

High Performance Computing Outreach to Industry



Public Presentation

www.hpcinnovationcenter.com/

Jeff Wolf

*Deputy Director and
Chief Business Officer
HPC Innovation Center*

LLNL-PRES-643607

This work was performed under the
auspices of the U.S. Department of Energy
by Lawrence Livermore National Laboratory
under contract DE-AC52-07NA27344.
Lawrence Livermore National Security, LLC



HIGH PERFORMANCE COMPUTING
INNOVATION CENTER

Leading Supercomputing Center Seeks Partners for High-Impact Collaborations

The High Performance Computing Innovation Center (HPCIC) is an outreach initiative to industry and academia by Lawrence Livermore National Laboratory under the auspices of the U.S. Department of Energy.

In support of LLNL's national security mission to:

- boost American competitiveness
- accelerate science and technology advancements
- develop a future HPC-skilled workforce




**Situated on the Livermore Valley Open Campus,
In Livermore, California**

The HPC Innovation Center seeks to . . .


- 1) ***Engage with companies*** to deliver transformational HPC-based solutions
... that combine HPC expertise with supercomputing technologies
- 2) ***Engage with HPC ecosystem partners*** (equipment suppliers, ISVs, other HPC centers, academia, expert users, etc.)
... to collaboratively deliver sustainable advanced computing innovations that broadly benefit the economy
- 3) ***Engage with academia and students*** to advance computational sciences
... and develop future HPC-skilled professionals



Offering companies benefits of scale: National labs lead in development and effective use of supercomputers



Industry historically lags HPC state-of-the-art by 5-15 years, and by 1 or more orders in performance




HPC Innovation Center: 5+ PF of supercomputing and decades of expertise available to benefit industrial partners



- Livermore Valley Open Campus offers a campus-like environment, permitting access by all partners
- Unclassified HPC Enclave allows outside users access to over 800 TF of shared LLNL cluster resources
- 5.0 PF IBM Blue Gene/Q “Vulcan” now available for industrial use
- Team approach coupled with exceptional resources enables creation of transformational solutions

~100,000 times more powerful than a desktop computer
= 100,000 times more realistic simulations

What does 100,000 times more realistic look like?



1980 Video Arcade game

1x ZiLOG Z80 @ 3.072 MHz

3.1
megaFLOPS

What does 100,000 times more realistic look like?



2011 PC Video Game

GeForce GT530

260
gigaFLOPS

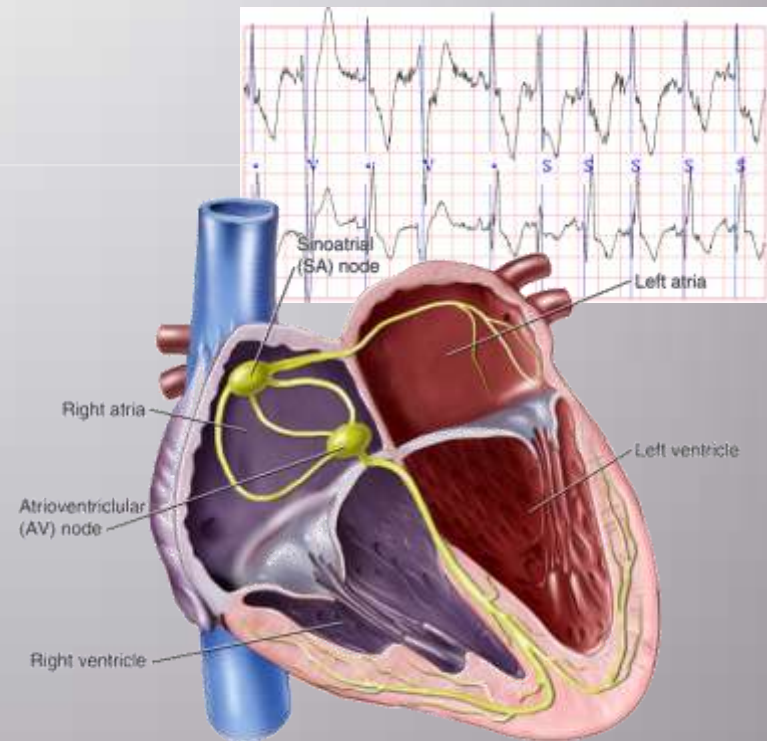
Another 100,000x faster: near-cellular resolution, real-time electro-physiological simulation of a beating human heart




Sequoia computer

20
petaFLOPS

Full scale simulation allows unprecedented investigation of arrhythmia in human heart



Heart wedge simulation in presence of d-sotalol





- Drug presence modeled as block to I_{Kr} channel
- S1-S2 stimulation introduces transmural reentrant pattern

The Cardioid Model on Sequoia

A joint capability development between LLNL and IBM Research

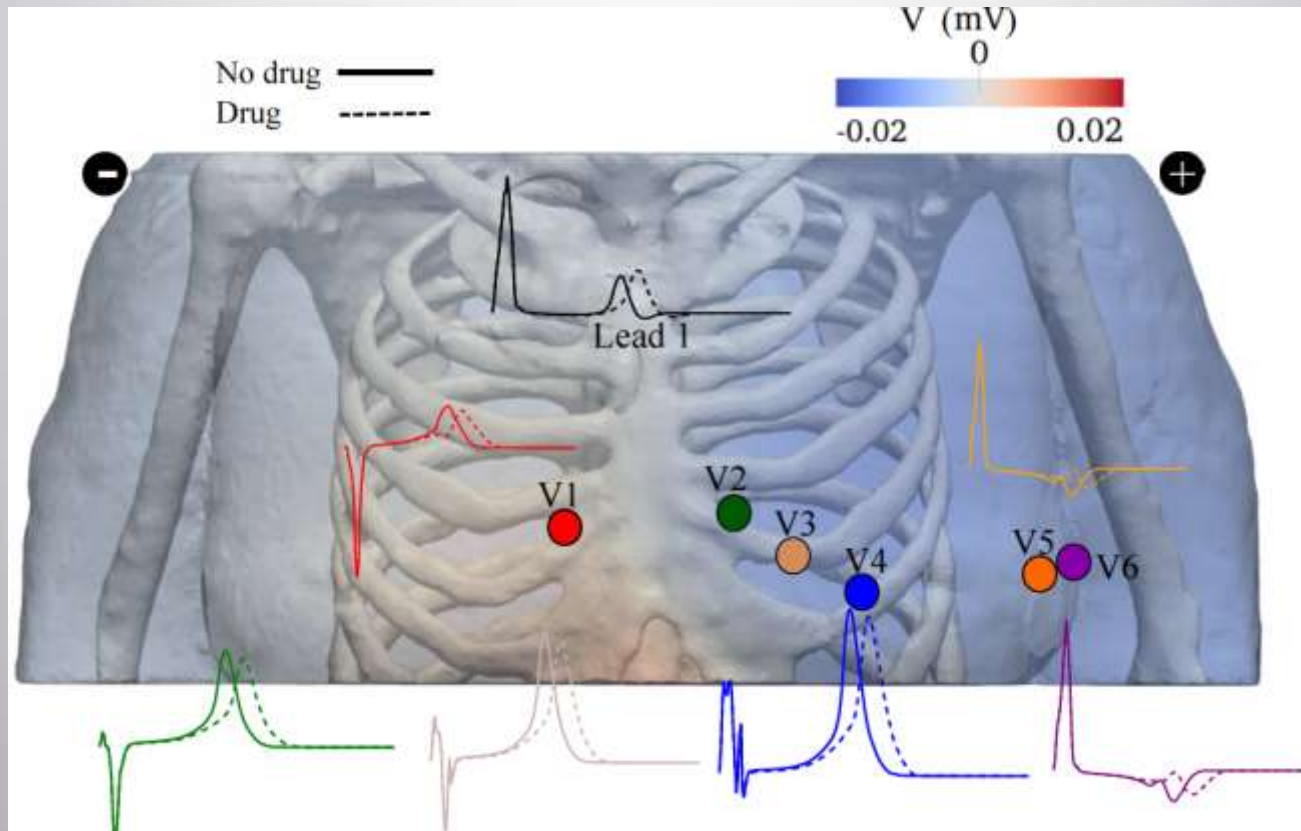
Anatomical model of human heart



96 racks of Blue Gene/Q in Sequoia

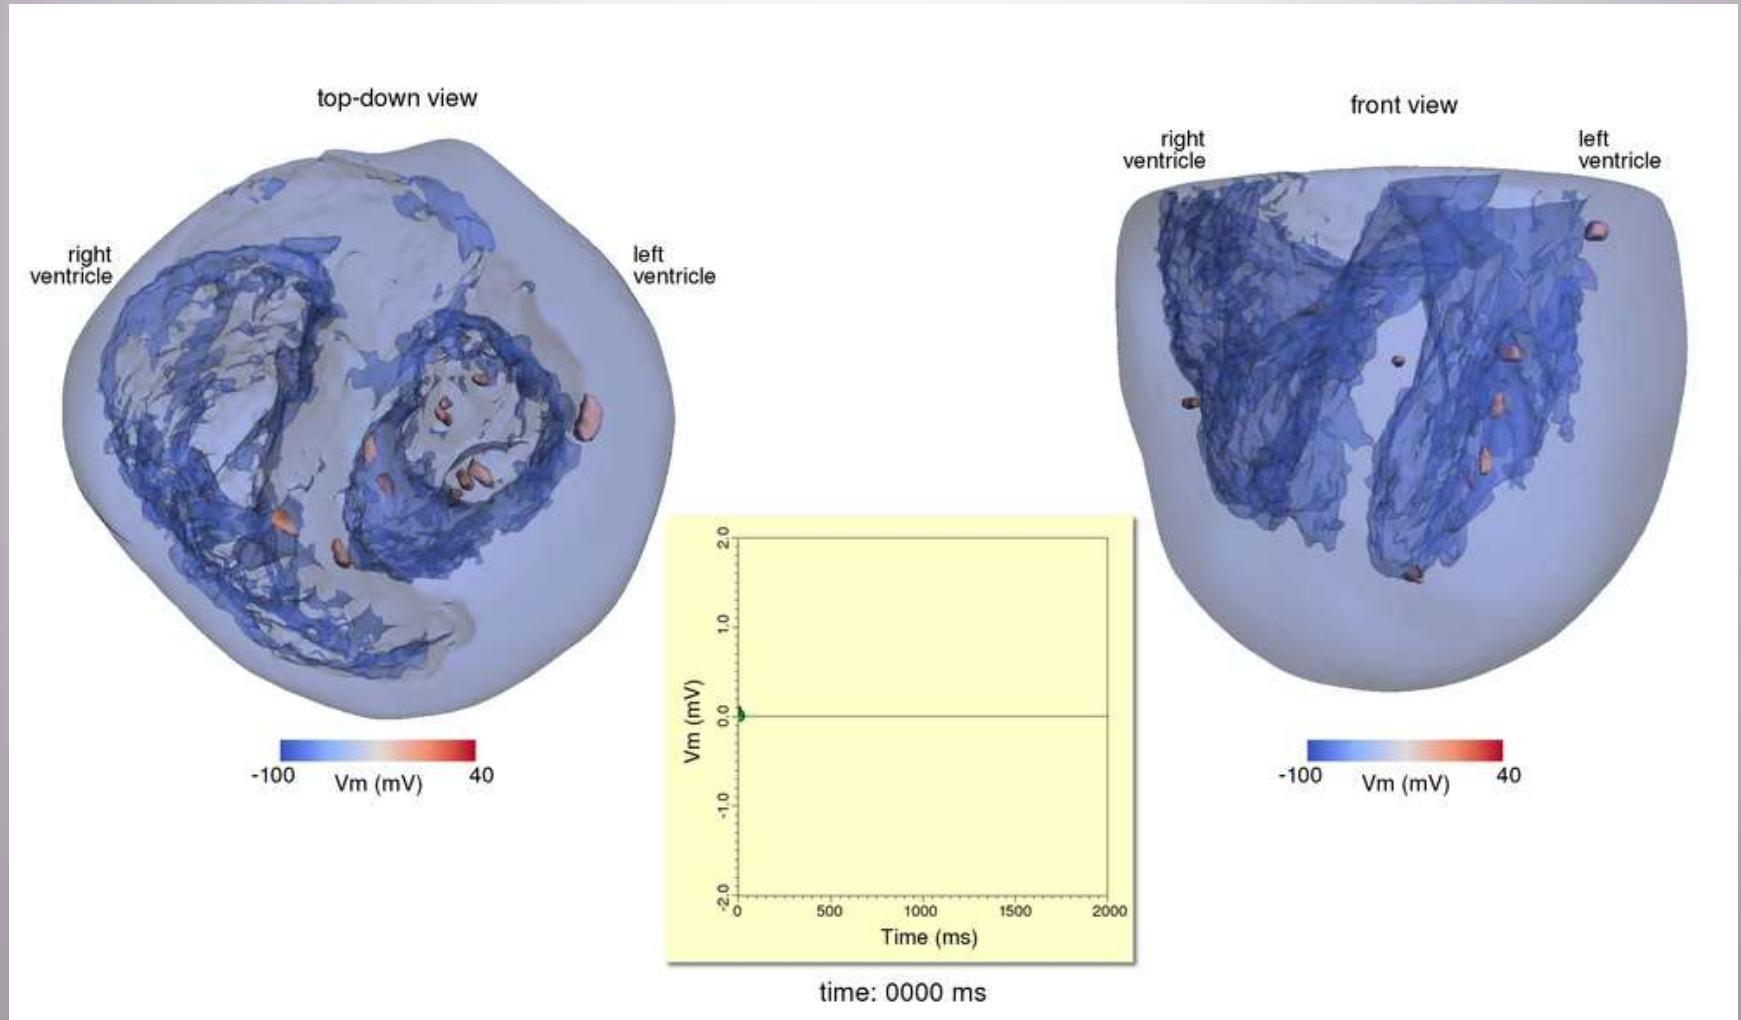
- Complete heart is 370-700 million cells in a complex geometry
- Sub-domains of ~200-500 cells must be mapped to each of 1.5 M cores
- Cell model must be computed on every cell for each time step including reaction and diffusion
- Sub-domains must exchange boundary data on each time step
- Code is written from scratch to take advantage of BG/Q hardware
- Extensive algorithmic improvements
- *1200x faster than previous record*
- 2012 Finalist in Gordon Bell Prize

Reconstruction of Torso to Enable Comparison of Simulation of ECGs



- Torso reconstruction includes bones, organs, skin, etc.
- For model validation and comparison to clinical signals
- Can simulate drug effects and pathology (e.g., infarct, heart failure, etc.)

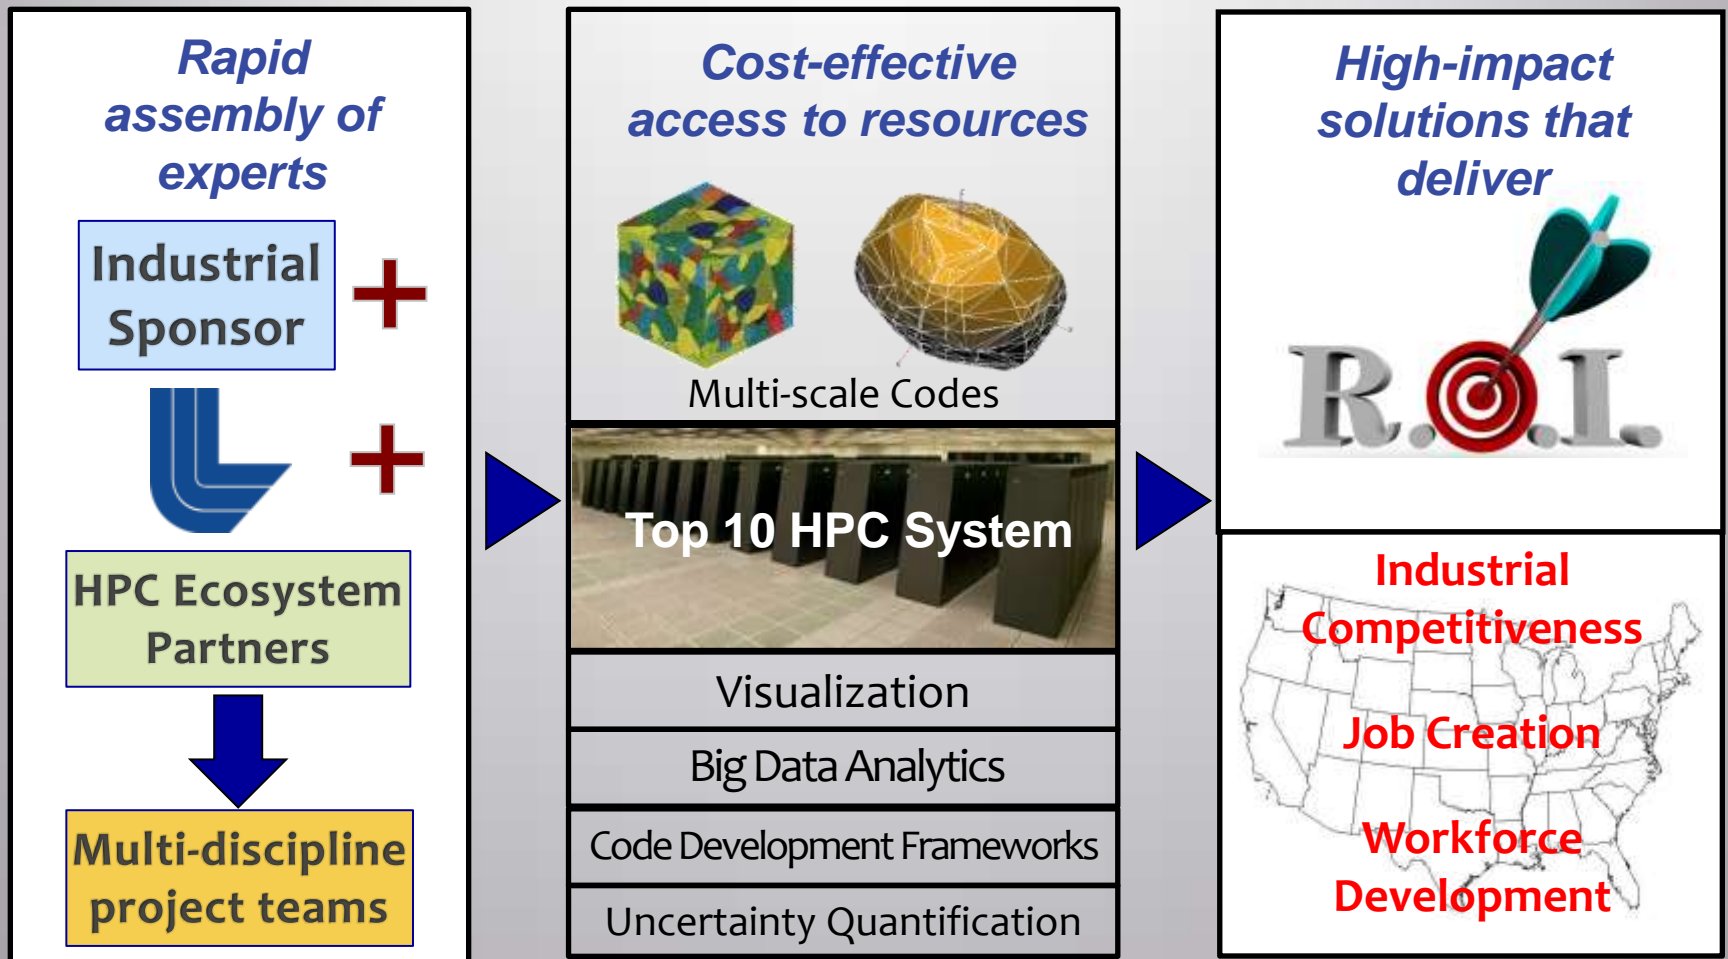
Capturing transmural reentrant activation



How scale helps: Cardioid model enables previously impossible research and new productivity tools

- Drug screening for adverse effects
- Drug discovery/design
- Optimum drug dosage
- Cardiac ablation therapies for arrhythmia
- Cardiac device placement and settings

HPCIC offers client companies depth to develop, prove and deploy transformational HPC solutions






Decades of investments in supercomputing expertise accessible to industry

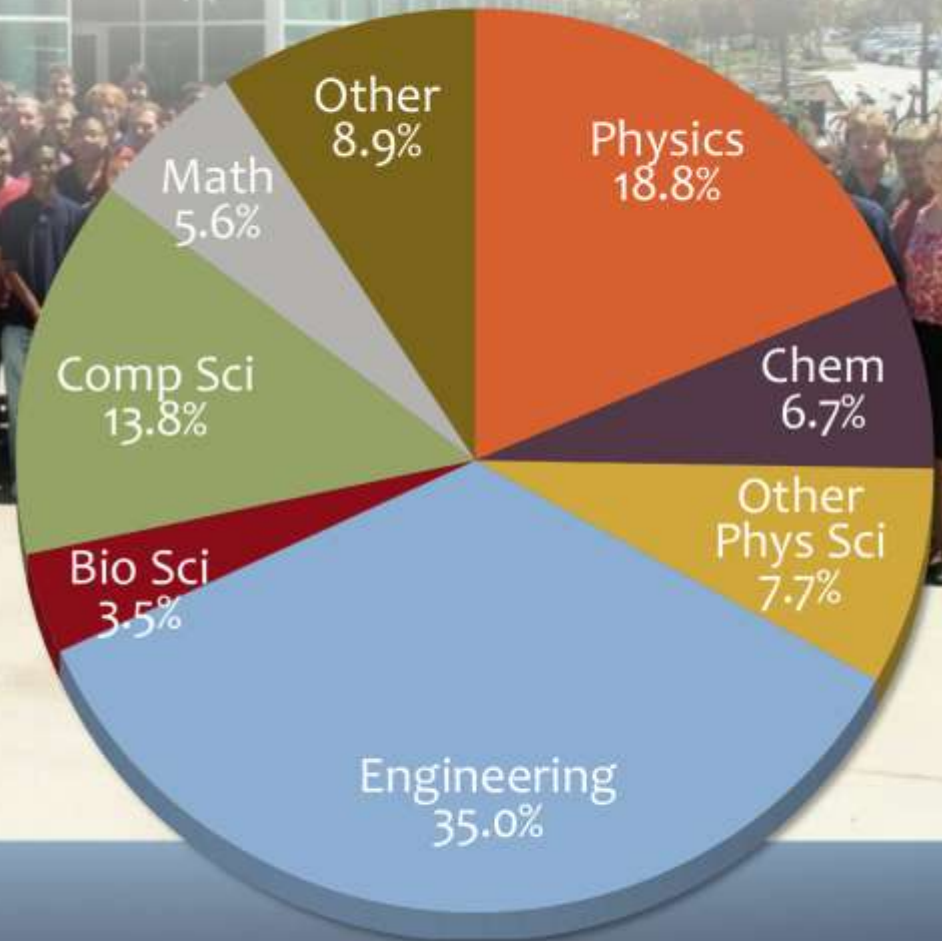
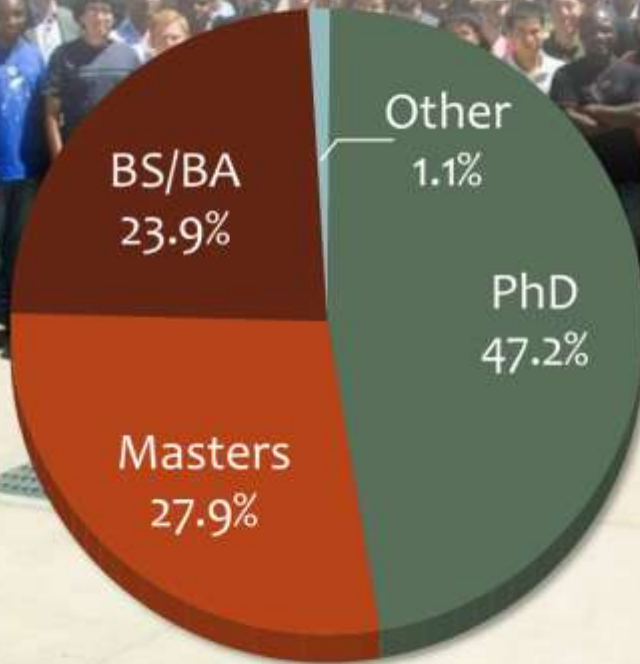
Lawrence Livermore National Laboratory: A National Science and Technology Resource

Award-winning multi-disciplinary computational science and engineering deep expertise in vital domains:

- Mathematics, physics, chemistry, biology, material science
- Energy management, alternative energy technologies, life sciences, earth sciences, transportation, green technologies
- Cyber security, networking, program management
- Massively parallel computing, data analytics, visualization



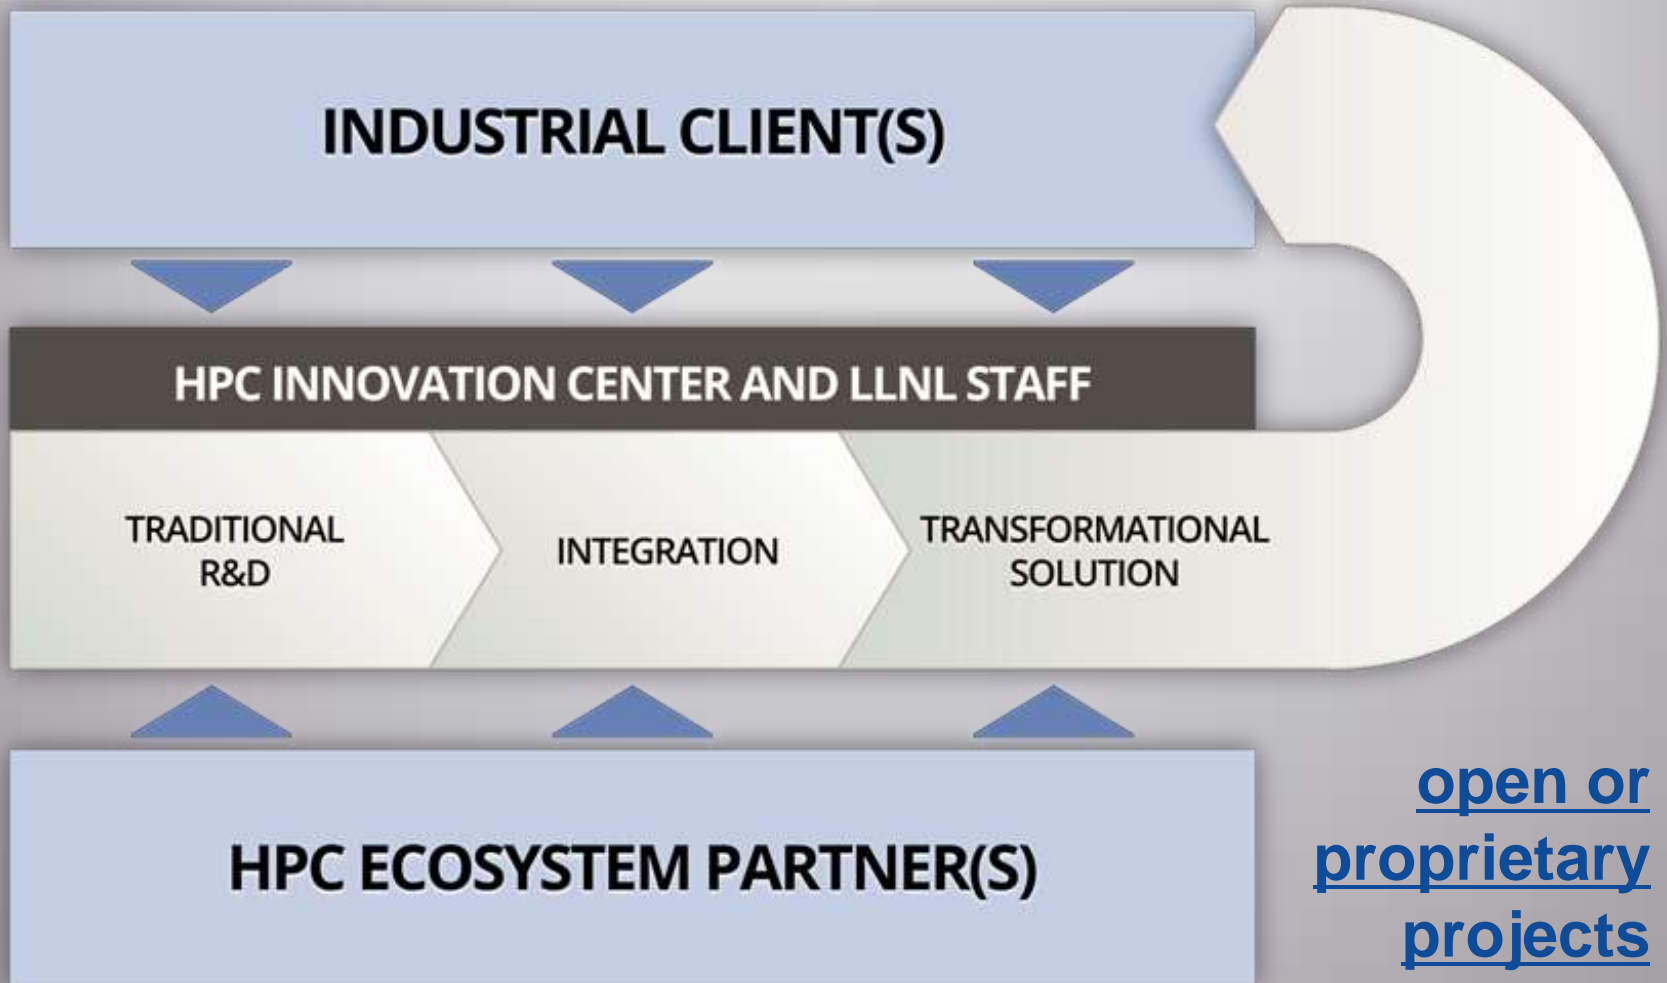
LLNL has a broad, multidisciplinary workforce



Total S&T staff: 2704


Annual budget ~\$1.5B

HPCIC's scalable engagement model goes beyond pure research to deliver integrated solutions . . . on-demand



Deep Computing Solutions:

a collaboration between IBM and LLNL



- Operates within LLNL's HPC Innovation Center as an HPC ecosystem partner option for industrial engagements
 - Joint staffing of industrial partner projects by experts from both IBM and LLNL
 - Extends HPCIC's range of industrial solutions to include integration, installation and management of enterprise-critical systems
 - Built on 20 year track record of collaborative innovation
 - Seven #1 ranked supercomputers, five Gordon Bell awards
- Shared commitment by IBM and LLNL to explore collaboration with other ecosystem partners to broaden adoption of advanced computing technologies by American industries

Partnership with STFC - Daresbury, UK

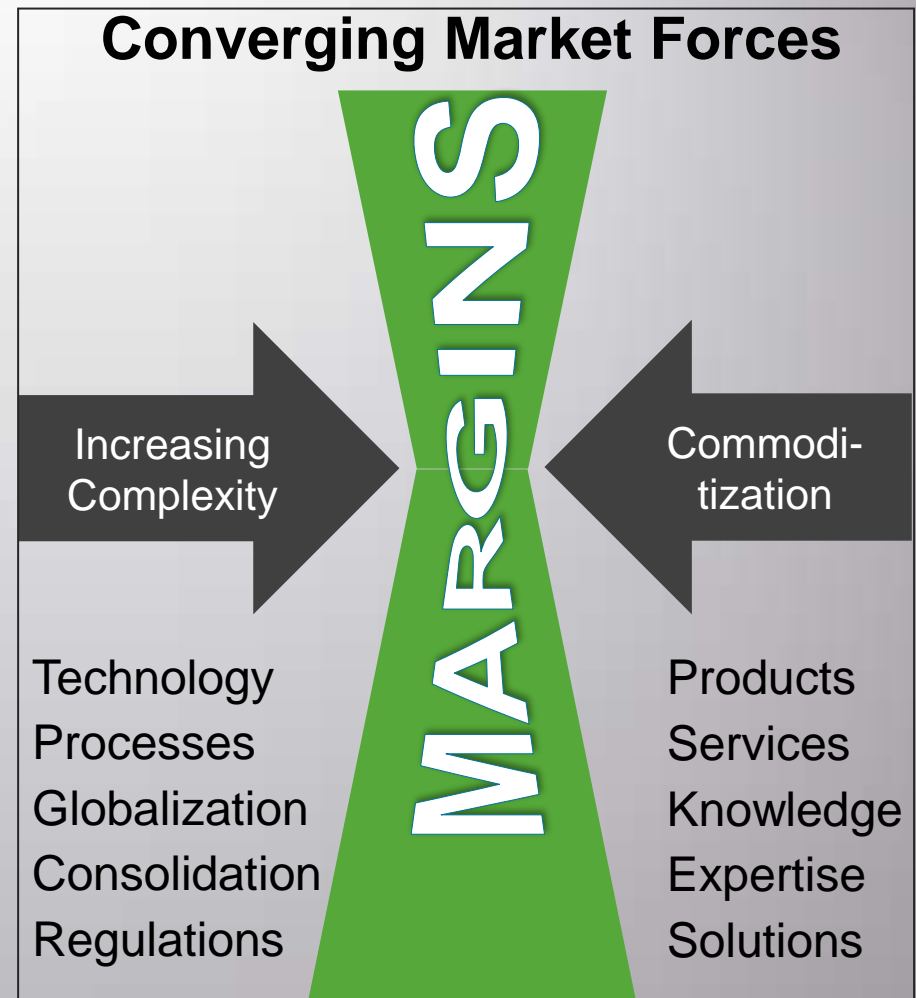


Hartree Centre
Science & Technology Facilities Council

- MOU signed August 29, 2013 between two similar organizations
 - Both Hartree Centre and HPC Innovation Center focus on boosting innovation and economic competitiveness through HPC engagements with industry and academia
 - Both centers utilize Blue Gene/Q supercomputers and Linux clusters
- Both parties see relationship as an accelerator to enable
 - Sharing: expertise, best practices, tools, software, solutions, successes
 - Rapid solution development on Blue Gene/Q, especially for data analytics – BG/Q is #1 architecture on Graph 500 list
 - Unique value propositions for large, multi-disciplinary projects, especially to to multi-national organizations

Advanced computing helps companies differentiate and better understand and manage complexity

- Complex technologies can be tamed using multi-scale, multi-physics modeling and simulation
 - Biomedicine
 - Alternative energy production
 - Engineered materials
 - Moore's Law *difficult challenges*
 - HPC-aided engineering
- Complex systems, processes and regulatory compliance can be more effectively managed using data analytics
- Solutions need to produce improved business metrics
 - Margins, revenues, valuation



Thull, J. *Mastering the Complex Sale: How to Compete and Win When the Stakes are High*. 2nd ed. John Wiley and Sons. 2010. p.10.

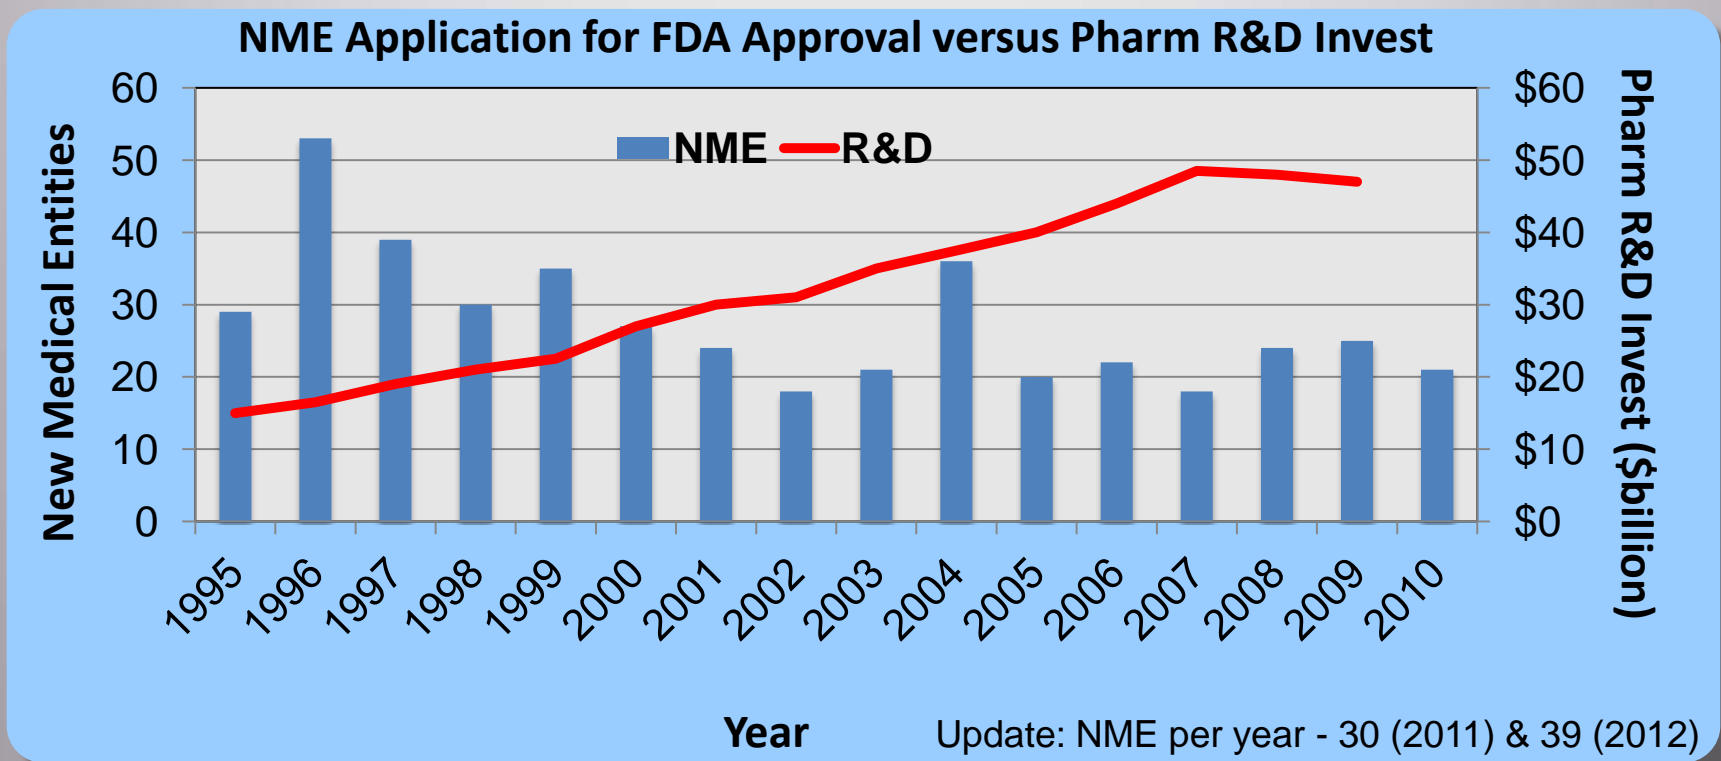
Helping an SME do HPC-enabled drug design



- Biopharmaceutical company focused on discovery, development and commercialization of innovative antibiotics for serious infections
 - Founded 2004, <100 employees
 - Tough, but necessary business: long regulatory cycle, short product life-cycle
- Engaged with LLNL since 2008 to identify and optimize lead molecules, and test them against Gram-negative, multi-drug-resistant biodefense pathogens using modeling and simulation
- Outcomes from the work:
 - Advances in Trius drug pipeline linked to increased valuation – Trius stock price has nearly tripled in 2013, followed by recent offer-to-acquire by Cubist
 - Validation of LLNL’s petascale parallel molecular docking program - VinaLC

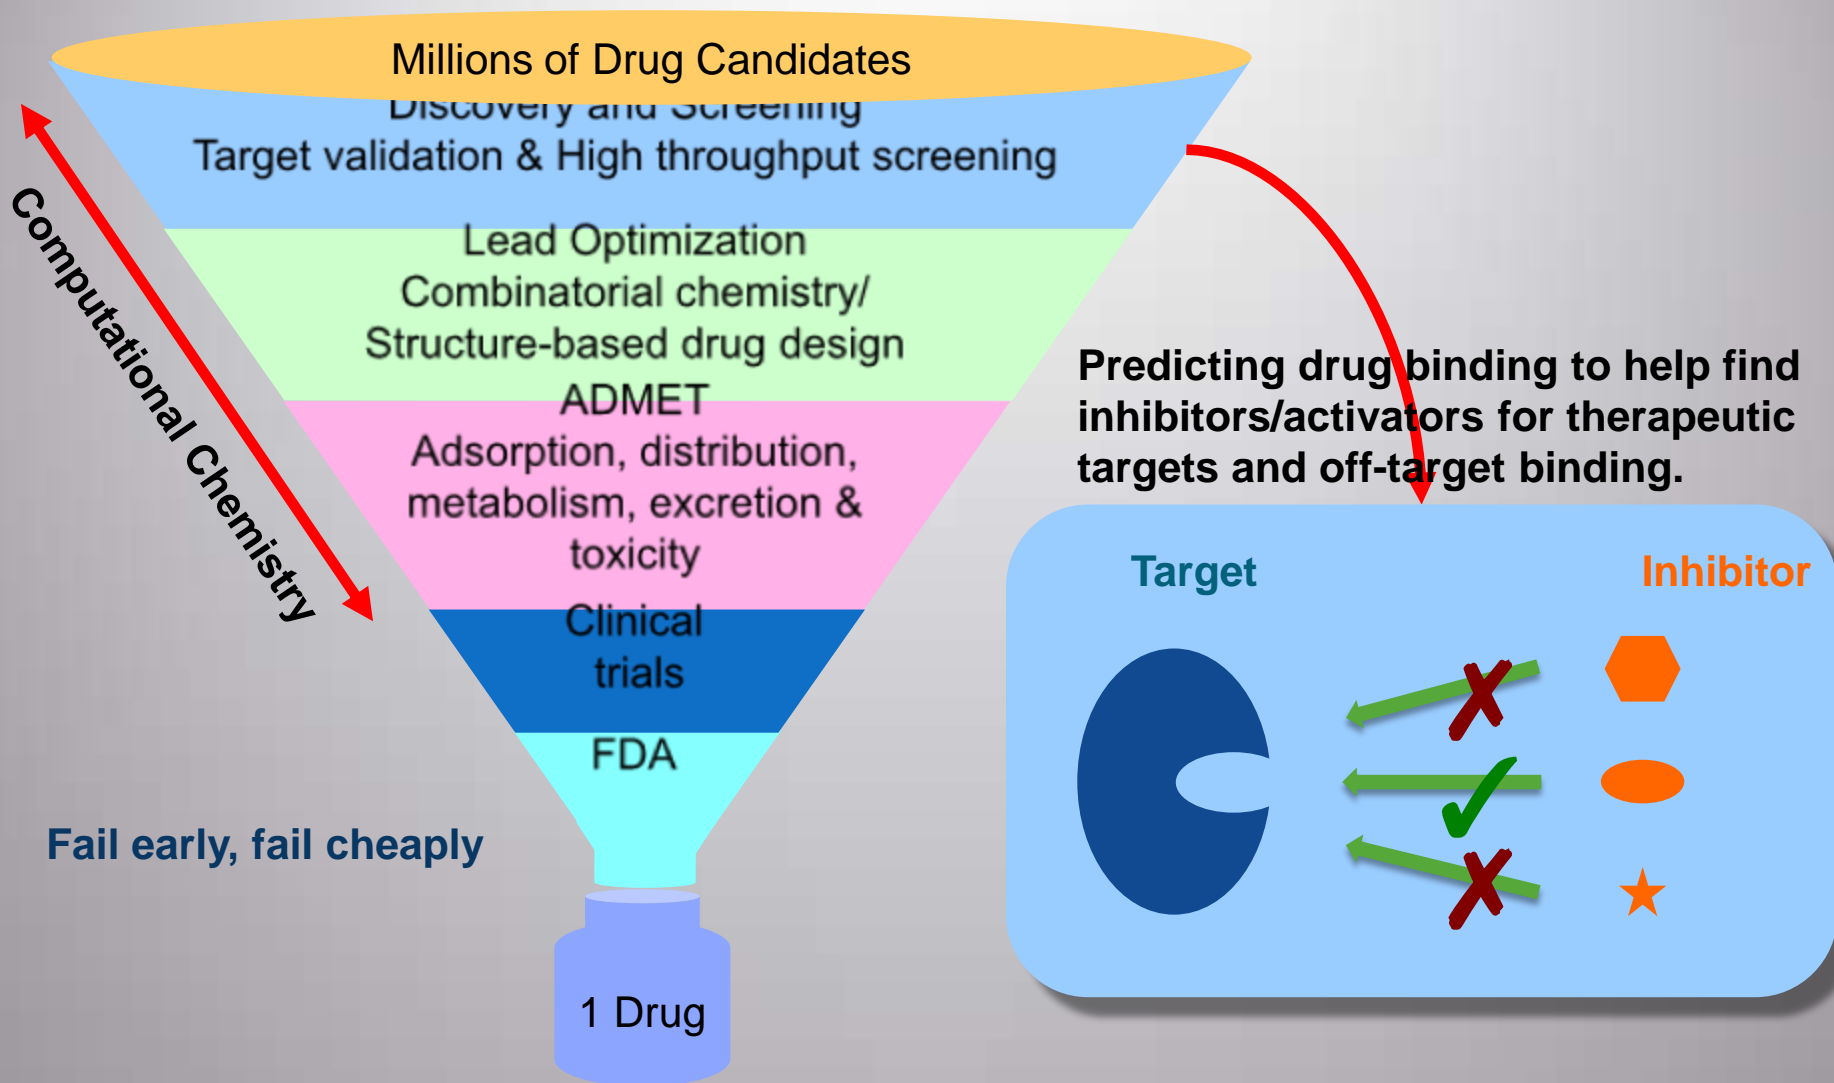
Urgent need for better *in-silico* prediction tools to reduce the cost and failure rate of drug discovery

- 7-15 years and >1\$B USD per drug
- >90% failure at late stage
- Post-market failure due to adverse drug effect (e.g. Vioxx, Lipobay).



Cook et al Pharmacoeconomics 2010, 28 (5) 355-363

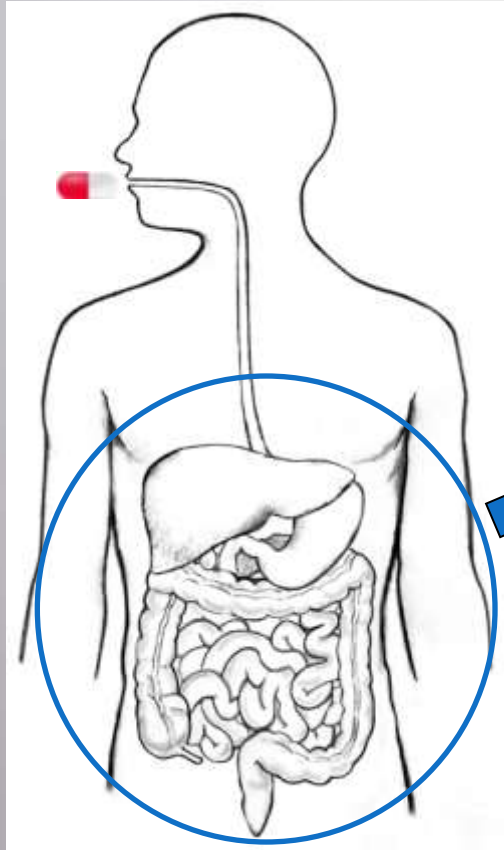
What can computational chemistry do for drug design?



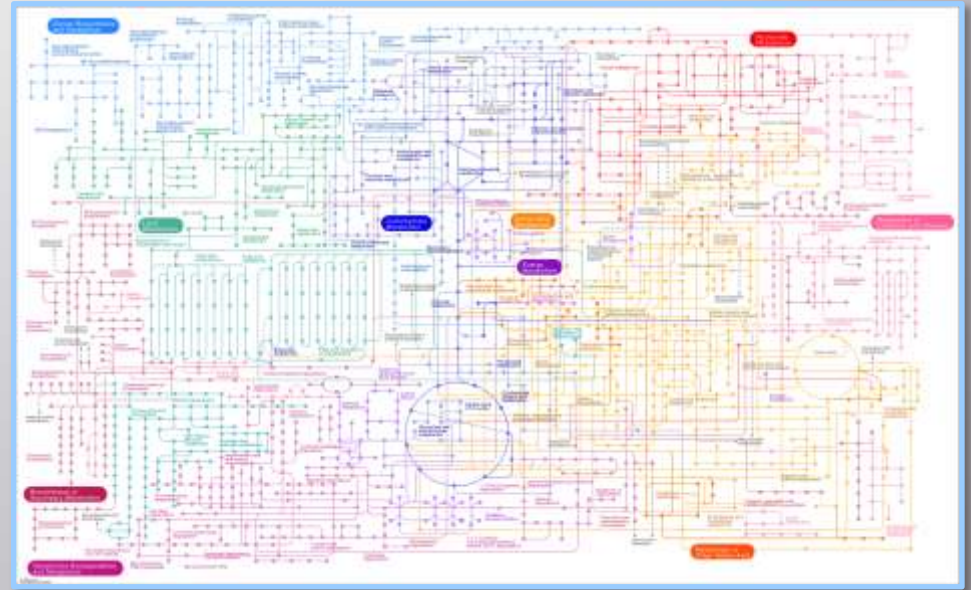
Fail early, fail cheaply

1 Drug

Drug on- and off-target binding along its metabolic pathway



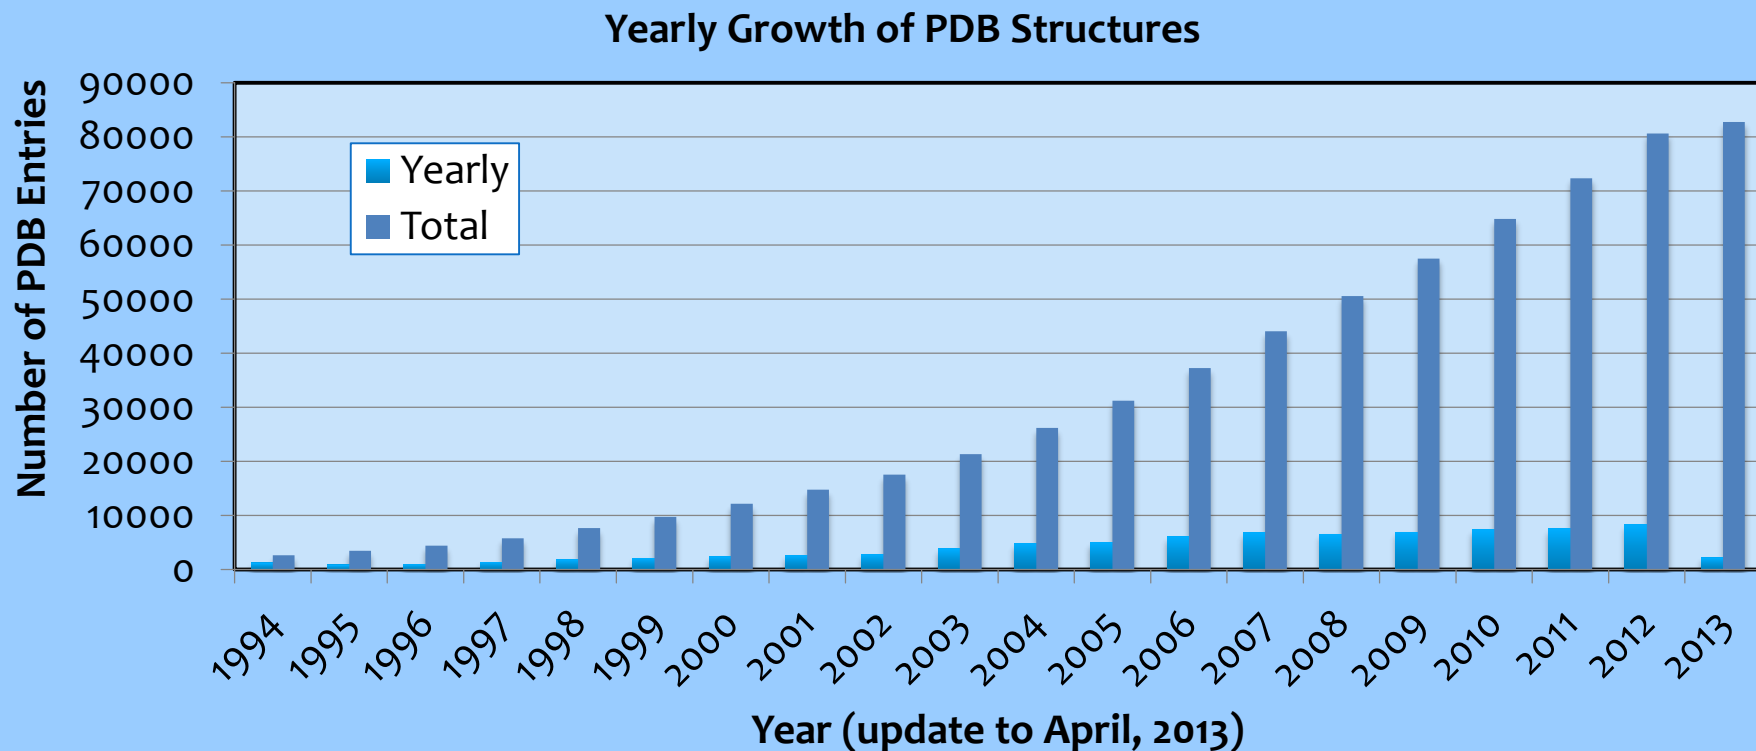
Metabolic Pathways



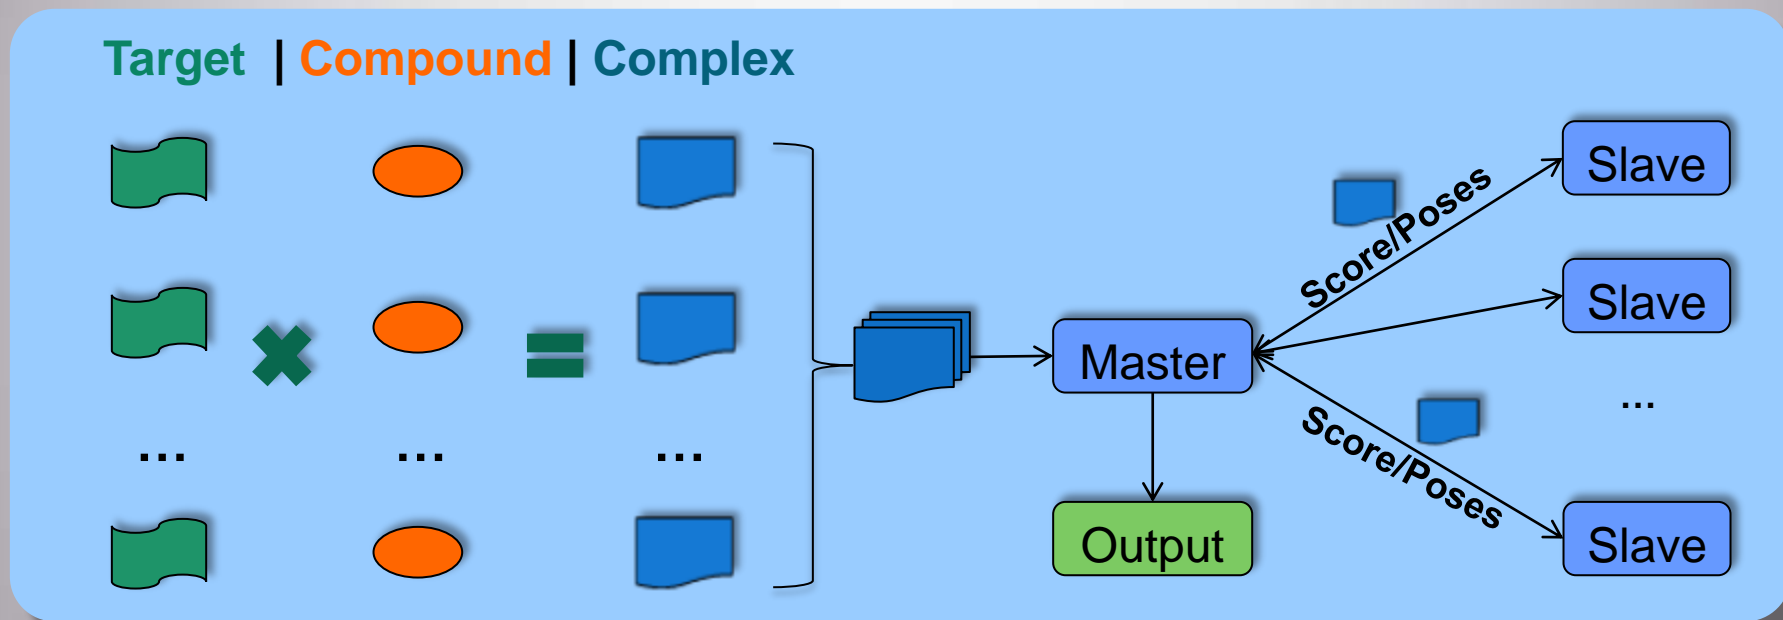
Drug could hit unexpected off-targets and cause severe side effects

Need for high performance computing in drug design to help manage growing complexity and improve productivity

- 1 ligand docking into 1 target takes ~1 min. Rescore 1 pose takes >10 min.
- Drug-like compounds $\sim 10^{60}$ possibilities.
- Drug-able targets $\sim 10\%$ human genome / off-target prediction.



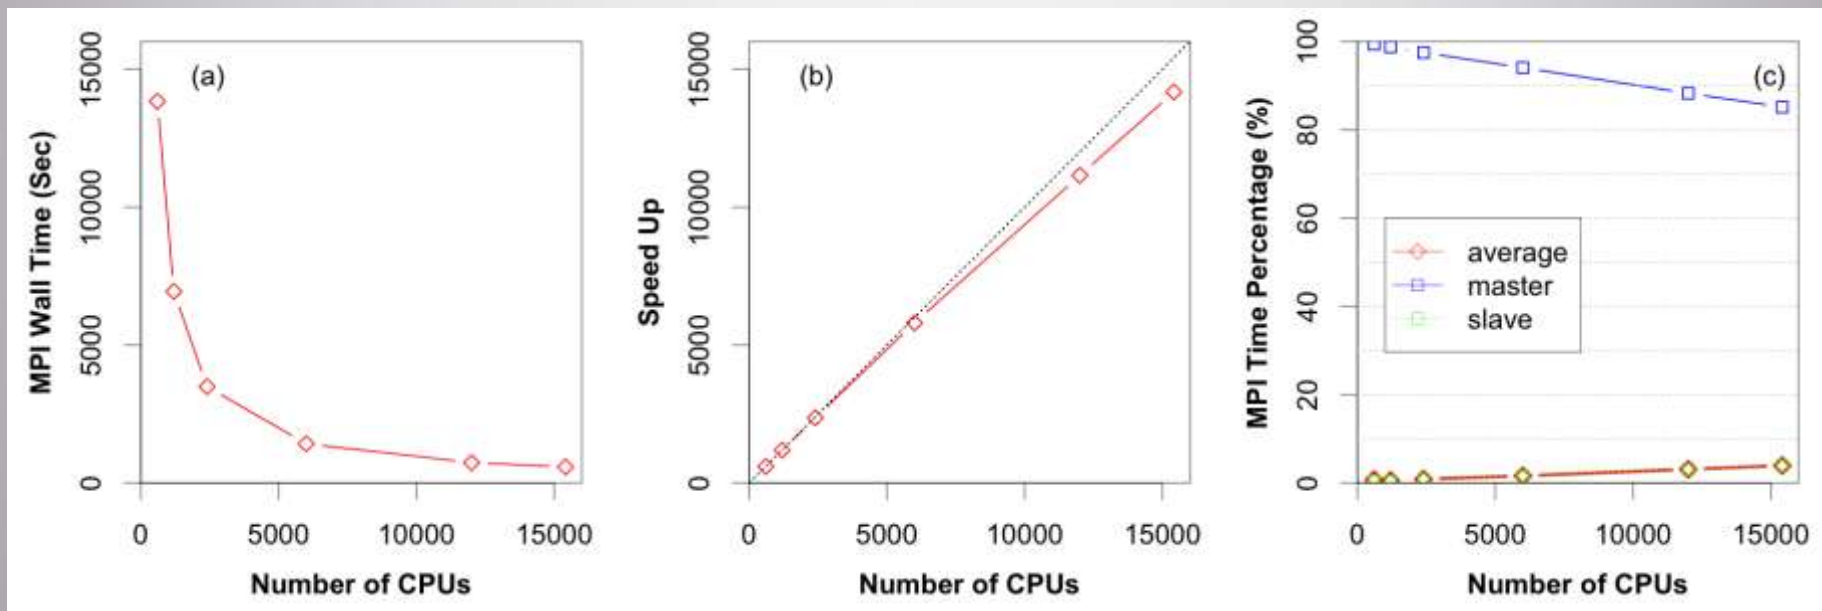
Start with a popular PC-based molecular docking program, AutoDock Vina, and parallelize it



HPC:	~10 min @ 1 CPU	~1 hr @ 100 CPU	~1 hr @ 700 CPU	~5 hr @ 15K CPU
PC:	~10 min	~4 day	~1 month	~10 year

An example for a target (~4K atoms) with ~40K ligands

Resultant VinaLC docking program has great parallel performance



Scale to a large set of CPUs (> 15K cores) with low overhead (<4%).
1 million flexible compound docking calculations took only 1.4 hours.

The VinaLC docking study has been published, highlighted on the JCC cover and reported on HPC wire



Leading HPC Solution Providers

HPC Wire

Since 1986 - Covering the Fastest Computers in the World and the People Who Run Them

Select Language | Transaction Decoder | Subscribe | Sign In | Search

Home | News | Topics | Sectors | Resources | Special Features | Market Watch | Events | Job Bank | About

People to Watch 2013
January 31, 2013
The Week in HPC Research
Page 1 | 2 | 3 | 4 | 5

MPI and Multithreading for Molecular Docking

A group of researchers from Lawrence Livermore Lab has published a new paper in the January 2013 edition of Journal of Computational Chemistry on the subject of MPI and multithreading for molecular docking. More specifically, they are focused on the parallel molecular docking of large databases on petascale computers.

The researchers adopted a mixed parallel scheme of both message passing interface (MPI) and multithreading implemented in the AutoDock Vina molecular docking program. The resulting program, named VinaLC, was run on the petascale systems at Lawrence Livermore National Laboratory.

"To exploit the typical cluster-type supercomputers, thousands of docking calculations were dispatched by the master process to run simultaneously on thousands of slave processes," the authors write. "Each docking calculation takes one slave process on one node, and within the node each docking calculation runs via multithreading on multiple CPU cores and shared memory."

Using this approach, the code was able to scale to more than 15K CPUs with a low overhead cost of 5.84%. One million docking calculations were processed in just 1.4 hours.

Speed Up

Number of CPUs

Ethernet
The high performance alternative to InfiniBand. Need Proof? Click Here. Chelsio Communications

Short Takes
Pittsburgh Supercomputing Center Welcomes 'Sharlock' Feb 01, 2013 | The Pittsburgh

VinaLC now available for download at <http://catsid.llnl.gov/vinalc>

What can we do for you?



For more information on the HPC Innovation Center and to explore engagement options, please visit www.hpcinnovationcenter.com and contact:

Jeff Wolf

*Deputy Director and
Chief Business Officer
HPC Innovation Center
925-422-7894
wolf28@llnl.gov*

LLNL-PRES-643607

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC



HIGH PERFORMANCE COMPUTING
INNOVATION CENTER