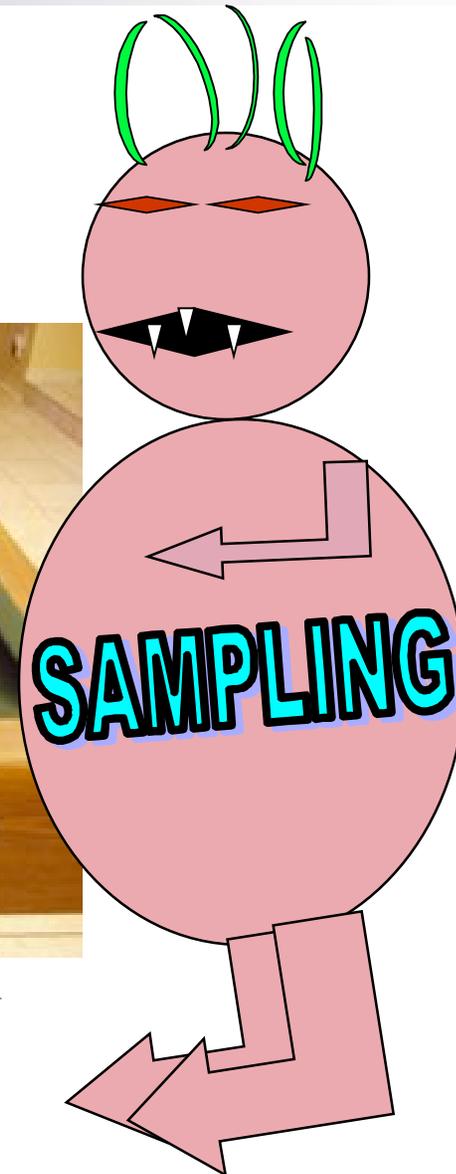
Where does that power come from? Using understanding instead of physical trial and error.



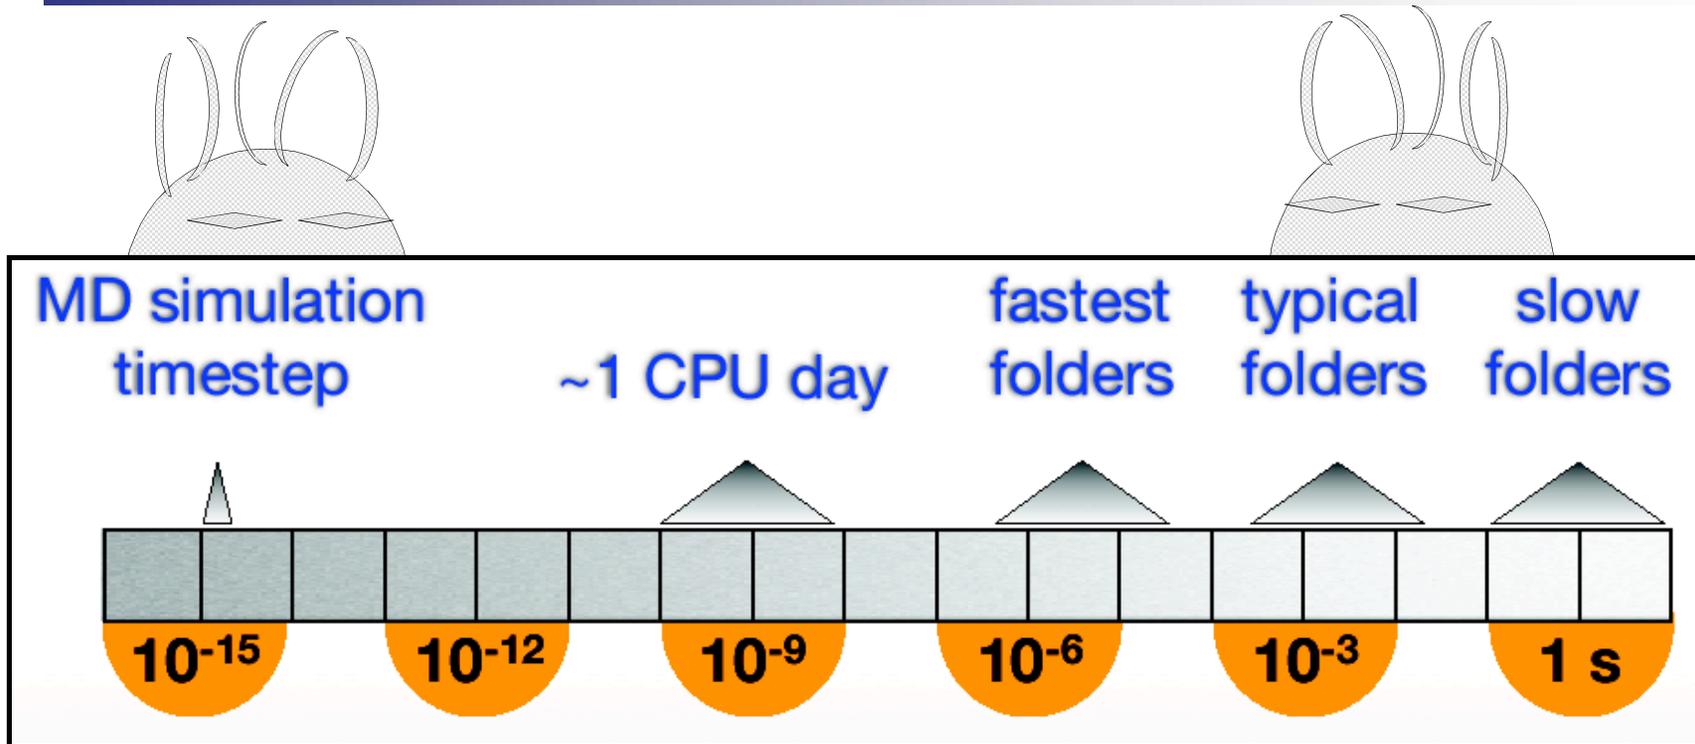
Difficulties in Simulation



Molecular dynamics is based on integrating Newton's laws of motion for each atom



Difficulties in Simulation



*If you could simulate 1 ns/day,
then a 1 ms simulation would take
3000 years!*

Difficulties in Simulation

Need:

**More computational
power and better
computational
techniques**

MD simulation
timestep



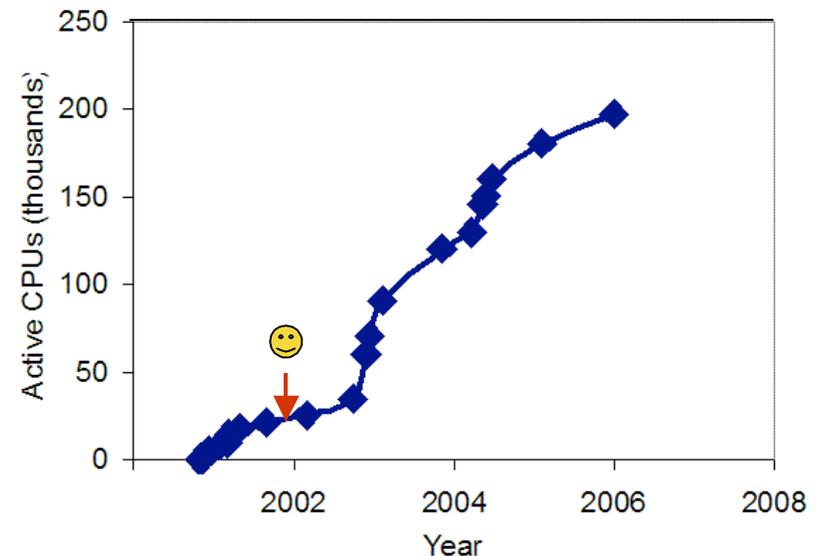
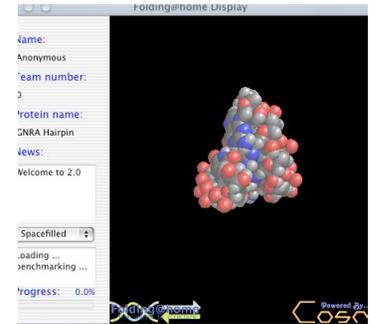
slow
folders



Folding@Home

Project Summary

- Distributed computing project (<http://folding.stanford.edu>): anyone can run our software (Linux/Windows/Macintosh)
- Clients run assigned simulations with unused processor cycles
- Steady growth in install base
- Current strength:
 - >250,000 active clients
 - ~800 TFLOPS sustained



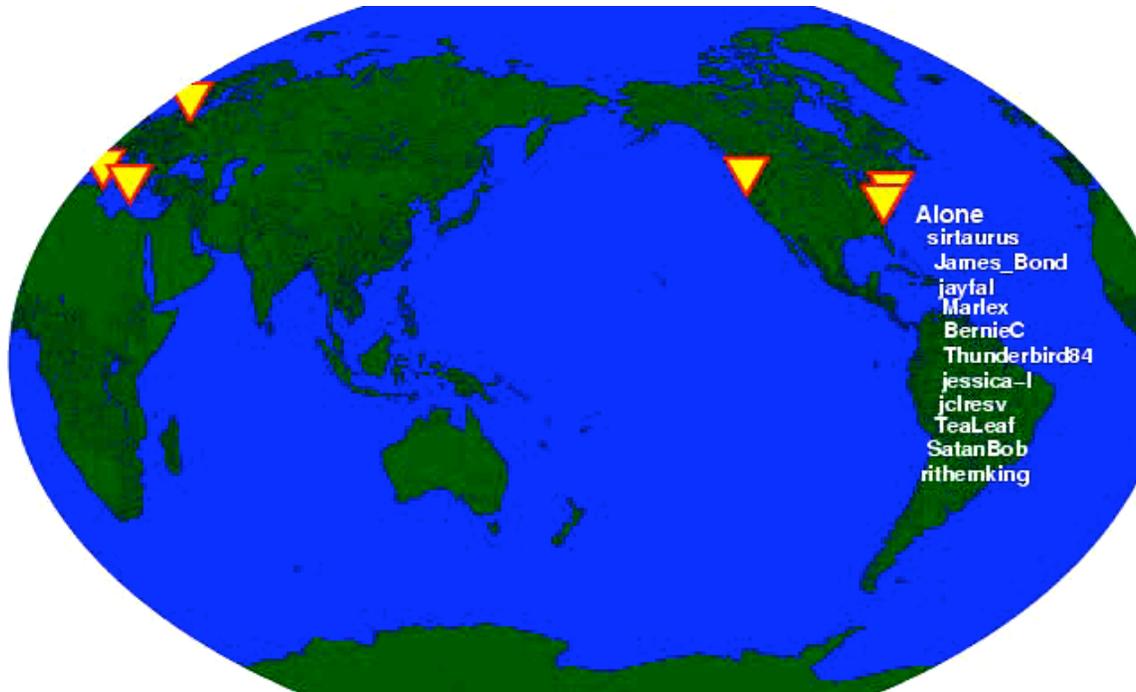
Global Client/Server Architecture

Server side



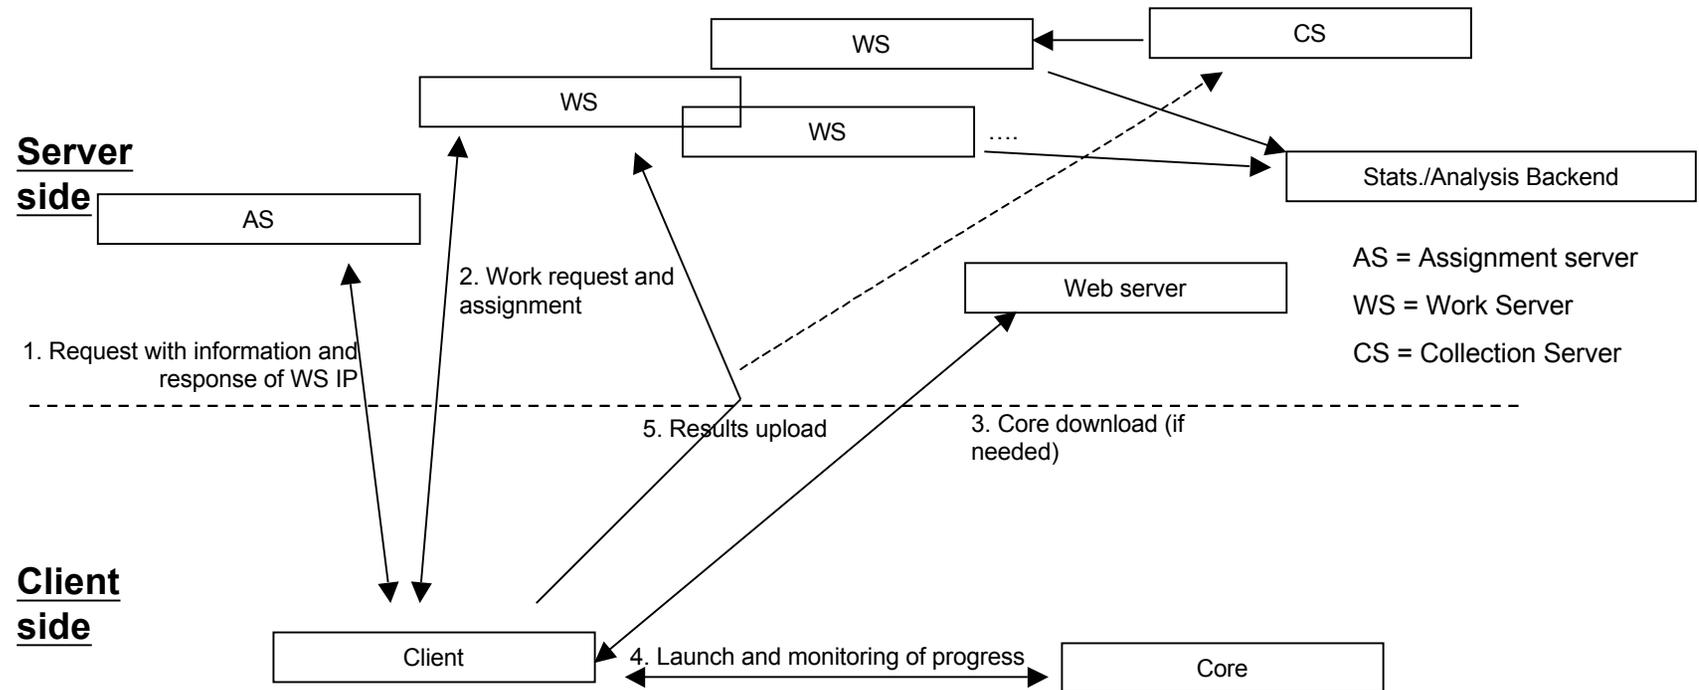
Servers are located in 4 campus buildings and are of varying ages, powers.

Client side

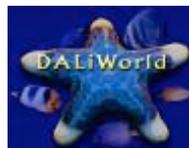
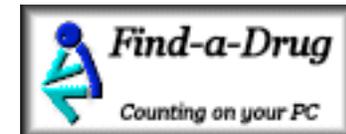


Client computers are in 6 continents and upgrade on their own.

Global Client/Server Architecture



Other Distributed Projects



CycleTraders



yeti @ home

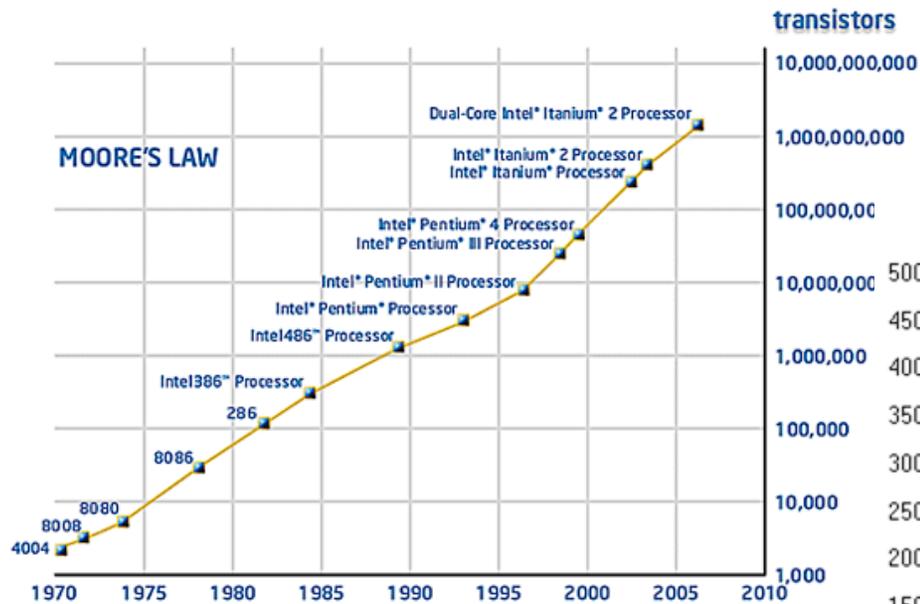
THE GREAT HAIRY DEAL

“there's a small, but
captivating possibility
that **your computer**
will detect the
faint footsteps of a
hairy monster...”

SIC@HOME

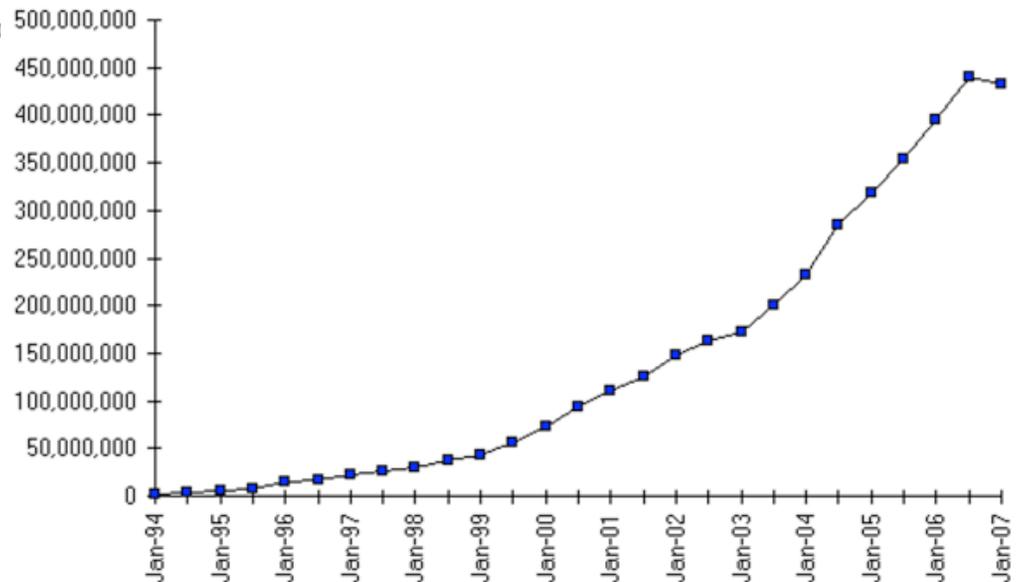
“The SIC project searches for just such incredible coincidences. If we can find such a coincidence, we can prove that the universe is truly infinite!”

Moore's Law + Internet Growth



From www.intel.com

Internet Domain Survey Host Count



Source: Internet Systems Consortium (www.isc.org)

Participants

Online statistics let participants track production and compete for top rankings. People also form teams and online community.

Donor Statistics Last updated: Monday December

Number	Name	Credit	Total	Te
Top 100 users				
1	anonymous	108926492	1224577	0
2	eastms.edu	44979648	407172	17
3	OC-AMD	43431574	356322	30
4	I920xI080i-How???	41185582	228140	64
5	FLECOM	20326537	120354	33
6	Silicon_Wars_-_a_New_Hope	18200167	133777	24
7	Scott_H	16835547	137364	26
8	Christian_Bargmann	12145552	73996	19
9	nikhsbl	12114344	129562	32
10	Sceggs.nsw.edu.au	10371612	80322	24
11	newtownhigh	9406884	79979	37
12	KWSN_MillenniumKnight	9090336	81014	11
13	ulf.van.laak	8641758	38310	28
14	larrybpsu	8443546	35575	39
15	ITAP_RCAC	8060743	47344	15
16	[LEENFAT]dapo	7692747	39915	51
17	marty9876	7664570	38040	33
18	wayne	7649388	36977	46
19	Wedo	7637471	131927	32
20	KWSN_Sir_CADCAM	7509438	73622	11
21	belloq	7293691	61165	18
22	University_of_Alabama_Computer_Science	7257739	71821	10
23	DorXtar_(Team_Rage3D)	7233542	62773	64
24	Netsphere	7144986	88572	32
25	Bollocks	6869277	52238	24
26	ChasR	6622742	26733	32
27	Toasty[H]ardOCP	6363649	50035	33
28	fcapes	6315845	72204	4
29	Bhurak_AU	6313517	39477	24
30	Mustanley	6223423	58366	32
31	Trinedy	6094696	38303	24
32	Robert_Ameeti	6065403	32253	33

Rankings screenshot.



Some team logos.



Folding-community.org Forum Index

The time now is Mon 25 Dec, 2006 20:42

Folding-community.org Forum Index

Board Announcements	Replies	Author
Announcement Please ask to be added to the Native-FAH beta team here. Description : This is for requests to the [Native] FAH Beta Team only... [Folding Community » General non-FAH]	7	ww

Pandegroup	Topics	Posts	Last Post
Pandegroup Pandegroup, the team lead by Vijay Pande, visits here :)	2962	23204	Not getting credit Today at 20:20 JDifuria
Pandegroup News News by Pandegroup about the project(s). Moderators can start topics and polls, registered users can reply and vote. Moderators: Pandegroup, Site Moderator	272	3361	SMP client now beta test... Yesterday at 10:44 codysluder
Data Server Status Please post here if you have trouble with a particular data server, and if a problem has been resolved. Please check out http://vspx27.stanford.edu/serverstat.html as well ;) Also, general "can't upload" or "can't download" questions go here. Moderators: Pandegroup, Site Moderator	964	9286	Missing credits Sat 23 Dec, 2006 01:36 Bruce
Stats/Web Server Status			

User run support Forum.

Partial Change Log

5.04:

- Allows configuration of the amount of memory to indicate available, especially useful when running multiple "big unit" clients on a single machine.
- Allows for the setting of -advmethods as a persistent option through the Configuration.

5.02:

- Error checking on queue, related to big packet option.
- OpenBSD branding fix. (5.01)

5.00:

- Collection server support. If a client finishes a work unit from a Work server with this feature enabled, then if 2 upload attempts to that server fail, the client will return to the work unit to a Collection server. Pending credit will be given for the unit, becoming final when the Work server comes back up and verifies what the Collection server has received. All work servers will eventually be associated with a Collection server, and many already are.
- Configuration option to allow downloads and uploads greater than 5 MB in size, which some new projects will require due to the size of what they are simulating.
- G@H configuration option reworded to be "No deadline" option.
- Print executable name and launch directory to log.
- Fix for bug where queue slot could get stuck "Fetching"
- MyFolding links updated.
- -forceSSE flag removed, since even with older clients, the flag had become meaningless because recent Core_78's give SSE priority over 3DNow automatically.
- OpenBSD support: If give flag -openBSD, client will automatically make elf2of system call on downloaded Cores.
- Win. only: Configuration option to pause work when the computer is using battery power (for laptops).
- Win. console: Built in service installation support through Configuration, to have the program automatically begin at startup.

4.00:

- Configurable minimum checkpointing frequency (3-30 minutes).
- If client run with verbosity of at least 5, FahCore_78 will print out a message whenever a timered checkpoint occurs.
- -configonly flag added to configure then exit.
- -oneunit flag added to exit after completing one unit.
- -forceSSE flag added.
- Warnings printed if -forceasm/-forceSSE flags given.
- ~~GitHub~~ [GitHub](https://github.com/Stanford.edu) at start of normal operation (for example, not if only configuring, listing queue, deleting unit, etc.).

- Warning if attempt to run client with Machine ID of already running client
- Web site printed updated.
- Upon finishing a unit, if fails to return the result but succeeds in getting new work, immediately tries sending all results again.
- Checks user name and team number for legality of use in statistics.
- Queue info properly outputted to FAHlog.txt.
- "Ask before connecting" change on Linux -- ignores typeahead and asks for "C"
- Proxy password written out in binary (upon reconfiguration)
- Assignment/work server request authentication
- Send all message altered ("sent x of y" if x not 0, or "no units...")
- CPU usage default of 100 printed out during Configuration.
- Tags in MyFolding page changed to better adhere to HTML standard
- Maximum machine IDs raised to 8.
- Upload and download bandwidths estimated to send to AS.
- If fails to post data, includes server and port in message. Also outputs HTTP status code (before this was only done if I.E. settings used).
- Bug fixed where performance fraction could exceed 1.
- Batch networking option for G@h (Network when less than 2 units ready to work on or more than 5 results to return, and when network, both get up to 10 units and send any results).
- If set G@h preference, reflected in acknowledgment message.

3.24 (Dec. 23, 2002):

- Benchmarking improved in consistency and uniformity.
- Support for sending a measure to the assignment server of time taken on past work units (a weighted average). Can view this measure using verbosity 4 after finished units.
- If bad arguments given, will exit only after a key press.
- If upon startup, FAHlog.txt exceeds 50Kb and a is to be started, the old one is retained under the name FAHlog-Prev.txt.
- Deadline information included with work (Printed to unitinfo.txt and available in queueinfo). Active units checked for expiration periodically all others checked for expiration at autosend time. Configuration option of ignoring deadline information (to use if system clock is bad).
- MyFolding.html page written out by client with useful links (which Linux and Win. console users did not have easy access to before)
- If specify -freeBSD flag, downloaded cores will automatically be branded as Linux (should make usage seamless on FreeBSD platform).
- In making request to server for work, inform on whether have units done waiting to be uploaded to server. If so, servers running new code will give a new unit instead of giving same.
- If a work unit is rejected for any reason, server now returns reason to client, which prints it out. If WU fatally flawed, deleted from queue.
- Checks for bad URL, port on reception of work unit, deleting it before even starting work if bad.

- "Could not transmit unit" statement now also includes date unit was completed ("Could not transmit unit x (Completed xxxx)").
- Time stamp on every line of FAHlog.
- Arguments to client printed to log.
- Some log messages modified (like the requested, and deserved, "Thank you" added!).
- File "unitinfo.txt" printed and updated, with information and progress for current work unit. (Had been requested for use by 3rd party people.)
- Answers to Configuration questions now validated. One letter answers (i.e. "y" or "Y" instead of "yes") allowed.

3.14:

- fixed bug in configuration of core priority.
- -forceasm command line option supported to force assembly optimizations (could be offed again by core if problems) IF current work unit is a Core_78 WU.
- If fails to connect to assignment server, client will no longer just go back to previous work server, since that server may no longer be appropriate for this client.
- Fixed message where if run "-send all," user was incorrectly told that failed.
- If FAHlog.txt is locked by OS, FAHlog2.txt started.
- -advmethods option added for users wishing to try the newest cores and/or WU's
- Advanced Configuration option of disabling optimized assembly code added.
- Send/autosend conflict resolved.
- If CORE_OUTDATED situation arises, and remains even after a new core is downloaded, the current work unit is deleted.

3.12:

- Rather than in the /tmp directory, the User ID is stored in the same directory as the client, in a file called machinedependent.dat.

3.11:

- If Client-core communications error occurs, current work unit and work files deleted.
- If 5 consecutive cores that have been downloaded during the current execution of FAH fail to complete a work unit, the client will sleep for 1 day. When it wakes up, it give one more downloaded core a chance to run and complete a WU before sleeping again.
- FAHlog.txt is appended to (until it reaches 50 Kb) instead of overwritten.

3.1:

- Preference of project added to Configuration's Advanced tab.

Partial Change Log

5.04:

- Allows configuration of the amount of memory to indicate available, especially useful when running multiple "big unit" clients on a single machine.
- Allows for the setting of -advmethods as a persistent option through the Configuration.

5.02:

- Error checking on queue, related to big packet option.
- OpenBSD branding fix. (5.01)

5.00:

- Collection server support. If a client finishes a work unit from a work server with this feature enabled, then if 2 upload attempts fail, the client will return to a Collection server. Pending credits will be given and verifies what the Collection server has received. All work units that fail to upload will be returned to the work server. If a client finishes a work unit and 2 upload attempts fail, the client will return to the work unit to a Collection server.

- G@H configuration option reworded to be "No deadline" option.
- Print executable name and launch directory to log.
- Fix for bug where queue slot could get stuck "Fetching"
- MyFolding links updated.
- -forceSSE flag removed, since even with older clients, the flag had become meaningless because recent Core_78's give SSE priority over 3DNow automatically.
- OpenBSD support: If give flag -openBSD, client will automatically make elf2of system call on downloaded Cores.
- Win. only: Configuration option to pause work when the computer is using battery power (for laptops).
- Win. console: Built in service installation support through Configuration, to have the program automatically begin at startup.

4.00:

- Configurable minimum checkpointing frequency (3-30 minutes). Core_78 will print out a message whenever a timed checkpoint occurs.
- -oneunit flag added to exit after completing one unit.
- -forceSSE flag added.
- Warnings printed if -forceasm/-forceSSE flags given.
- GuhaJ@cs.stanford.edu start of normal operation (for example, not if only configuring, listing queue, deleting unit, etc.).

- Warning if attempt to run client with Machine ID of already running client.

Assignment/work server request authentication

- Upon finishing a unit, if fails to return the result but succeeds in getting the result, results again.
- Checks user name and team number for legality of use in statistics.
- Queue info properly outputted to FAHlog.txt.
- "Ask before connecting" change on Linux -- ignores typeahead and asks for "C"
- Proxy password written out in binary (upon reconfiguration)
- Assignment/work server
- Send all message altered ("sent x of y" if x not 0, or "no units...")
- CPU usage default
- Tags in MyFolding page changed to better adhere to HTML standard
- Maximum machine
- Upload and download bandwidths estimated to send to AS.
- data
- HTTP status code (before this was only done if I.E. settings used).
- fixed where performance fraction could exceed 1.

- Batch networking option for G@h (Network when less than 2 units to work on or more than 5 results to return, and when network, both get up to 10 units and send any results).
- If set G@h

Support for sending a measure to the assignment server of time taken on past work units

- 3.24 (Dec. Benchmarking improved in consistency and uniformity.
- Support for assignment server of time taken on past work units (a weighted average). Can view this measure using verbosity 4 after finished units.
- If bad arguments given, will exit only after a key press.
- If the old one is retained under the name FAHlog_Prev.txt.
- and available in queueinfo). Active units checked for expiration periodically all others checked for expiration at autoseed time. Configuration option of ignoring deadline information (to use if system clock is bad).
- MyFolding.html page written out by client with useful links (which Linux and Win.

- If specify -freeBSD flag, downloaded cores will automatically be branded as Linux (should not be necessary on BSD systems).
- In making request for server, will download several work units, doing waiting to be uploaded.
- If a work unit is rejected, server now returns specific reason to client, which prints it out. If the unit is fatally flawed, then it's deleted from the client queue.
- Checks for bad URL, port on reception of work unit, deleting it before even starting work if bad.

- "Could not transmit unit" statement now also includes date unit was completed ("Could not transmit unit x (Completed xxxx)").
- Time stamp on every line of FAHlog.

Provide unitinfo and MyFolding pages

- Answers to client printed to log.
- Some messages requested, and deserved, "Thank you" added.
- File, with information and progress for current work unit. (Had been requested for use by 3rd party people.)
- Answers to Configuration questions now validated. One letter answers (i.e. "y" or "Y" instead of "yes") allowed.

Batch networking option. Network

when less than 2 units ready to work

- 3.14: configuration of core priority.
- -forceasm core and line option supported to force assembly to be offered again by core if problems) IF current work unit is a Core_78 WU.

- If fails to connect to assignment server, client will no longer just go back to previous work server, since that server may no longer be appropriate for this client.

- Fixed message where if run "-send all," user was incorrectly told that failed.

- If FAHlog.txt is locked, -advmethods option added for users wishing to try the newest cores and/or WU's

- Advanced Configuration option of disabling optimized assembly code added.
- Send/autosend conflict resolved.

- If CORE_OUTDATA and remains even after a new core is downloaded, the current work unit is deleted.

-advmethods option added for users wishing to try the newest cores and/or work units

- 3.12: Rather than in the /tmp directory, the User ID is stored in the same directory as the client, in a file called machinedependent.dat.

3.11:

- If Client-core communications error occurs, current work unit and work files deleted.

- If a work unit is rejected, server now returns specific reason to client, which prints it out. If the unit is fatally flawed, then it's deleted from the client queue.

- If Client-core communications error occurs, current work unit and work files deleted.
- If a work unit is rejected, server now returns specific reason to client, which prints it out. If the unit is fatally flawed, then it's deleted from the client queue.

3.11:

- Preference of project added to Configuration's Advanced tab.

Scientific Results

Local structure formation in simulations of two small proteins

Guha Jayachandran, V. Vishal, Angel E. Garcia and V. S. Pande. Journal of Structural Biology (2006)

Kinetic Definition of Protein Folding Transition State Ensembles and Reaction Coordinates
C. Snow and V. S. Pande. Biophysical Journal (2006)

Parallelized Over Parts Comp. of Absolute Binding Free Energy with Docking and Mol. Dynamics

Guha Jayachandran, M. R. Shirts, S. Park, and V. S. Pande. Journal of Chemical Physics (2006)

Folding Simulations of the Villin Headpiece in All-Atom Detail

Guha Jayachandran, V. Vishal, and V. S. Pande. Journal of Chemical Physics (2006)

Ensemble molecular dynamics yields submillisecond kinetics and intermediates of membrane fusion
P. Kasson, N. Kelley, N. Singhal, M. Vrljic, A. Brunger, and V. S. Pande. PNAS, USA

Electric Fields at the Active Site of an Enzyme: Direct Comparison of Experiment with Theory
Ian T. Suydam, Christopher D. Snow, Vijay S. Pande, Steven G. Boxer. Science (2006)

A novel approach for computational alanine scanning: application to the p53 oligomerization domain
L.T. Chong, W. C. Swope, J. W. Pitera, and V. S. Pande. Journal of Molecular Biology (2006)

On the role of chemical detail in simulating protein folding kinetics
Young Min Rhee and Vijay S. Pande. Chemical Physics (2006)

Nanotube confinement denatures protein helices
Eric J. Sorin and Vijay S. Pande. JACS (2006)

The solvation interface is a determining factor in peptide conformational preferences
Eric J. Sorin, Young Min Rhee, Michael R. Shirts, and Vijay S. Pande. J. Molecular Biology (2006)

How large is alpha-helix in sol.? Studies of the radii of gyr. of helical peptides by SAXS and M.D.

Bojan Zagrovic, Guha Jayachandran, I. S. Millett, S. Doniach and Vijay S. Pande. J. Mol. Biology (2005)

Error Analysis in Markovian State Models for protein folding
Nina Singhal and Vijay S. Pande. Journal of Chemical Physics (2005)

A New Set of Molecular Mechanics Parameters for Hydroxyproline and Its Use in Molecular Dynamics Simulations of Collagen-Like Peptides
Sanghyun Park, Randall J. Radmer, Teri E. Klein, and Vijay S. Pande. J. Computational Chemistry (2005)

Solvation free energies of amino acid side chain analogs for common molecular mechanics water models
Michael R. Shirts and Vijay S. Pande. Journal of Chemical Physics (2005)

Foldamer dynamics expressed via Markov state models. I. Explicit solvent molecular-dynamics simulations in acetonitrile, chloroform, methanol, and water
Sidney Elmer, Sanghyun Park, & Vijay S. Pande. Journal of Chemical Physics (2005)

Guha Jayachandran, V. Vishal, Angel E. Garcia and V. S. Pande. Journal of Structural Biology (2006)

Unusual compactness of a polyproline type II structure

Bojan Zagrovic, J. Lipfert, E.J. Sorin, I. S. Millett, W.F. van Gunsteren, S. Doniach & V.S. Pande. PNAS (2005)

Empirical Force-Field Assessment: The Interplay Between Backbone Torsions and Noncovalent Term Scaling
Eric J. Sorin and Vijay S. Pande. Journal of Computational Chemistry (2005)

Exploring the Helix-Coil Transition via All-atom Equilibrium Ensemble Simulations
Eric J. Sorin and Vijay S. Pande. Biophysical Journal (2005)

Does Water Play a Structural Role in the Folding of Small Nucleic Acids?
Eric J. Sorin, Young Min Rhee, and Vijay S. Pande. Biophysical Journal (2005)

Dimerization of the p53 Oligomerization Domain: Identification of a Folding Nucleus by M.D. Simulations
Lillian T. Chong, Christopher D. Snow, Young Min Rhee, and Vijay S. Pande. J. Molecular Biology (2005)

Simulations of the role of water in the protein-folding mechanism

Young Min Rhee, Eric J. Sorin, Guha Jayachandran, Erik Lindahl, & Vijay S Pande. PNAS (2004)

Trp zipper folding kinetics by molecular dynamics and temperature-jump spectroscopy
Christopher D. Snow, Linlin Qiu, Deguo Du, Feng Gai, Stephen J. Hagen, & Vijay S Pande. PNAS (2004)

Does Native State Topology Determine the RNA Folding Mechanism?

Eric J. Sorin, B.J. Nakatani, Y.M. Rhee, Guha Jayachandran, V Vishal, & Vijay S Pande. J. Mol. Biology (2004)

Solvent Viscosity Dependence of the Folding Rate of a Small Protein: Distributed Computing Study
Bojan Zagrovic and Vijay S. Pande. Journal of Computational Chemistry (2003)

Insights Into Nucleic Acid Conformational Dynamics from Massively Parallel Stochastic Simulations
Eric J. Sorin, Young Min Rhee, Bradley J. Nakatani & Vijay S. Pande. Biophysical Journal (2003)

Multiplexed-Replica Exchange Molecular Dynamics Method for Protein Folding Simulation
Young Min Rhee & Vijay S. Pande. Biophysical Journal (2003)

The Trp Cage: Folding Kinetics and Unfolded State Topology via Molecular Dynamics Simulations
Christopher D. Snow, Bojan Zagrovic, and Vijay S. Pande. Journal of the American Chemical Society (2002)

Absolute comparison of simulated and experimental protein-folding dynamics
Christopher D. Snow, Houbi Ngyen, Vijay S. Pande, and Martin Gruebele. Nature (2002)

Native-like Mean Structure in the Unfolded Ensemble of Small Proteins
Bojan Zagrovic, C.D. Snow, Siraj Khaliq, M.R. Shirts, and Vijay S. Pande. J. Mol. Biology (2002)

Sim. of Folding of a Small Alpha-helical Protein in Atomistic Detail using Worldwide Distributed Computing
Bojan Zagrovic, Christopher D. Snow, Michael R. Shirts, and Vijay S. Pande. J. Molecular Biology (2002)

Atomistic protein folding simulations on the submillisecond timescale using worldwide distributed computing
Vijay Pande, et al. Peter Kollman Memorial Issue, Biopolymers (2002)

New Hardware Platforms

- Multi-core machines
 - Use MPI for communication
 - Crucial for simulating large systems



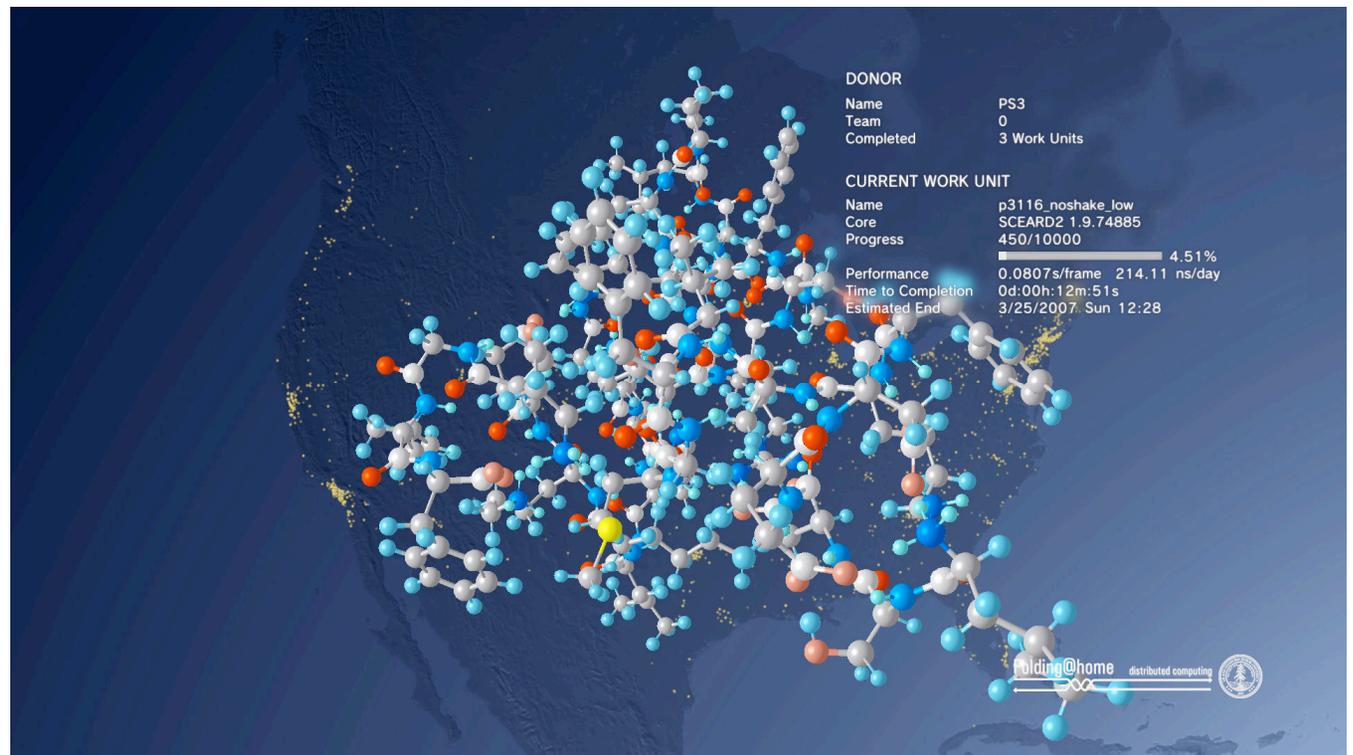
New Hardware Platforms

- Graphical processing unit (V. Vishal)
 - 20-30x speed up for certain simulations, not fully general
 - Only fairly minor changes to client and server needed



New Hardware Platforms

- Sony PlayStation 3
 - Cell processor
 - Client/core port by SCEA (Redwood City)
 - Nearly 150,000 users signed up since mid-March, ~30,000 active

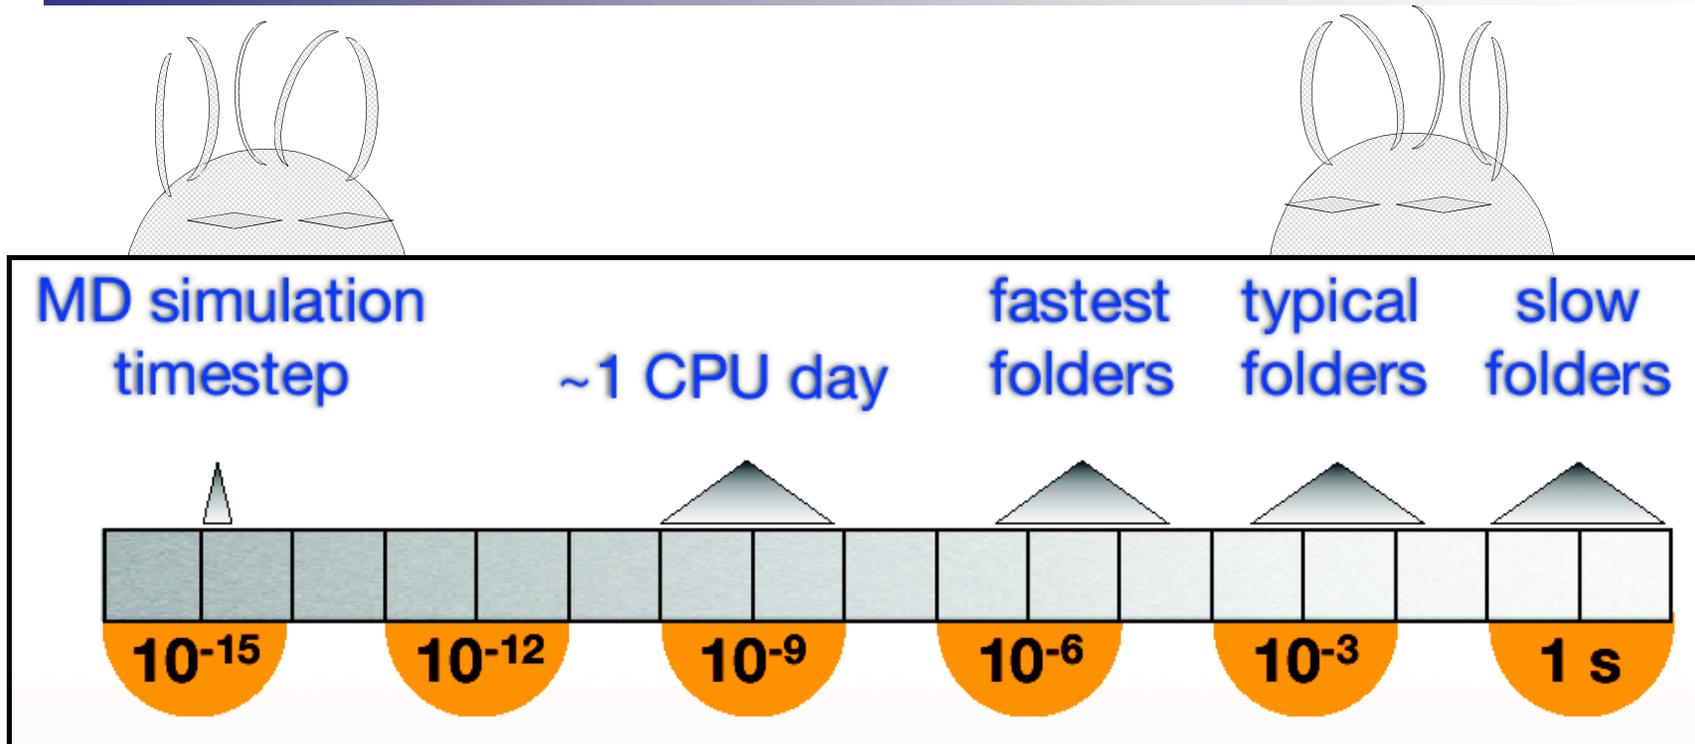




Example Parallel Techniques



So What if Many Computers?



*If you could simulate 1 ns/day,
then a 1 ms simulation would take
3000 years!*

Stochastic Processes



Movie by Calvin Co.

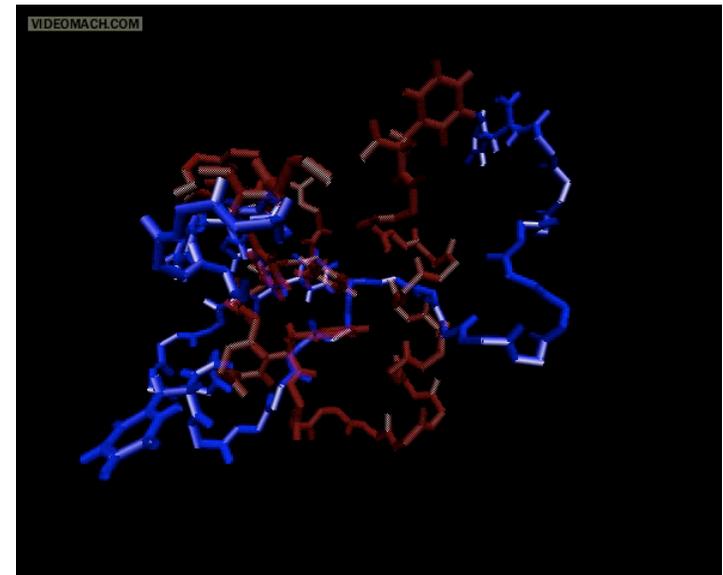
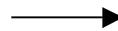
One takes too long

Out of large ensemble,
expect some to do faster

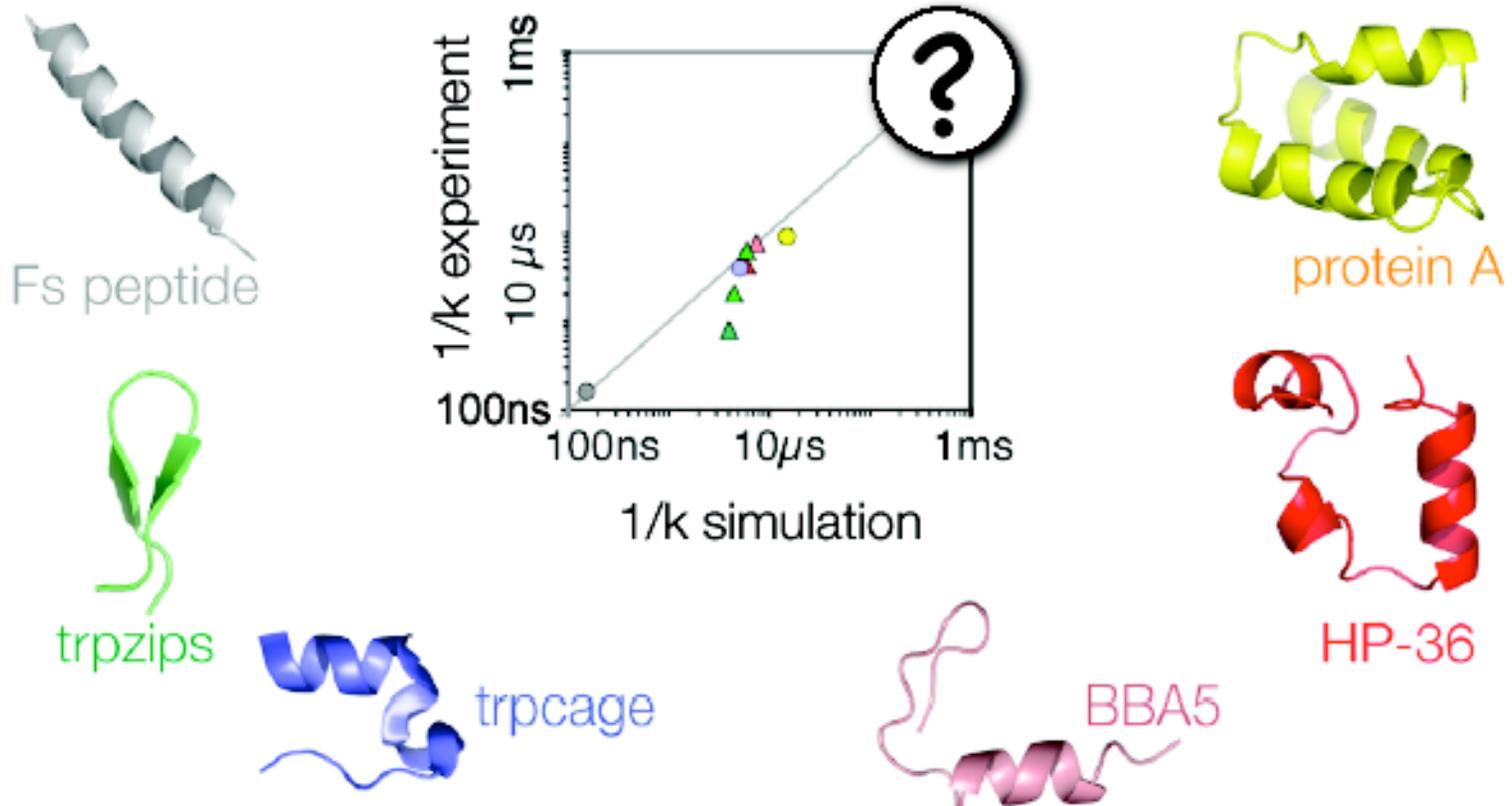
Villin and Protein A Simulations

	Structure	Trajectories	Total sampling
Villin headpiece		>20,000	~500 μ s
10-55 fragment B of protein A		>5,000	>100 μ s

A trajectory (blue) of villin reaching the folded state from the unfolded state. The native structure is shown fixed in red.



Folding Rate Validation



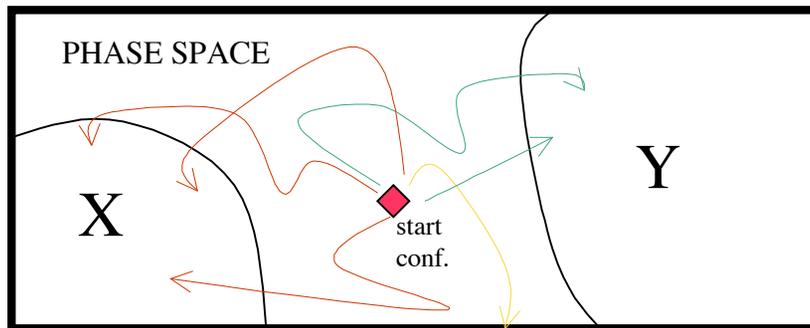
Perturbation Analysis

- Often interested in effect of a perturbation to a system
 - Perturbation can be physically meaningful (like changing temperature) or methodological
- Assessing such perturbations requires accounting for stochastic nature of dynamics
 - This requires much sampling



P(X,Y) Definition

P(X,Y) equals probability of reaching state X before state Y

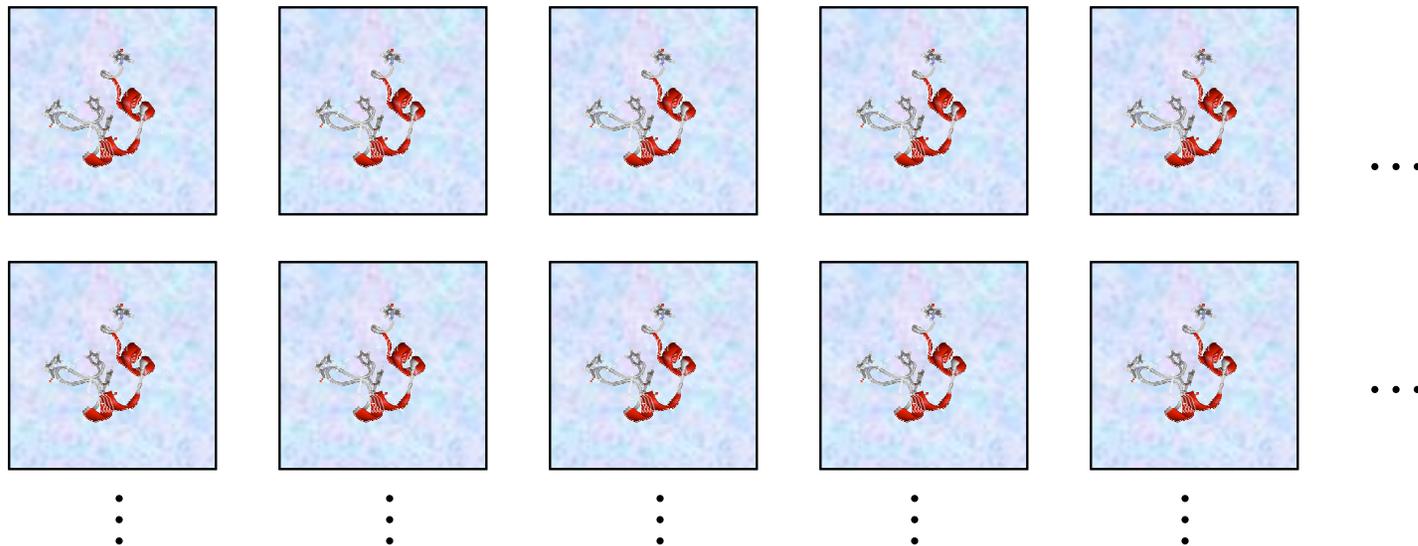


$$P(X,Y) = \frac{N(\rightarrow)}{N(\rightarrow) + N(\leftarrow)}$$

Is a generalization of P_{fold} :

$P_{\text{fold}} = P(X,Y)$ where X='folded' and Y='unfolded'

P(X,Y) Use

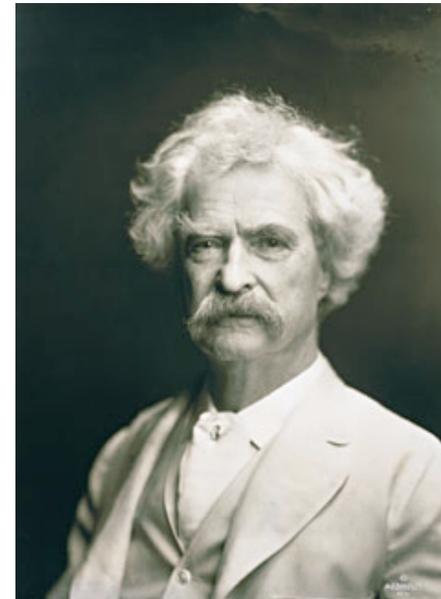


- Calculate value for a given conformation by running an ensemble of simulations and seeing what fraction reach state X before state Y
- Compute value again after a given perturbation and see if the values change, for various X and Y

Water

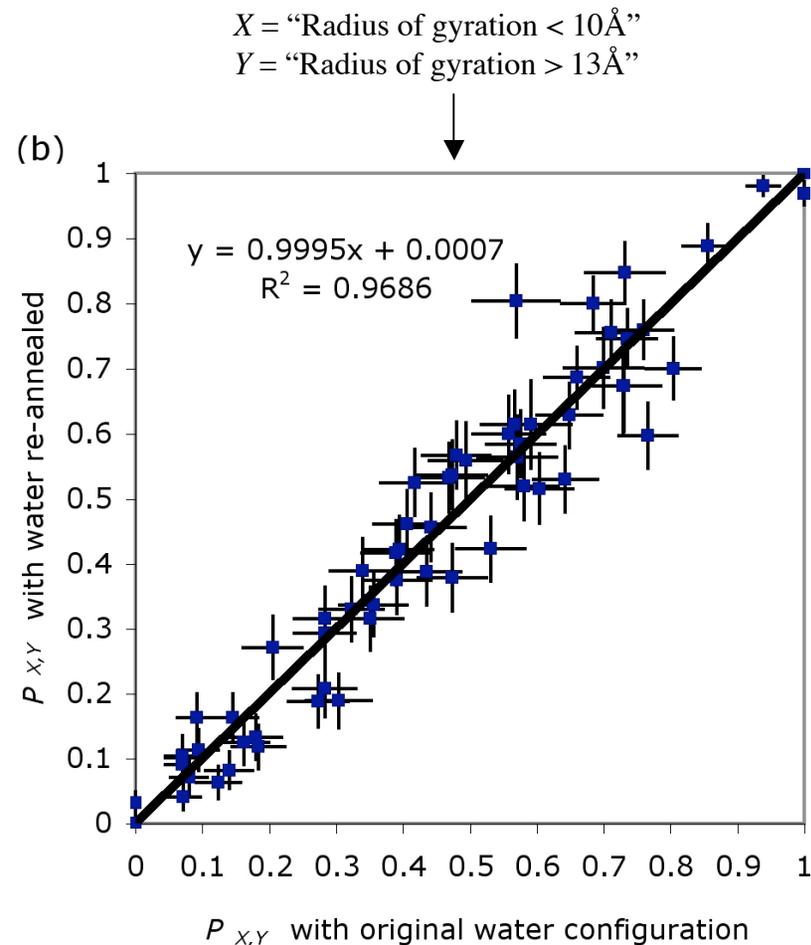
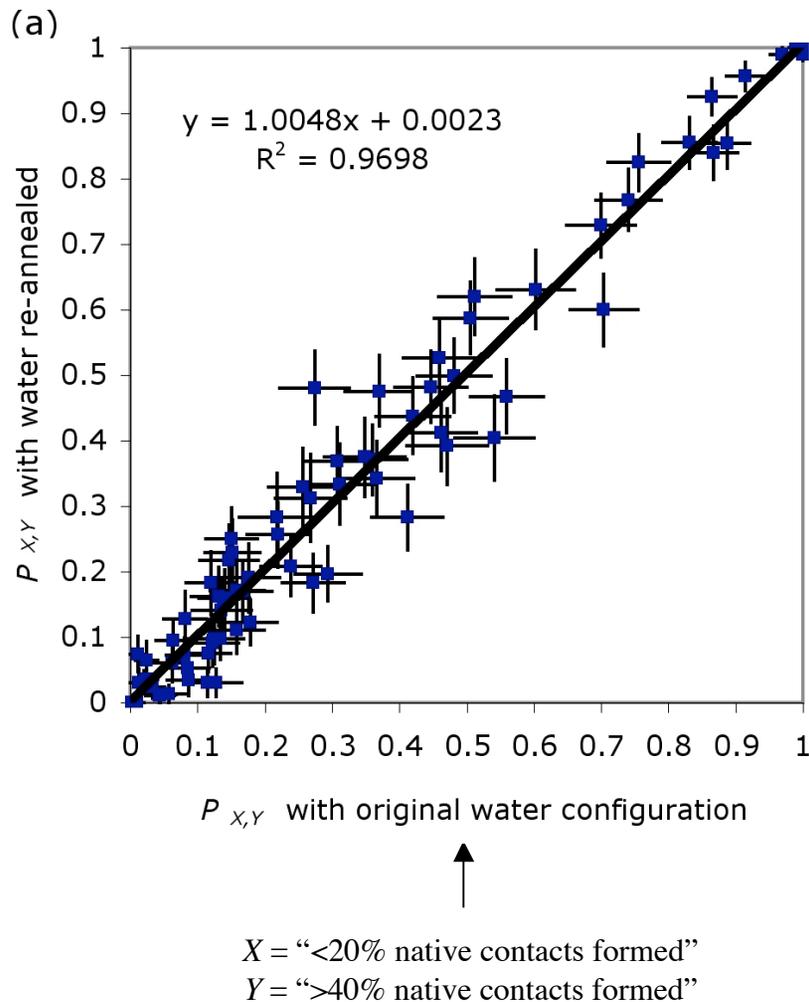
“Water,
taken in moderation,
cannot hurt anybody.”

-Mark Twain
(who didn't do simulations)

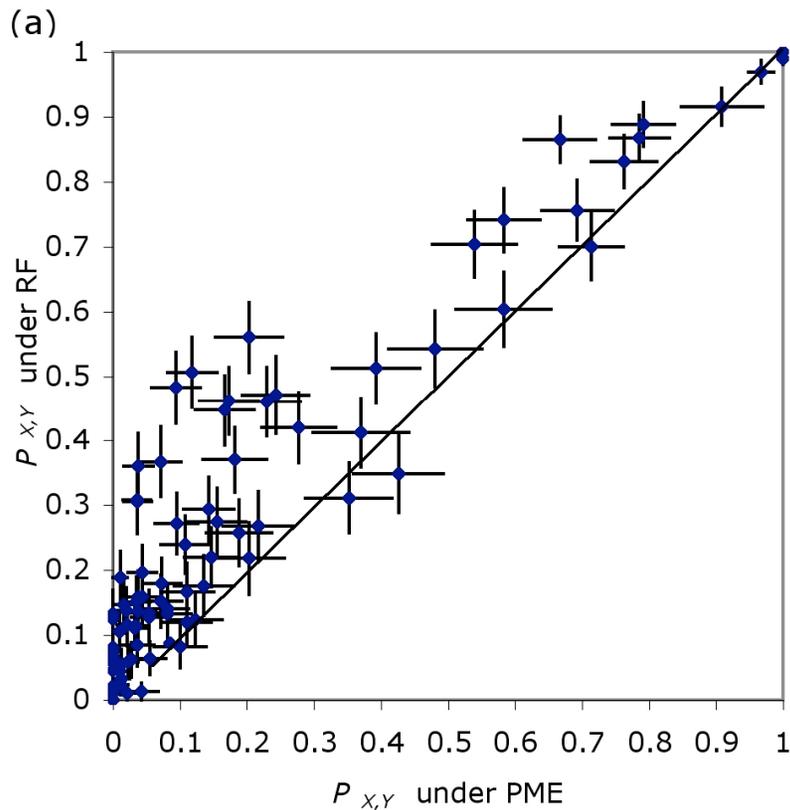


P(X,Y) and Water Configuration

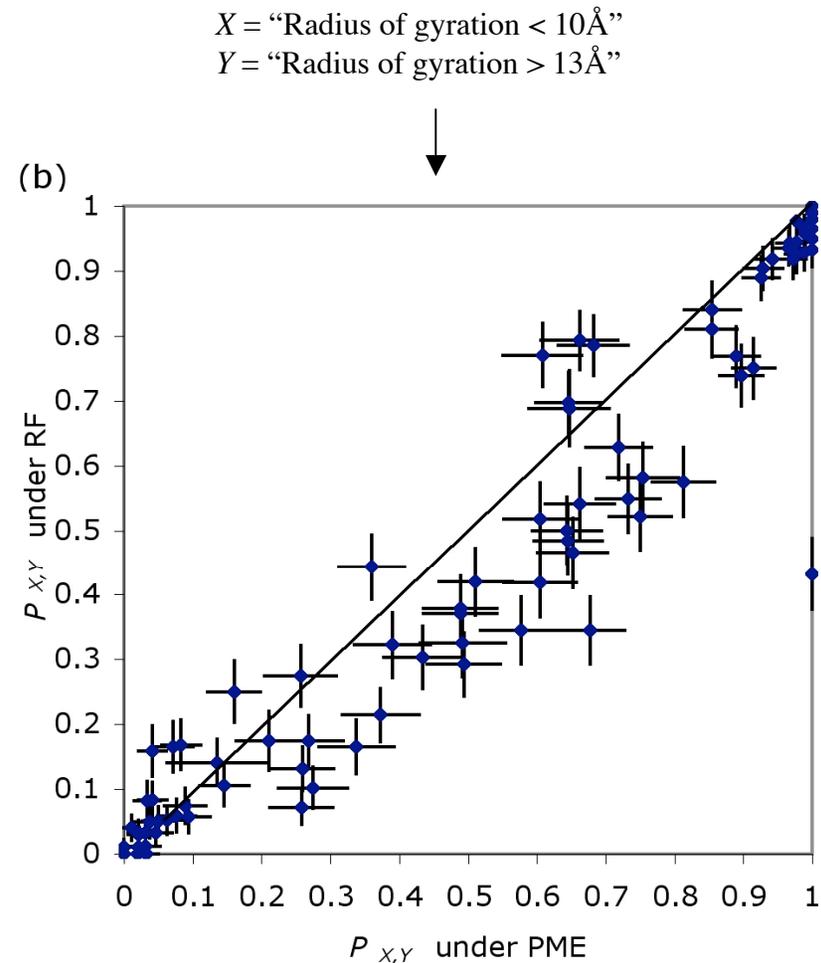
Each point on graph is associated with a different protein conformation.



P(X,Y) and Electrostatic Treatment



$X = <20\%$ native contacts formed”
 $Y = >40\%$ native contacts formed”



Learning

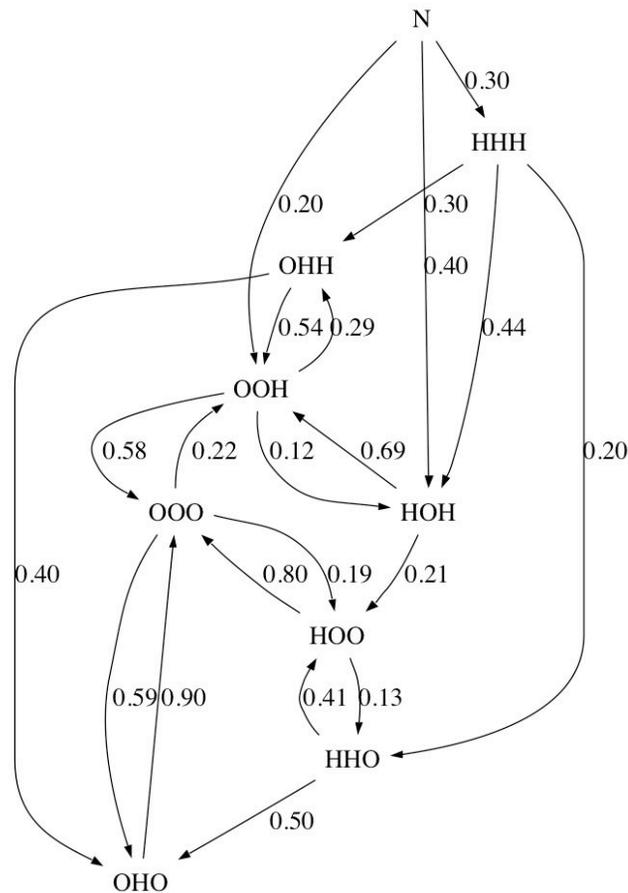
← Instruction

Observation



Protein A Helix Formation

■ Fractions normalized by source



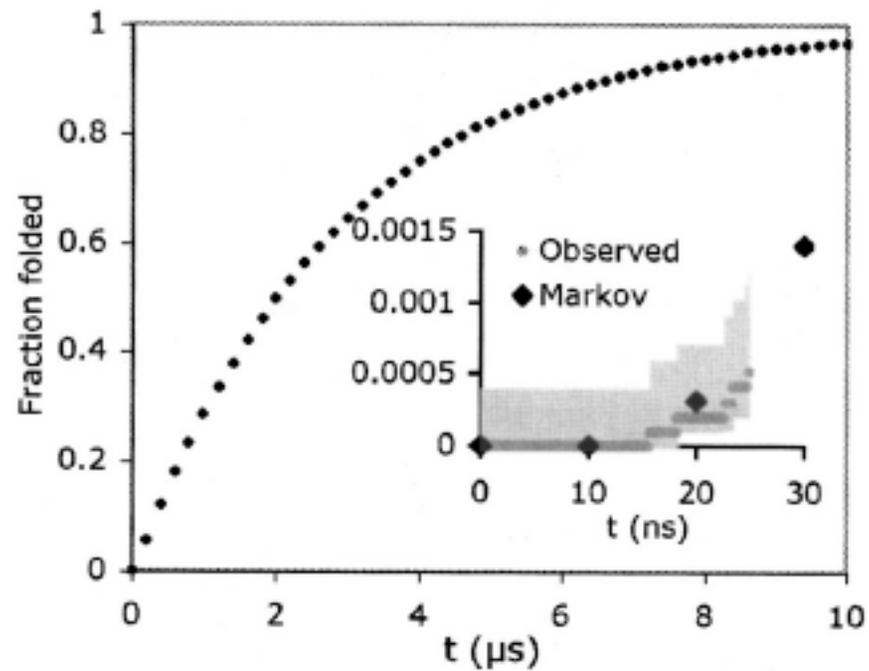
- Like for villin, observed the middle helix most commonly (approximately 60% of time) transiently forming first.
- If either the first or second helix existed in isolation, it was most often lost—in 80% and 90% of cases, respectively—before any additional helices formed. If the third helix existed in isolation, it gained the second helix in nearly 30% of cases.
- For each helix combination, the majority of observed transitions involved loss of a helix rather than gain of an additional one.

Markov Model

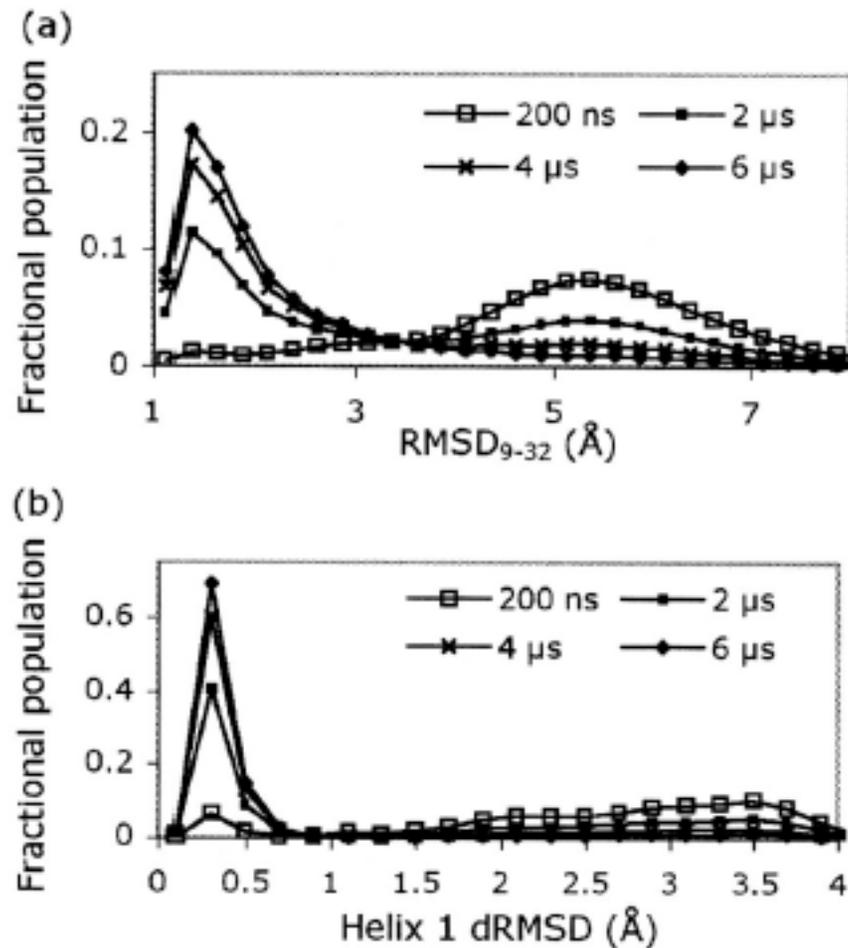
- Cluster conformations from simulations into states connected with weighted edges
- Element (i,j) of transition matrix P gives probability of going to state j given currently in state i
- Then, vector $\pi(n) = \pi(0) P^n$ gives fractional population of each state over time
- Several applications

Background shows cutout from a cluster transition map constructed from simulation data

Long Timescale Dynamics



Long Timescale Dynamics



Free Energy

- $\Delta G = \Delta H - T\Delta S$

↑ ↑ ↑ ↙
Change in Change in Temperature Change in
free energy enthalpy entropy

- Associated with probability

- More negative the ΔG of a change, higher the probability of the change

- State function (pathway independent)

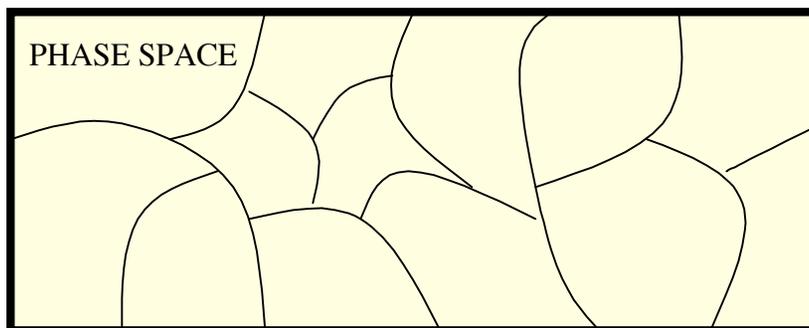


New Method

- Aims:
 - Accurate
 - Fast
 - Work without known structure for bound complex

- Basis:

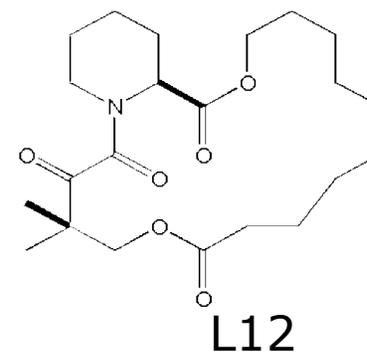
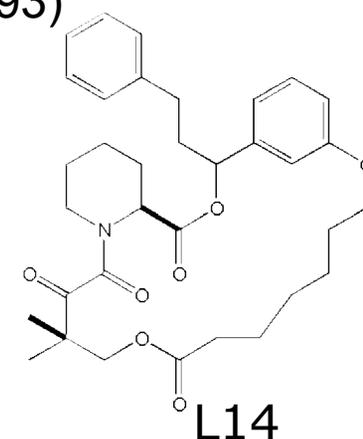
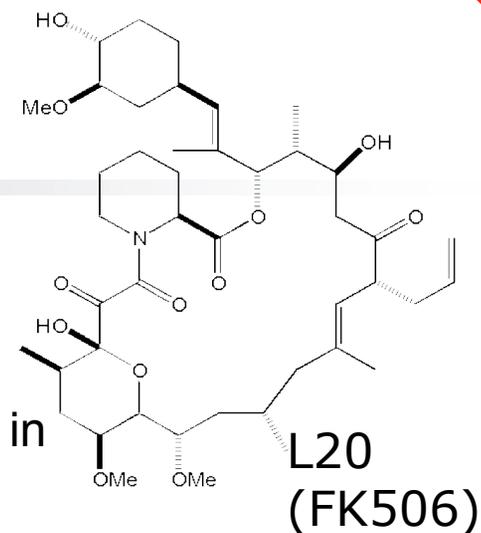
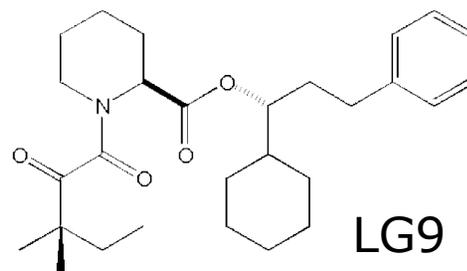
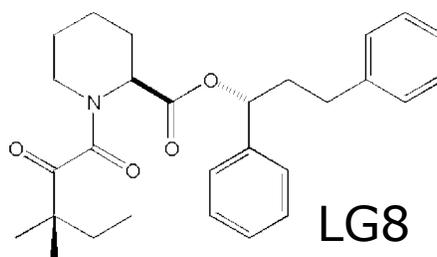
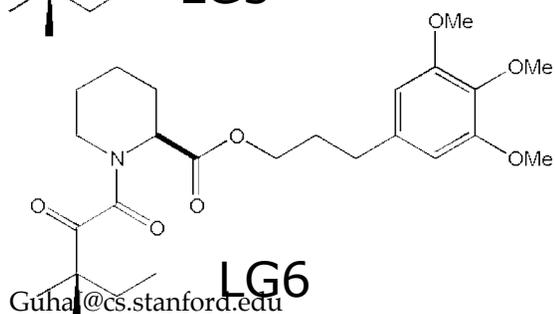
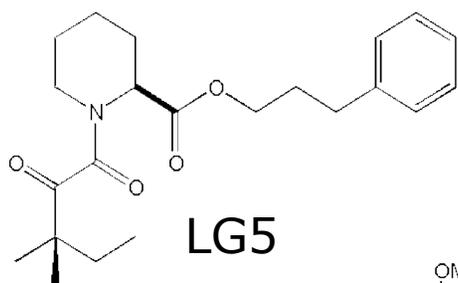
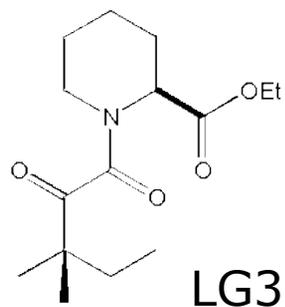
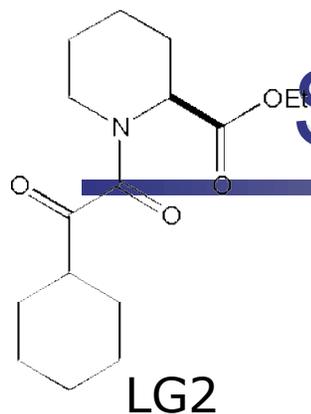
Divide configurational space into states and find free energy change in each independently



$$\Delta G(0, U \rightarrow 1, U) = -kT \ln \sum_{s \in R} \left[\exp \left(\left(-\Delta G(0, U \rightarrow 0, s) - \sum_{i=0}^N \Delta G(\lambda_i, s \rightarrow \lambda_{i+1}, s) \right) / kT \right) \right]$$

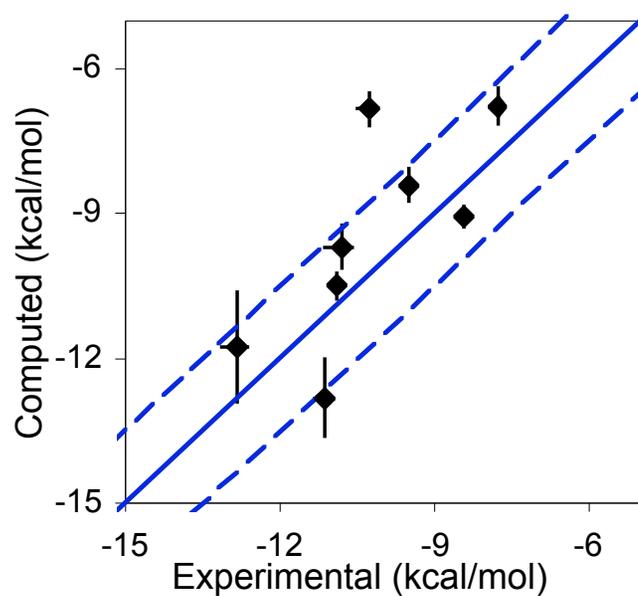
Guha Jayachandran, Michael Shirts, Sanghyun Park, Vijay Pande.
"Parallelized over parts computation of absolute binding free energy with docking and molecular dynamics." J.C.P. (2006).

Sample System



- FKBP12: a canonical system for binding calculations and experiments
- 107 residues, stable, well studied, involved in immune response
- Experimental data from Holt *et al*, *JACS* (1993)

Results

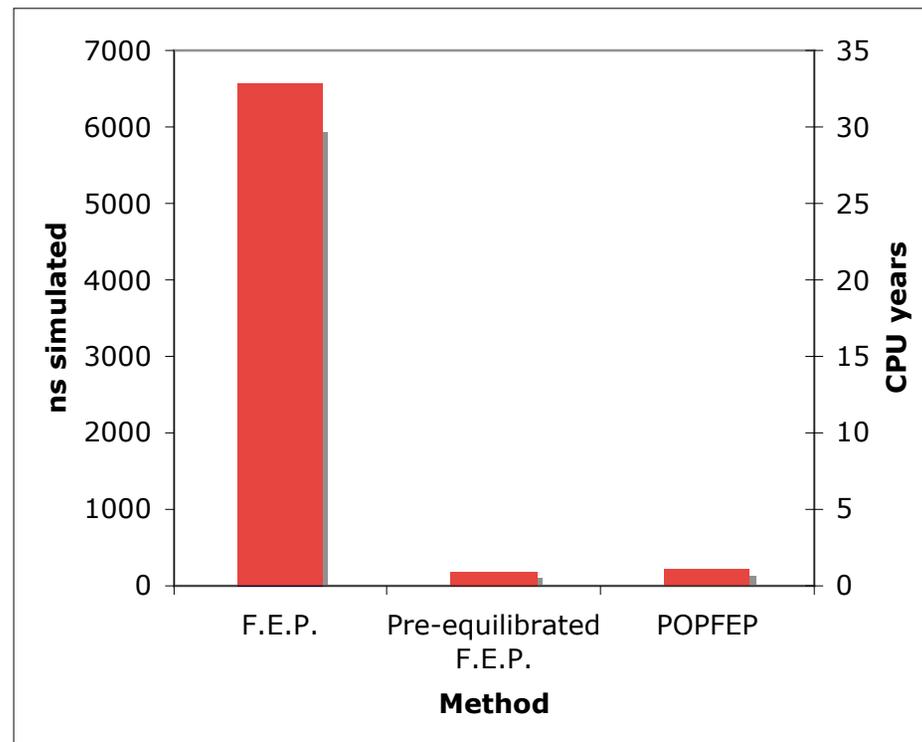


$y=x$ and $y=x\pm 1.5$
drawn as guides

1.6 kcal/mol RMSD between experimental and computed values

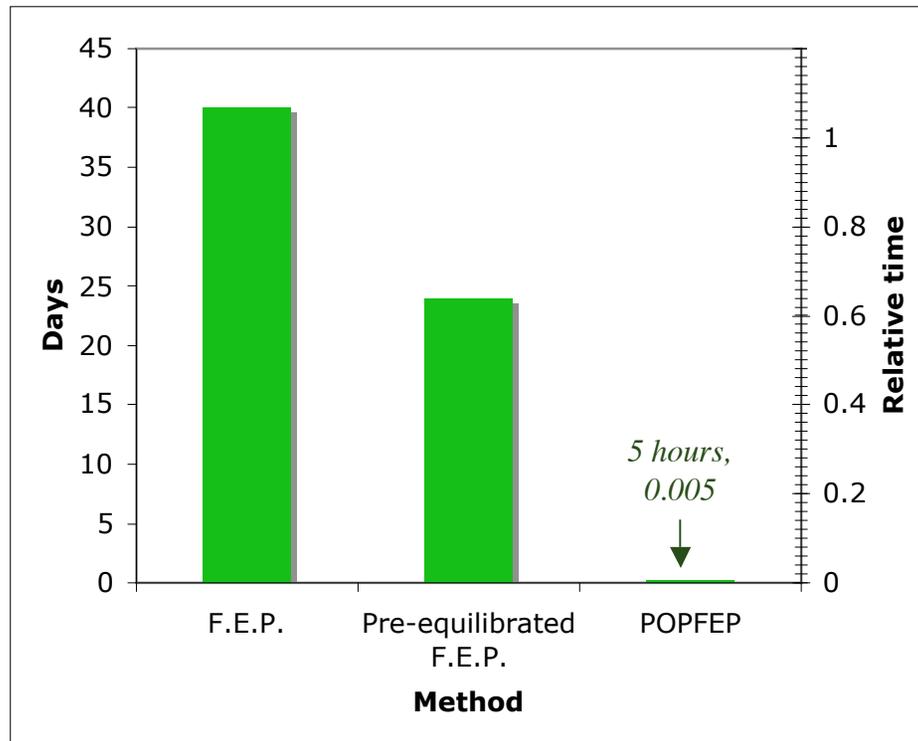
1.1 kcal/mol excluding L12

Simulation and CPU Time



For 1 binding free energy.

Rate Limiting Time



Limiting step sets a lower bound on wall clock time.

With 2000 processors, wall clock time approximately equal to above figures.

Future

Methods to Applications

Biochemical understanding, applications, etc.

Protein folding (kinetics)

Free energy calculation
(thermodynamics)

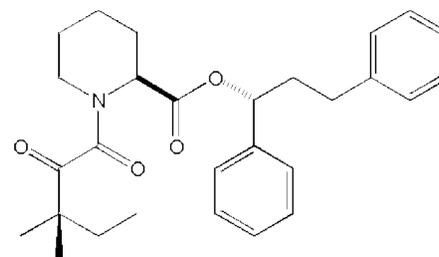
Same underlying issues of
models, sampling, and
domain context

Massively parallel
computing

New, more efficient
methods

Great Possibilities

- Modeling and simulation of varied disease/agent-relevant proteins (folding, aggregation, and binding)
- Large scale virtual screening of small molecule drugs
- Lead discovery and refinement



Acknowledgments

- Prof. Vijay Pande



- Former group members

- Bojan Zagrovic
- Christopher Snow
- Eric Sorin
- Lillian Chong
- Mark Engelhardt
- Michael Shirts
- Sanghyun Park
- Sidney Elmer
- Siraj Khaliq
- Stefan Larson
- Sung-Joo Lee
- Young Min Rhee

- www.cruxlux.com

- Current group members

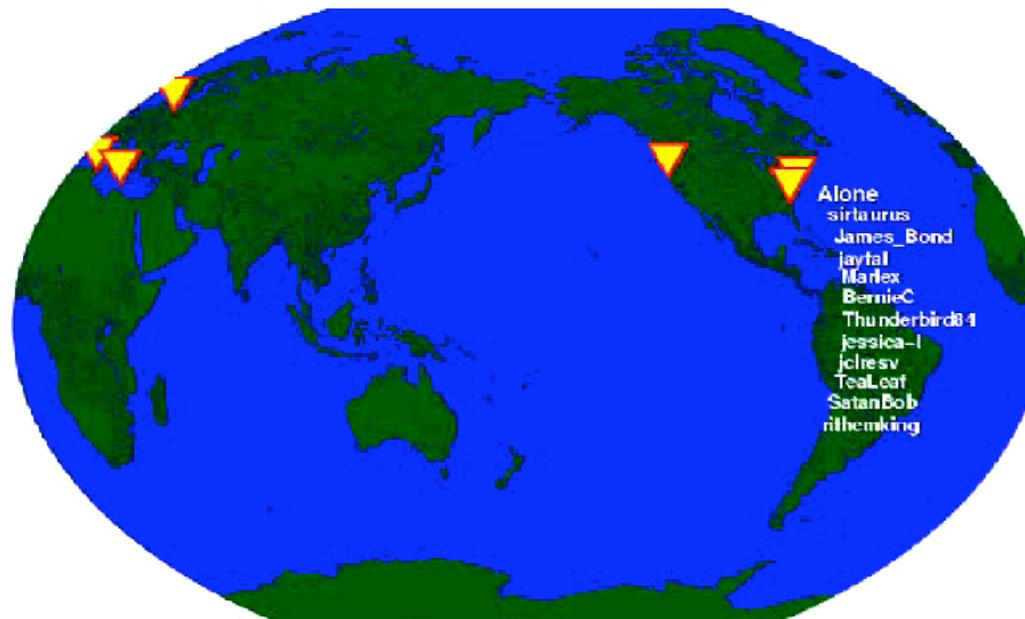
- Adam Beberg
- Dan Ensign
- Del Lucent
- Edgar Luttmann
- Jason Wagoner
- Jeremy England
- John Chodera
- Kim Branson
- Nick Kelley
- Nina Singhal
- Paula Petrone
- Relly Brandman
- Vishal Vaidyanathan

- Folding@Home participants

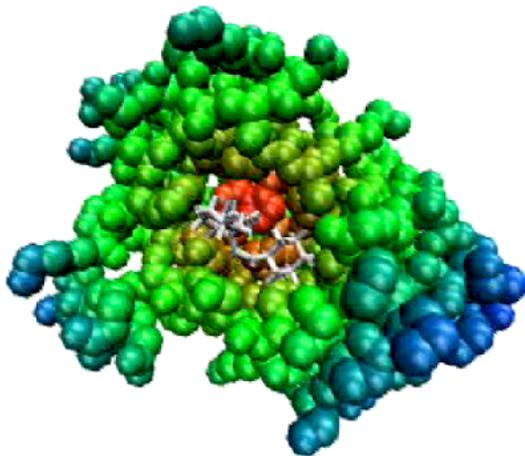
- *Listing at folding.stanford.edu*

- Funding

- N.D.S.E.G. Fellowship
- N.S.F./N.I.H. grants



Examining Biomolecular Phenomena with Massively Parallel Computation



Guha Jayachandran, Ph.D.
GuhaJ@cs.stanford.edu
June 20, 2007