IBM Research

Computational Science Challenges and Approaches

Ulisses Mello & Jim Sexton
Computational Sciences Center
IBM T. J. Watson Research Center, NY

Some Quotes:

- **A Science Based Case for Large-Scale Simulation**
  – Office of Science, US Department of Energy, July 2003
  – "Computational simulation offers to enhance, as well as leapfrog, theoretical and experimental progress in many areas of science … Successes have been documented in such areas as advanced energy systems (e.g. fuel cells, fusion), biotechnology (e.g. genomics, cellular dynamics), nanotechnology (e.g. sensors, storage devices), and environmental modeling (e.g., climate prediction, pollution remediation)."
- **"The New Secret Weapon"**
  – U.S. Council for Competitiveness
  – "Supercomputing is part of the corporate arsenal to beat rivals by staying one step ahead of the innovation curve"
  – "Modeling, simulation and massive data analysis are the next huge game changing drivers for innovation".

Mark Seager: “The Scientific Method has changed as simulation has become the integrating element with Theory and Experiment”

Predictive simulation ENABLES

- **Detailed predictive assessment of complex models for overarching physical problems**
- **Design of experiments**
- **Impact assessment of policy choices**
- **Elimination of costly physical prototypes**

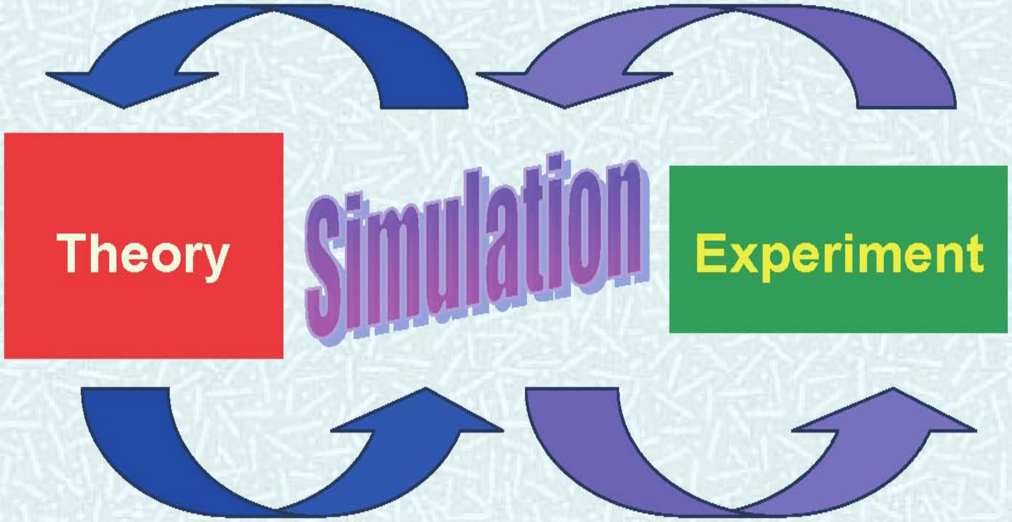

Predictive simulation REQUIRES

- Verification and validation of complex models (**experiment**)
- Development of science based models (**theory**)
- **Databases** of physical properties and catalogues of scientific data
- **Petascale simulation environments**

Revolution in the making: BlueGene/L at LLNL

IDC Technical Computing Systems Forecast

Bio Sci	Genomics, proteomics, pharmacogenomics, pharma research, bioinformatics, drug discovery.
Chem Eng	Chemical Engineering: Molecular modeling, computational chemistry, process design
CAD	Mechanical CAD, 3D Wireframe – mostly graphics
CAE	Computer Aided Engineering – Finite Element modeling, CFD, crash, solid modeling (Cars, Aircraft, …)
DCC&D	Digital Content Creation and Distribution
Econ Fin	Economic and Financial Modeling, econometric modeling, portfolio management, stock market modeling.
EDA	Electronic Design and Analysis: schematic capture, logic synthesis, circuit simulation, system modeling
Geo Sci	Geo Sciences and Geo Engineering: seismic analysis, oil services, reservoir modeling.
Govt Lab	Government Labs and Research Centers: government-funded R&D
Defense	Surveillance, Signal Processing, Encryption, Command, Control, Communications, Intelligence, Geospatial Image Management. Weapon Design
Software Engineering	Development and Testing of Technical Applications
Technical Management	Product Data management, Maintenance Records management, Revision Control, Configuration Management
Academic	University Based R&D
Weather	Atmospheric Modeling, Meteorology, Weather Forecasting

IDC Forecasts Graphed

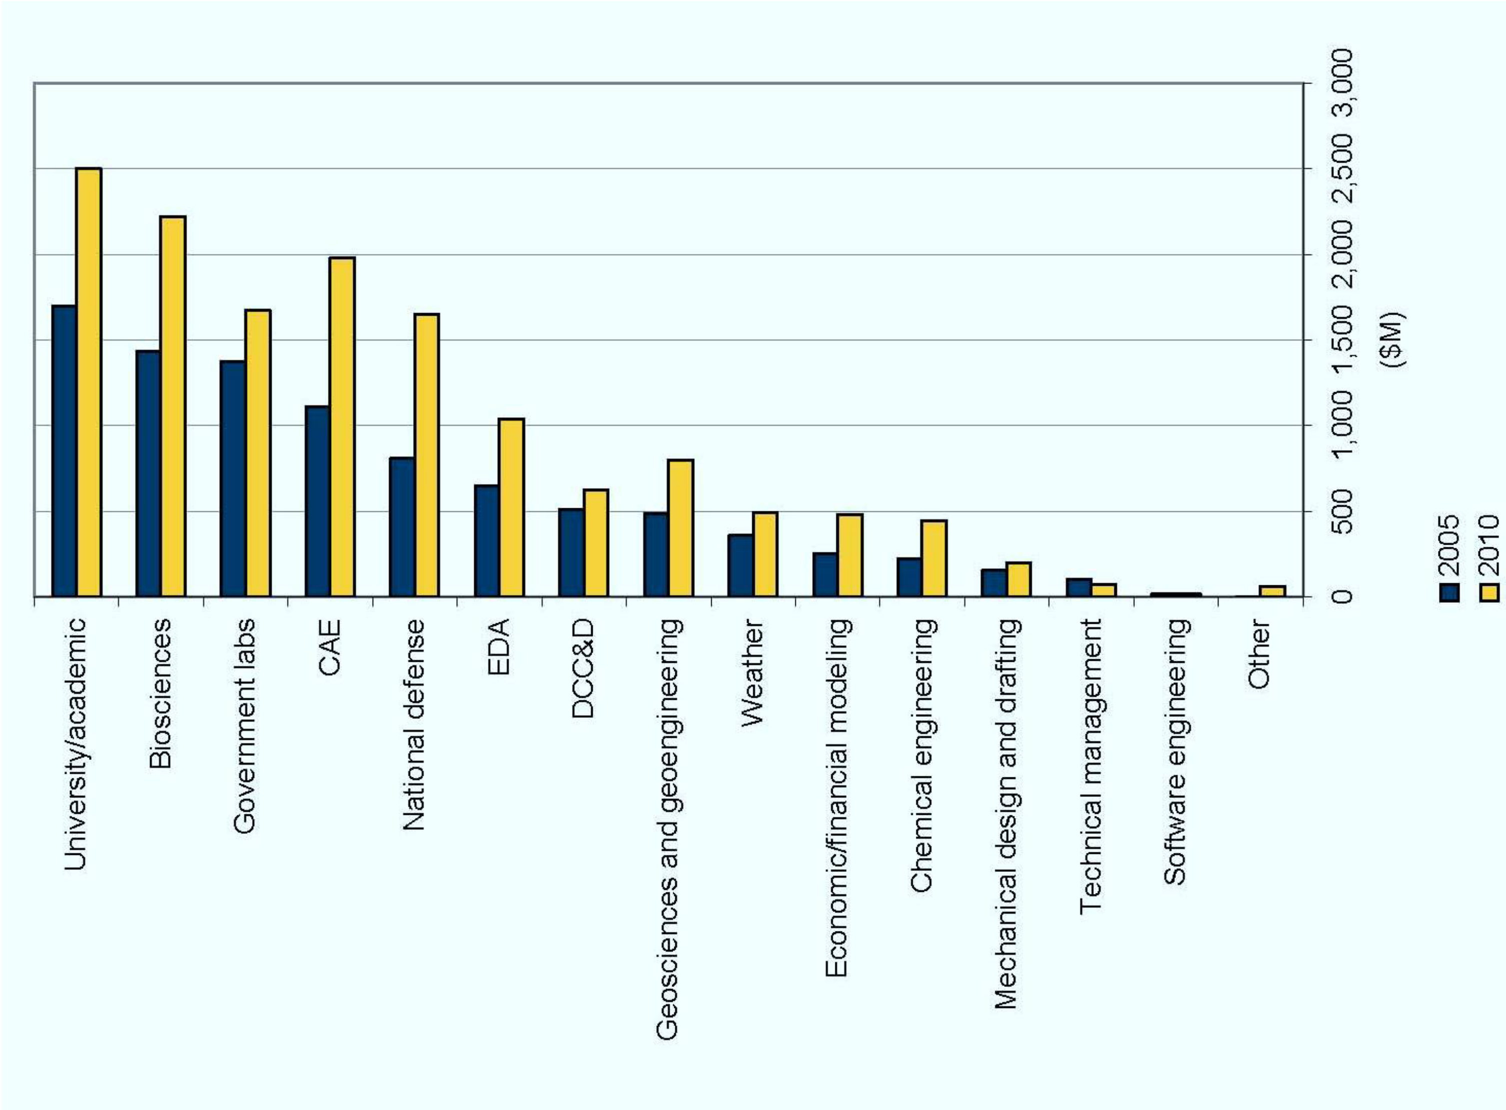

HPC Use Cases

- **Capability**
  - Calculations not possible on small machines
  - Usually these calculations involve systems where many disparate scales are modeled.
  - One scale defines required work per "computation step"
  - A different scale determines total time to solution.
- **Examples**
  - Protein Folding:
    - 10-15.secs – 1 sec
  - Refined grids in Weather forecasting:
    - 10km today -> 1km in a few years
  - Full Simulation of Human Brain

Useful only as proofs of concept

- **Complexity**
  - Calculations which seek to combine multiple components to produce an integrated model of a complex system.
  - Individual components can have significant computational requirements.
  - Coupling between components requires that all components be modeled simultaneously.
  - As components are modeled, changes in interfaces are constantly transferred between the components
- **Examples**
  - Water Cycle Modeling in Climate/Environment
  - Geophysical Modeling for Oil Recovery
  - Virtual Fab
  - Multisystem / Coupled Systems Modeling

Critical to manage multiple scales in physical systems

- **Understanding**
  - Repetition of a basic calculation many times with different model parameters, inputs and boundary conditions.
  - Goal is to develop a clear understanding of behavior / dependencies / and sensitivities of the solution over a range of parameters
- **Examples**
  - Multiple independent simulations of Hurricane paths to develop probability estimates of possible paths, possible strength,
  - Thermodynamics of Protein / Drug Interactions
  - Sensitivity Analysis in Oil Reservoir Modeling
  - Optimization of Aircraft Wing Design,

Essential to develop parameter understanding, and sensitivity analysis

Capability

Computational Needs of Technical, Scientific, Digital Media and Business Applications Approach or Exceed the Petaflops/s Range

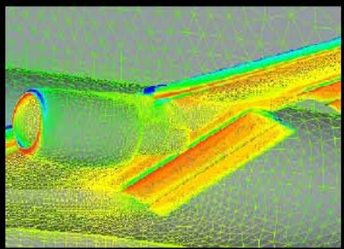

CFD Wing Simulation

512x64x256 Grid
(8.3 x 10e6 mesh points)
5000 FLOPs per mesh point,
5000 time steps/cycles
2.15 x 10e14 FLOPs

Source: A. Jameson, et al

CFD Full Plane Simulation

512x64x256 Grid
(3.5 x 10e17 mesh points)
5000 FLOPs per mesh point
5000 time steps/cycles
8.7x 10e24 FLOPs

Materials Science

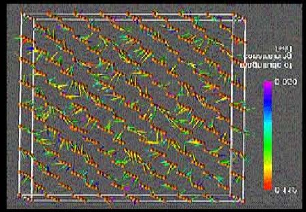

Magnetic Materials:
Current: 2000 atoms; 2.64 TF/s, 512GB
Future: HDD Simulation – **30TF/s, 2 TBs**
Electronic Structures:
Current: 1000 atoms;**0.5 TF/s, 250GB**
Future: 10000 atoms;**100TF/s, 2.5TB**

Source: D. Bailey, NERSC

Digital Movies and Special Effects

~ 1E14 FLOPs per frame
50 frames/sec
90 minute movie

- 2.7E19 **FLOPs**
-

~150 days on 2000 1 GFLOP/s CPUs

Source: Pixar

Spare Parts Inventory Planning

Modeling the optimized deployment of 10,000 part numbers across 100 parts depots and requires:

- 2 x 10e14 FLOP/s
  (12 hours on 10, 650MHz CPUs)
- 2.4 PetaFlop/s sust. performance
  (1 hour turn-around time)

Industry trend for rapid, frequent modeling for timely business decision support drives higher sustained performance

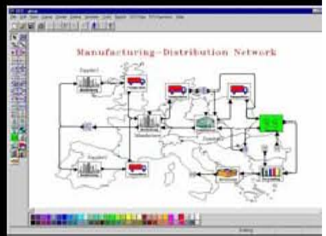

Source: B. Dietrich, IBM

Complexity: Modern Integrated Water Management

Adv Water Mgmt Reference IT Architecture

Partner Ecosystem
- Climatologists
- Environmental Observation Systems Companies
- Sensors Companies
- Environmental Sciences Consultants
- Engineering Services Companies.
- Subject Matter Experts
- Universities

Sensors
- Physical
- Chemical
- Biological
- Environmental
- In-situ
- Remotely sensed
- Planning and placement

Enabling IT
- HPC
- Visualization
- Data management

Physical Models
- Climate
- Hydrological
- Meteorological
- Ecological

Model Strategy
- Selection
- Integration & coupling
- Validation
- Temporal/spatial scales

Analyses
- Stochastic model & stats
- Machine learning
- Optimization

Historical – Present – Near future – Seasonal – Long term – Far future

Understanding: Hurricane Forecasts

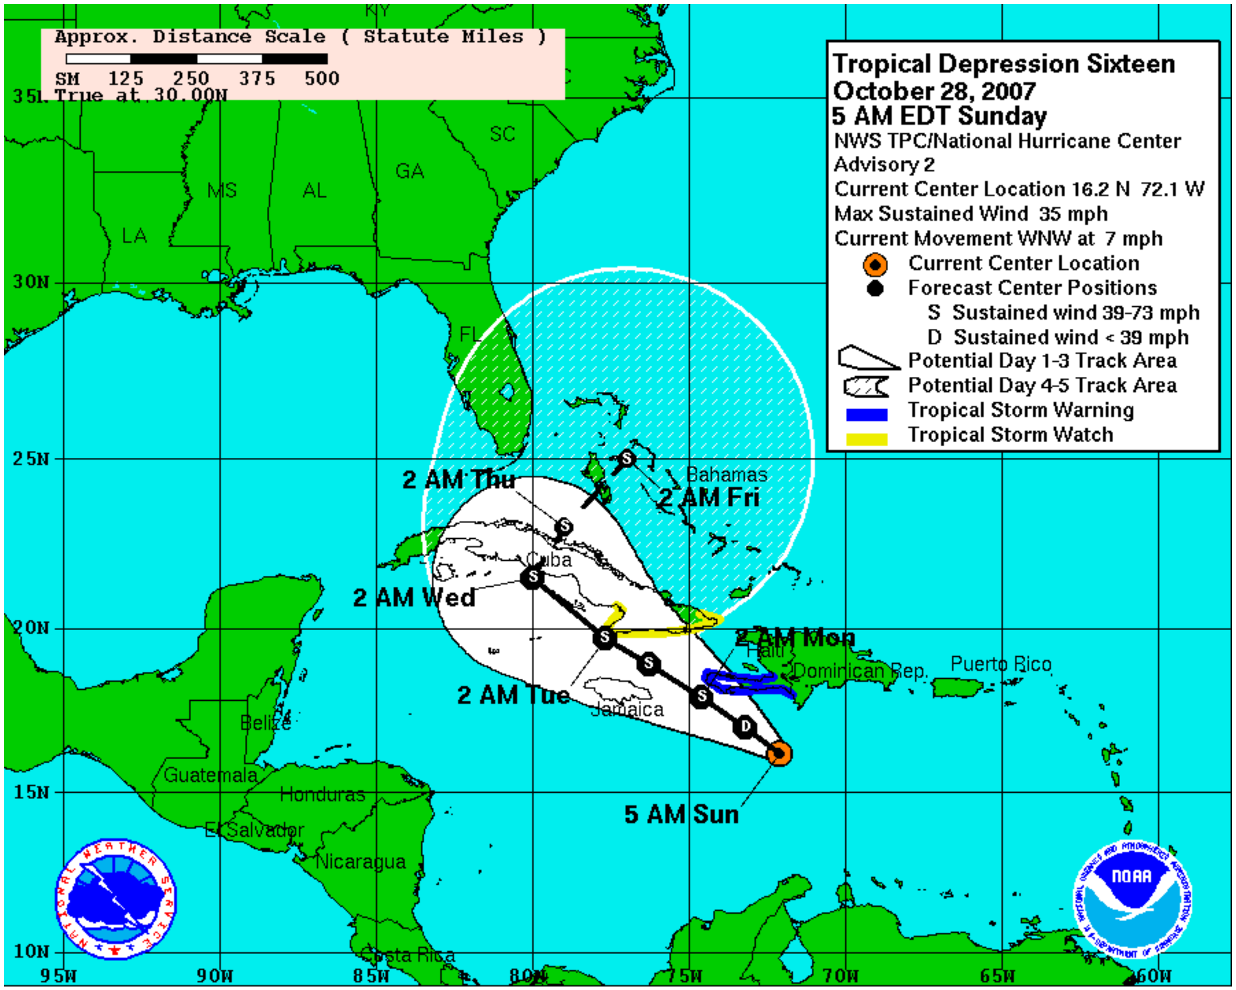

High Performance Computing Trends

- **Three distinct phases .**
  - Past: Exponential growth in processor performance mostly through CMOS technology advances
  - Near Term: Exponential (or faster) growth in level of parallelism.
  - Long Term: Power cost = System cost ; invention required
- **Curve is not only indicative of peak performance but also performance/$**

1PF: 2008

10PF: 2011

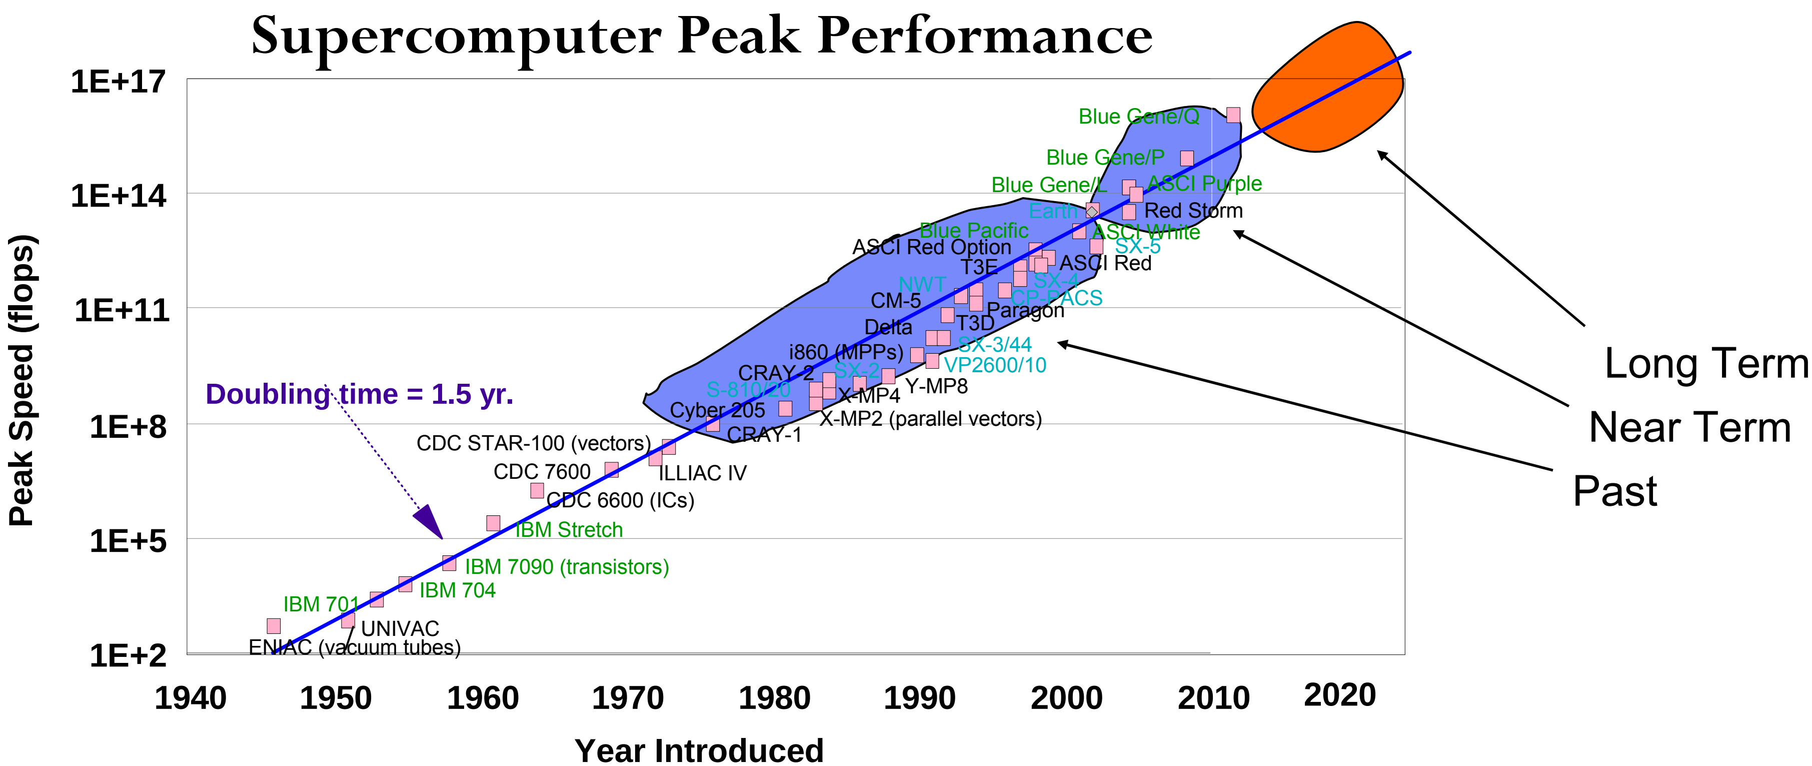

Architecture Challenges

- **Core Frequencies ~**
  - 2-4 GHz, will not change significantly as we go forward
  - 100,000,000 Cores to deliver an Exaflop
- **Power**
  - At today's MegaFlops / Watt: 2 GW needed
  - Power reduction will force simpler chips, longer latencies, more caches, nearest neighbour network
- **Memory and Memory Bandwidth**
  - Much less memory / core (price)
  - Much less bandwidth / core (power / technology)
- **Network Bandwidth**
  - Much less network bandwidth per core (price / core)
  - Local network connectivity
- **Reliability**
  - Expect algorithms / applications will have to permit / survive hardware fails.
- **I/O Bandwidth**
  - At 1 Byte / Flop, an EXAFLOP system will have 1 EXABYTE of Memory.
  - No disk system can read / write this amount of data in reasonable time.

- **Exascale Computing**
  - O(100 M) compute engines working together
- **Capability delivered has the potential to be truly revolutionary**
- **However**
  - Systems will be complex
  - Software will be complex
  - Applications will be complex
  - Data Centers will be complex
  - Maintenance / Management will be complex

An integrated approach needed

Research, Industrial, Business Problems

Applications

Computational Sciences

Algorithms
Messaging
Programming
Libraries
Performance
Analysis

Deep Computing Computer Sciences

Data
Workflow
Visualization
Virtualization
Integration
Monitoring

Cyber Infrastructure

Integrated Research Programs

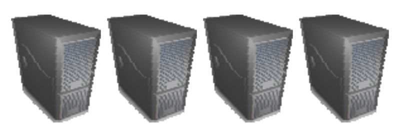

Research Themes

- **Applications**
  - Research to enable applications by engaging very closely with teams requiring deep computing solutions
  - In this application work, the Center will identify the computational and computer sciences research challenges required for enablement of these deep computing solutions
- **Computational Science**
  - Research in fundamental algorithms, in programming approaches, in messaging approaches, in library development, in applications performance for deep computing solutions
- **Deep Computing Computer Science**
  - Research into those aspects of computer science which can enable better management and usage of data, of work flow, and of systems for deep computing solutions

An Integrated Approach

- **Core research in deep computing solutions enablement**
  - To help develop deep computing ecosystem
  - Training for graduate students, research staff and cyberinfrastructure staff.
  - General systems and applications analysis and assessments to contribute to infrastructure planning, program planning,
  - A resource for expertise in deep computing solutions.
- **Work intimately with research teams requiring deep computing solutions**
  - Engage those research teams to help identify and implement best practice for development, deployment, production and analysis.
  - Engage in application enablement research to advance best practice deployment of deep computing solutions
  - Identify / extract general research themes
- **Work intimately with the core facilities teams providing Cyber Infrastructure.**
  - Provide expertise / training / best practice to those teams
  - Identify, assess, research and resolve challenges

Applications Component

- **Parallelization**
  - Parallelize a code or code base (both multicore and mpp)
- **Scaling**
  - Assess / restructure codes to improve scalability
- **Performance**
  - Identify performance bottlenecks, research solutions for those bottlenecks
- **Algorithm Development**
  - As necessary identify algorithmic issues and research solutions
- **Data management**
  - All issues associated with data management, loading / unloading data to an application, workflow restructuring to improve data flow,
- **Integration**
  - Identifying how best to optimize a work flow on available systems. Research possible infrastructure changes which might improve work flow
- **Design / Implementation**
  - If needed, consult with a given research team on design and implementation of codes

Examples

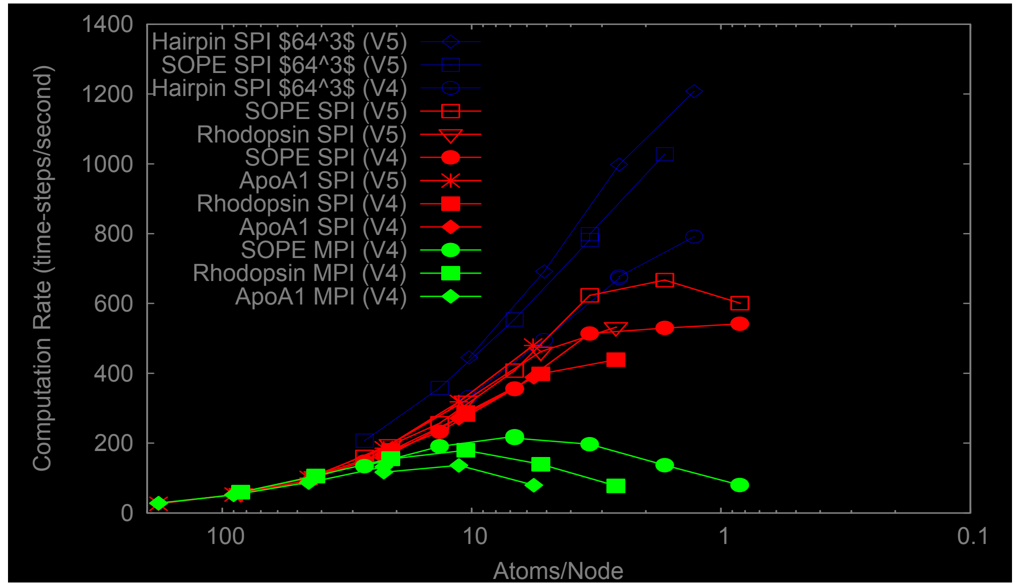

Blue Matter: Protein Molecular Dynamcs
Scalable to O(1) Atom / Node (Internal)

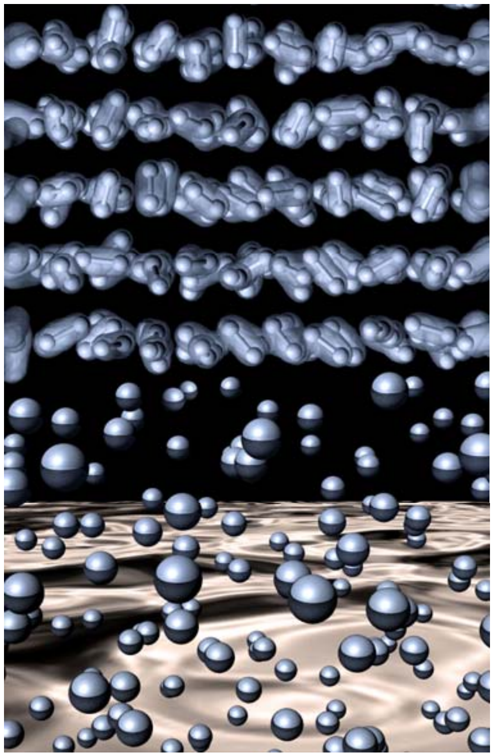

QBox: Quantum MD
200+ TF sustained
(With LLNL/Gigi)

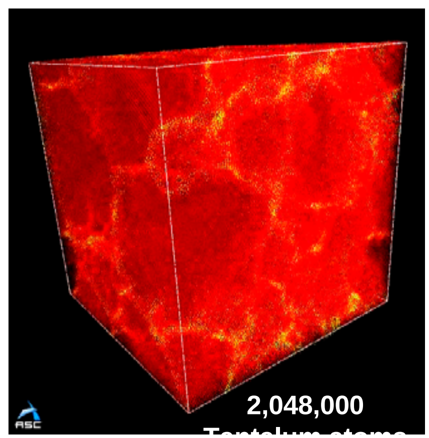

ddCMD: short range
Classical MD
100+ TF sustained

(With LLNL / Streitz)

Example Challenge: Basin Simulation Optimization

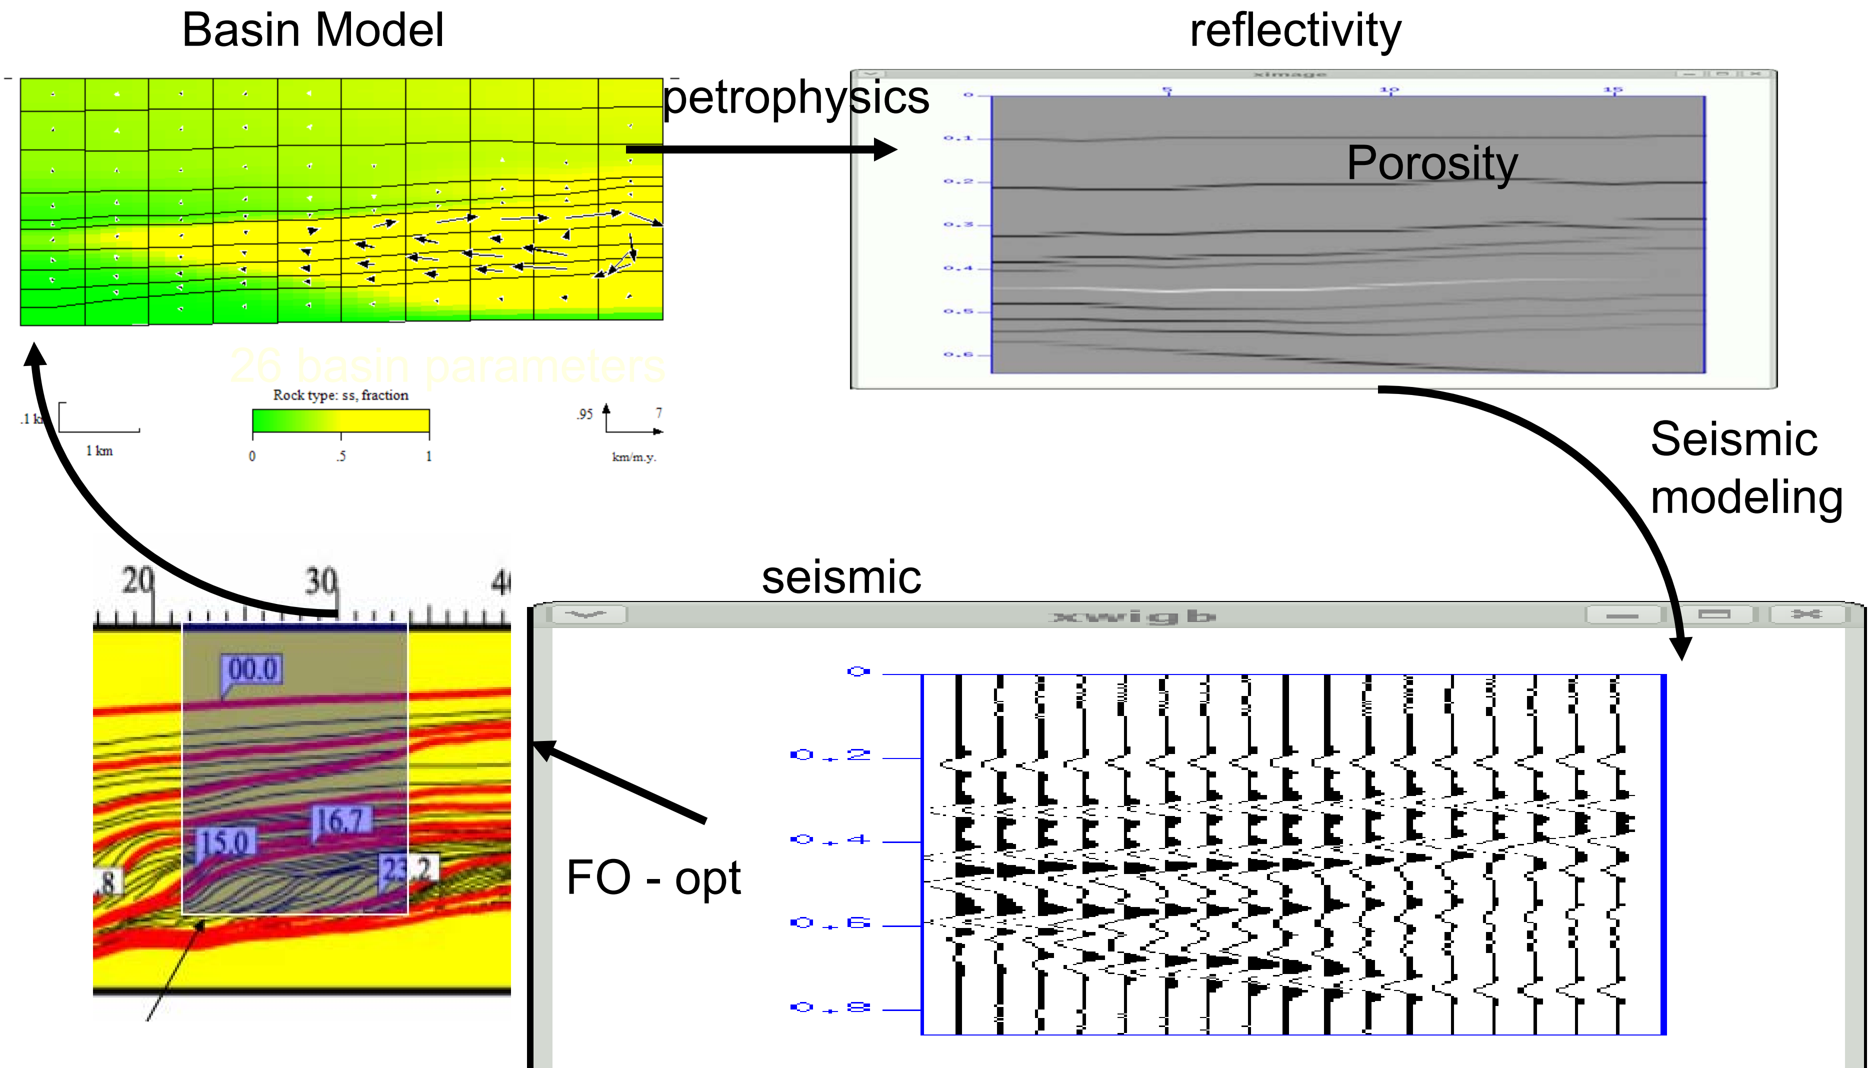

Basin Simulation

Primary variables $m_b = P, T, S_w, \ldots$

Petrophysics (Biot-Gassmann)

densities, velocities $m_p = \rho, Vp, Vs$

Seismic Modeling (AVO/2D FE)

amplitude, incidence $m_s = \{R, e\}$

Optimization

$$F = Min(R_{obs} - R_{sim})^2$$

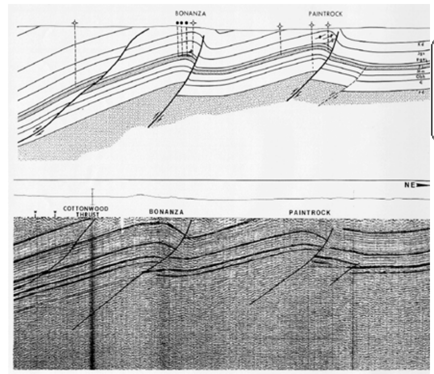

$$\frac{dm_p}{dm_b} = ?$$

$$\frac{\partial m_s}{\partial m_p} = ?$$

$$\frac{\partial F}{\partial m_s} = ?$$

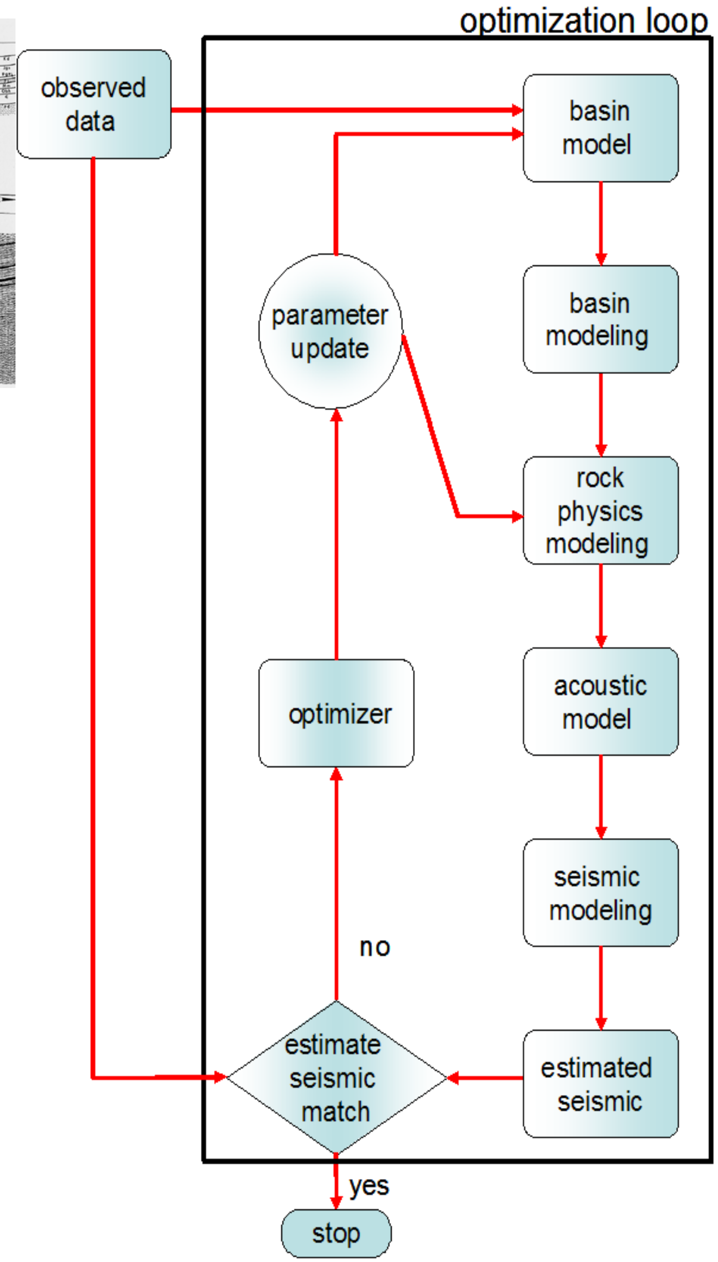

Computational Science Component

- **Algorithm Research**
  - Dense and Sparse Linear Algebra Algorithms
  - Optimization
  - Partitioning / Load Balancing
  - Graph Partitioning
  - Algorithms for Molecular Dynamics, CFD, …
- **Library Development for Multicore / Many Core**
- **Libraries for efficient, parallel IO**
- **Interprocessor communications paradigms**
  - MPI, ARMCI, Global Arrays, …
- **Platform agnostic programming approaches**
  - Is it possible to develop a code base which is portable across emerging architectures
- **Performance analysis for key applications.**
  - Study on current platforms
  - Project to future platforms

Examples

- **Dense Linear Algebra API's and Implementations for Multicore**
  - key benchmarks (e.g. Linpack),
  - in ab-initio quantum systems,
  - in all codes using BLAS
- **Communications Optimization**
  - Applications in QCD (bypass MPI and implement on Hardware APIs)

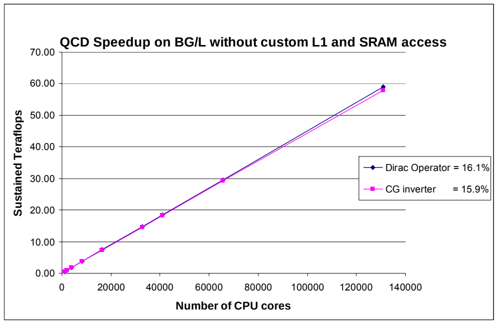

- **Development of new Lithography solution for 22nm Silicon Fabrication**
  - New code base in development
  - Don't yet know optimum computational platform, so working to design in a platform agnostic way

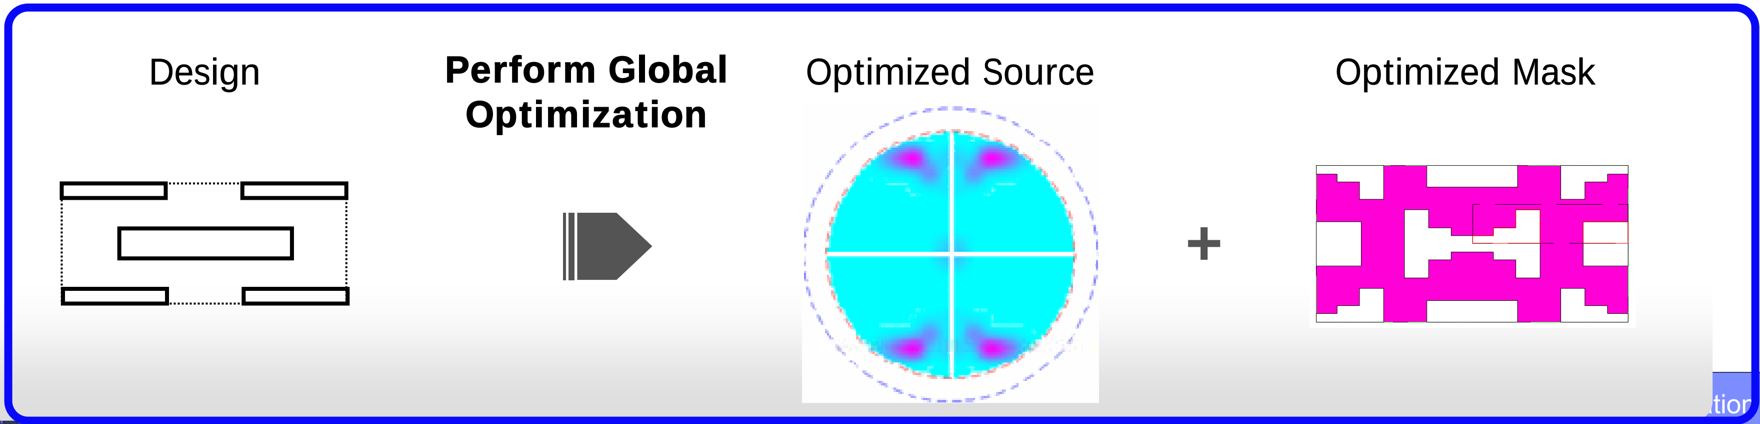

Deep Computing Computer Science Component

- **Use models**
  - Heterogeneous computing paradigms
  - Acceleration paradigms
  - Ultrascale High Throughput Computing
- **Data management**
  - In-memory work flow management solutions
  - As data sets get larger, IO between
  - In memory data bases
- **Visualization**
  - Deep Computing Visualization, Scalable Visual Networking, Remote Visual Networking
- **Management of compute center infrastructure**
  - Virtualization, automated service management, risk & resilience management, service delivery management

Examples

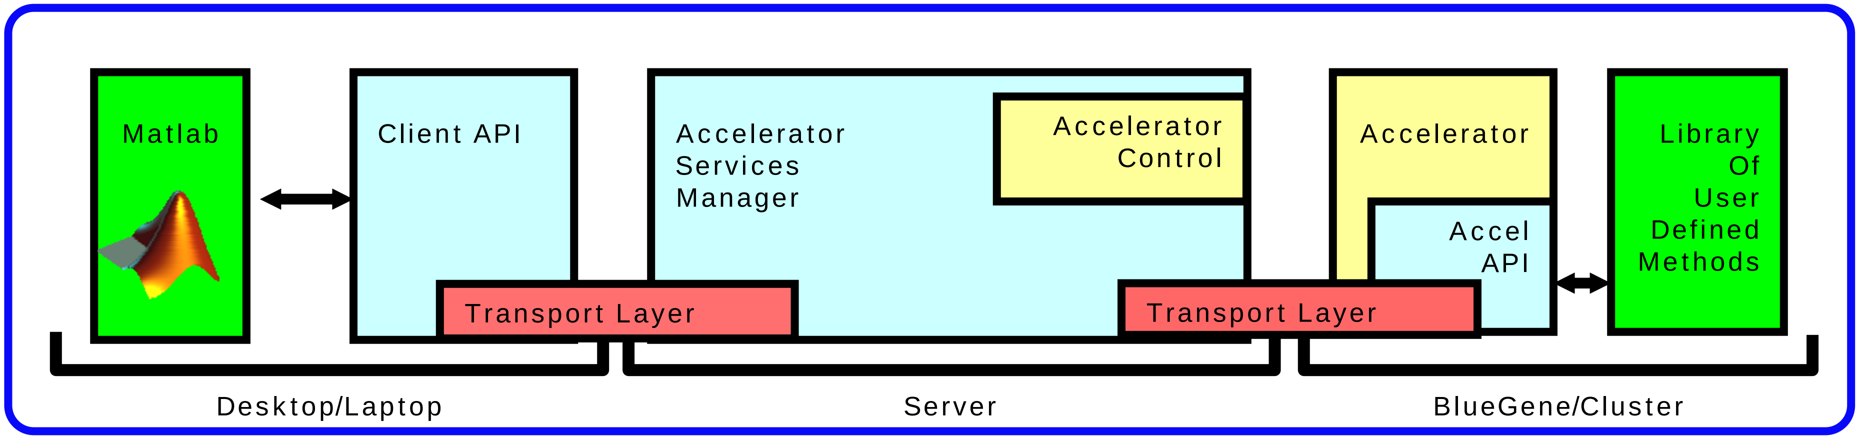

System Level Acceleration:
Desktop application transparent access to HPC Capability

In Memory DB Search Primitive:
Pointer Chasing
13Billion chases / sec on 8 TB
Distributed data set

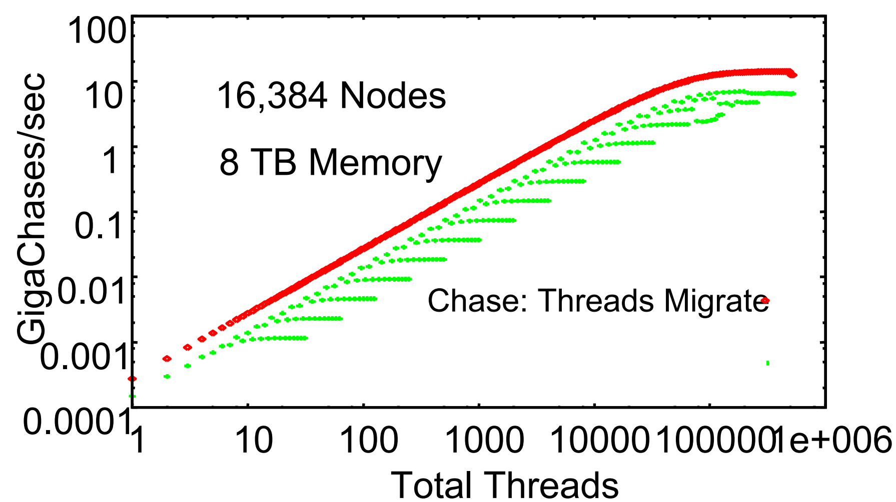

Summary

- **Emerging HPC landscape is extremely complex**
- **A time of extraordinary potential**
  – Game-changing capability is now available
- **A time of significant challenge**
  – Just as HPC starts to have real scientific and industrial impact - it gets extraordinarily hard.
- **A radical research and development approach required**
  – Multidisciplinary from domain to system design
  – Collaboration across research teams within universities
  – Collaboration between universities and industry