Chemistry beyond the petascale

Why is it necessary?
How do we get there?

Robert J. Harrison

harrisonrj@ornl.gov

robert.harrison@utk.edu

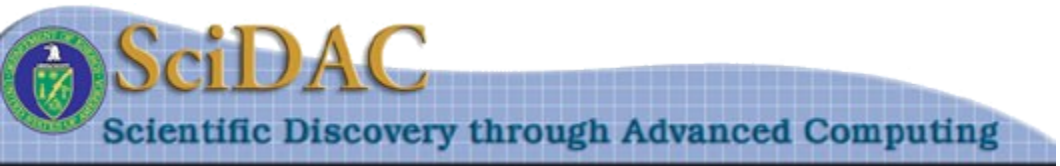

If you're not scared, you're not thinking big enough.

National Science Foundation
WHERE DISCOVERIES BEGIN

Robert J. Harrison
Oak Ridge National Laboratory,
University of Tennessee, Knoxville

SciDAC
Scientific Discovery through Advanced Computing

UT-BATTELLE
OAK RIDGE NATIONAL LABORATORY

Impact of sustained exponential growth

- We are only beginning to realize the transforming power of computing as an enabler of innovation and discovery.

- A characteristic of exponential growth is that we will make as much progress in the next doubling cycle as we've made since the birth of the field:
  - 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, ...

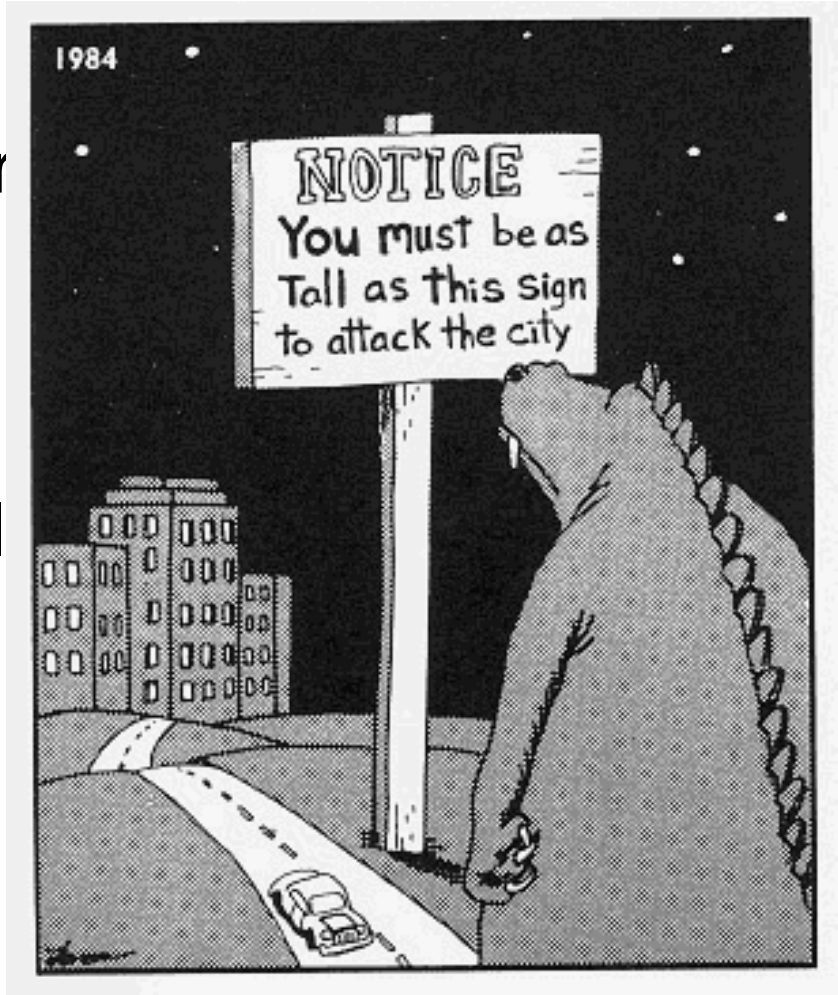

Computing now …

- The death of sequential computing
- Does anyone in the room still have a single cpu
  - Desktop computer?
  - Laptop?
  - Cell phone?

Computing in 2022 - I

- Looking back from 2007 to 1992
  - About 500x increase in desktop performance
    - 100MHz to 2x2 3GHz Core2
    - 30x from clock, 8x from parallelism
  - About 2500x increase in supercomputer speed
    - 100GF to 250TF
    - 30x from clock, 40x from parallelism

Computing in 2022 - II

- Looking forward to 2022
  - Expect same performance increases
  - Almost entirely from increased parallelism
  - Custom devices with much higher speed
  - Memory and I/O hierarchy much deeper

  - 20K * 2500 = 500M “processors”

Mission of the ORNL National Leadership Computing Facility (NLCF)

- field the most powerful capability computers for scientific research
- select a few time sensitive problems of national importance that can take advantage of these systems
- join forces with the selected scientific teams to deliver breakthrough science.

HPC Roadmap

Mission: Deploy and operate the computational resources needed to tackle global challenges	Vision: Maximize scientific productivity and progress on the largest scale computational problems
• Understanding the universe • Materials and nanoscience • Climate change and terrestrial sequestration of carbon • Sustainable energy • Clean and efficient combustion • Energy, ecology and security	• Providing world class computational resources and specialized services for the most computationally intensive problems • Providing a stable hardware/software path of increasing scale to maximize productive applications development

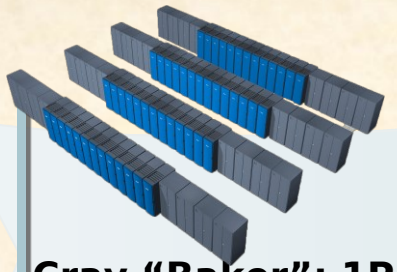

Cray "Baker": 1PF leadership class system for science

FY2009

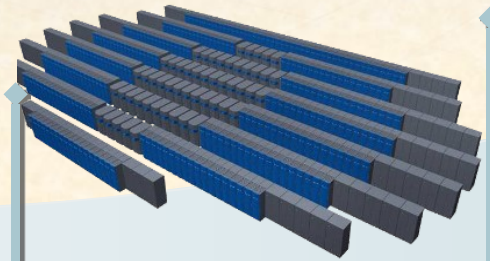

DARPA HPCS: 20 PF leadership class sustained PF system

FY2011

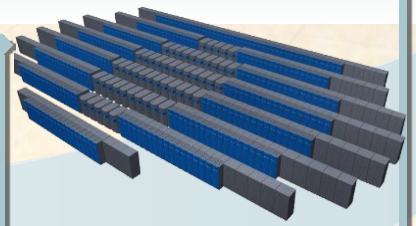

100-250 PF

FY2015

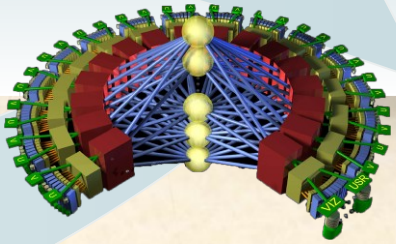

Future System: 1 EF

FY2018

10-year plan with required H/W, S/W, and "disruptive technologies" for applications

OAK RIDGE National Laboratory

The Jaguar Cray XT4 Leadership System

2007

- 11,508 compute nodes
  - Dual-core AMD Opteron processors with 4 GB memory
  - 23,016 compute cores
- 396 service & I/O nodes
- ~750 TB local storage
- 3D Torus interconnect
- 46 TB aggregate memory
- 119 TF peak performance

Today

- 7,832 compute nodes
  - Quad-core AMD Opteron processors with 8 GB memory
  - 31,328 compute cores
- 240 service & I/O nodes
- 900 TB local storage
- 3D Torus interconnect
- 63 TB aggregate memory
- >250 TF peak performance
- General availability to user community in May 2008

Cray “Baker” – 1 PF System

FY 2009: Cray “Baker”

- 1 Petaflops system
- 37 Gigaflops processor
- 27,888 quad-core processors Barcelona 2.3 GHz
- 2 GB per core; 223 TB total
- 200+ GB/s disk bandwidth
- 13,944 dual-socket 8-core SMP “nodes” with 16 GB
- 6.5 MW system power
- 150 Cabinets, 3,500 ft^2

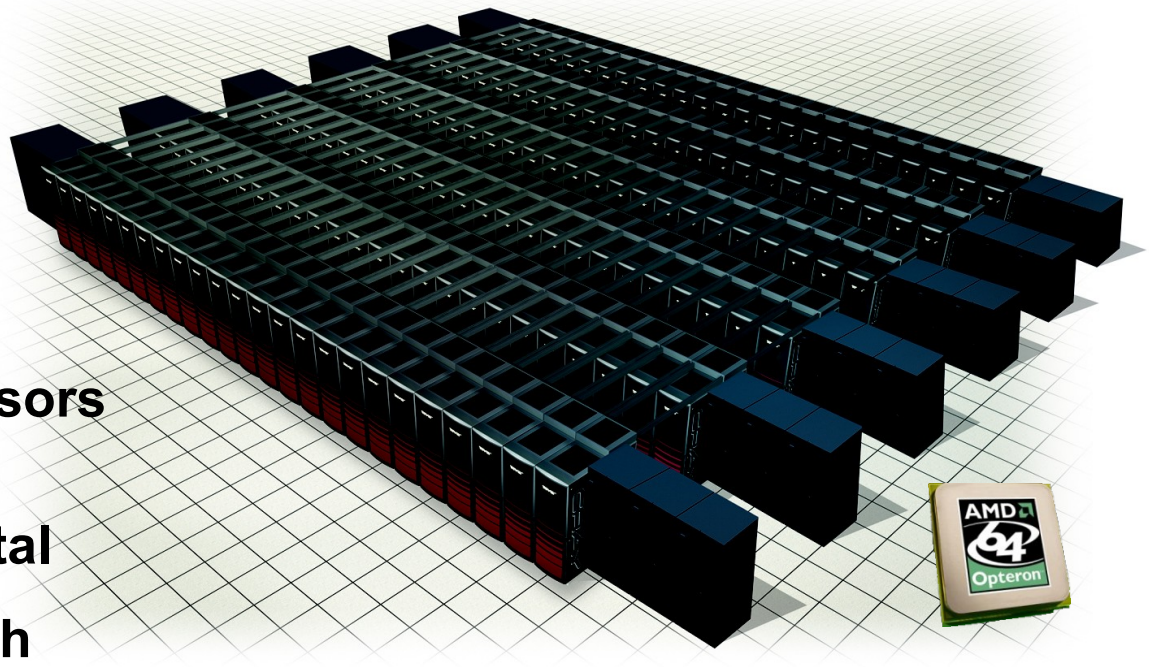

- Liquid cooled
- Compute node Linux operating system
- Torus interconnect

Univ. of Tennessee & ORNL Partnership
National Institute for Computational Sciences

- UT is building a new NSF supercomputer center from the ground up
  - Building on strengths of UT and ORNL
  - This center (NICS) goes operational by May 2008
- Series of computers culminating in a 1 PF system in 2009
  - Initial delivery (May 2008)
    - 4512 quad-core Opteron processors (170 TF)
  - Cray "Baker" (2009)
    - Multi-core Opteron processors; 100 TB; 2.3 PB of disk space

Current and Planned Data Centers

Open Science Center (40K ft²)

- Upgrading building power to 14 MW
  - Deploying a 6,600 ton chiller plant
  - Tripling UPS and generator capability

Multiprogram Research Facility (30K ft²)

- Capability computing for national defense
- Expanding to 20 MW of power and 6,000 ton chiller

Multiprogram Data Center (200K ft²)

- 100K ft² classified; 100K ft² unclassified
- Shared mechanical & electrical infrastructure
- Build out 25K ft² on each side as needed
- Lights out facility

ORNL Provides Leadership Computing to 2008 INCITE Program

- Allocation of computing resources to 30 programs in 2008 under the DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program (together with NERSC and PNNL).
- Leading researchers from government, industry, and the academic world will explore challenges including climate change, energy and alternative fuels on the center's leadership computers.
- This year's total allotment of processing hours nearly doubles that which ORNL provided in 2007.

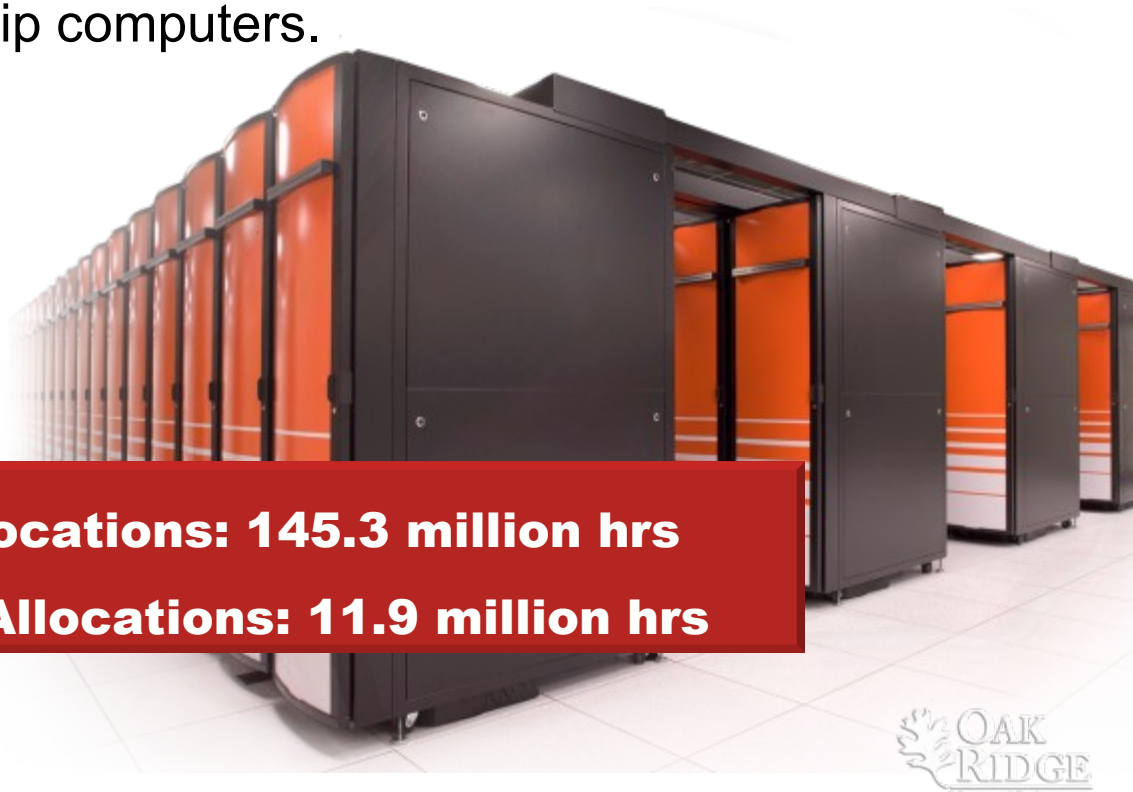

Project Allocations: 145.3 million hrs

Industrial Allocations: 11.9 million hrs

OAK RIDGE National Laboratory

Some Science Drivers

Science Domains	Science and Engineering Driver
Accelerator Physics	Optimize a new low-loss cavity design for the ILC
Astrophysics	Explosion mechanism of core-collapse supernovae and Type Ia supernovae
Biology	Can efficient ethanol production offset the current oil and gasoline crisis?
Chemistry	Catalytic transformation of hydrocarbons; clean energy & hydrogen production and storage
Climate	Predict future climates based on scenarios of anthropogenic emissions
Combustion	Developing cleaner-burning, more efficient devices for combustion.
Fusion	Plasma turbulent fluctuations in ITER must be understood and controlled
High Energy Physics	Find the Higgs particles thought to be responsible for mass, and find evidence of supersymmetry
Nanoscience	Designing high temperature superconductors, magnetic nanoparticles for ultra high density storage
Nuclear Energy	Can all aspects of the nuclear fuel cycle be designed virtually? Reactor core, radio-chemical separations reprocessing, fuel rod performance, repository
Nuclear Physics	How are we going to describe nuclei whose fundamental properties we cannot measure?

ORNL INCITE 2008 Allocations by Discipline

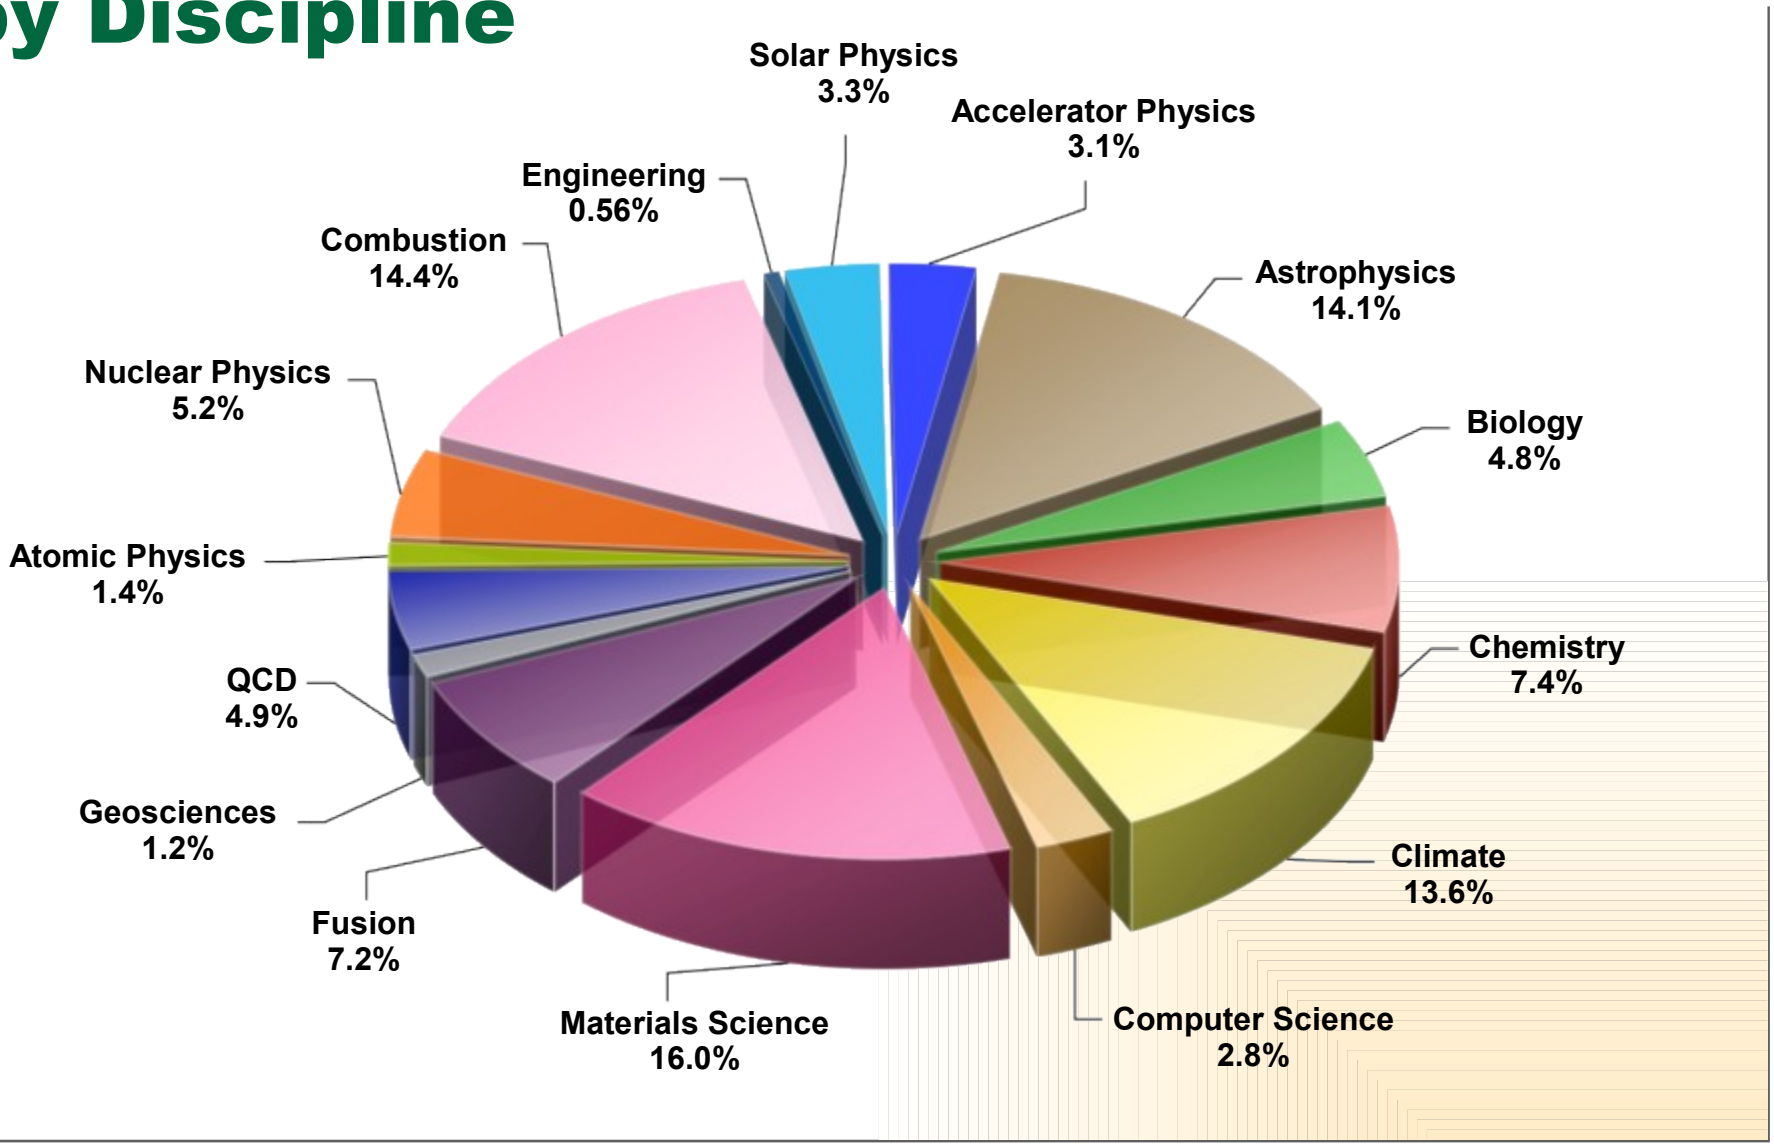

Chemical Computations on Future High-end Computers Award #CHE 0626354, NSF Cyber Chemistry

- NCSA (T. Dunning, PI)
- U. Tennessee (Chemistry, EE/CS)
- U. Illinois UC (Chemistry)
- U. Pennsylvania (Chemistry)
- Conventional single processor performance no longer increasing exponentially
  - Multi-core
- Many “corners” of chemistry will increasingly be best served by non-conventional technologies
  - GP-GPU, FPGA, MD Grape/Wine

Other technologies

- Field programmable gate arrays – multi TOP/s now

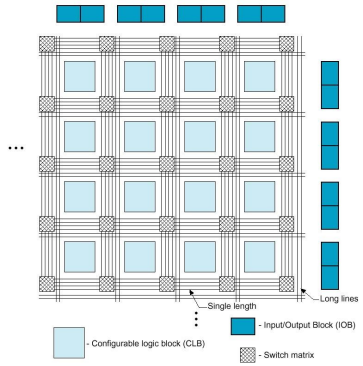

- General purpose graphical processor unit – 1TFLOP/s now
  - Lots of caveats on relevance

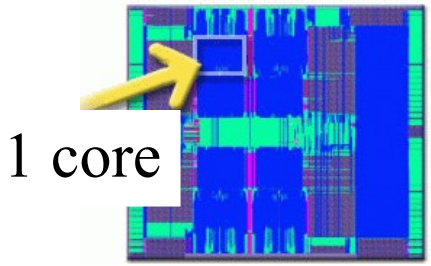

- Highly threaded devices
- FLOPs are cheap: bandwidth is expensive

O(1) programmers … O(10,000) nodes … O(100,000) processors … O(10,000,000) threads

- Complexity kills … sequential or parallel
- Expressing/managing concurrency at the petascale
  - It is too trite to say that the parallelism is in the physics
  - Must express and discover parallelism at more levels
  - Low level tools (MPI, Co-Array Fortran, UPC, …) don't discover parallelism or hide complexity or facilitate abstraction
- Management of the memory hierarchy
  - Memory will be deeper ; less uniformity between vendors
  - Need tools to automate and manage this, even at runtime

The way forward demands a change in paradigm

- by us chemists, the funding agencies, and the supercomputer centers

- A communal effort recognizing the increased cost and complexity of code development for modern theory at the petascale

- Re-emphasizing basic and advanced theory and computational skills in undergraduate and graduate education

Frameworks

- NWChem
- MADNESS
- ChemES – chemistry end-station
- Frameworks
  - Increase productivity; hide complexity
  - Interface disciplines
  - Capture knowledge
  - Open HPC to wider community
  - Expensive, communal projects

Computational Chemistry Endstation

International collaboration spanning 7 universities and 6 national labs

- Led out of UT/ORNL
- Focus
  - Actinides, Aerosols, Catalysis
- ORNL Cray XT, ANL BG/L

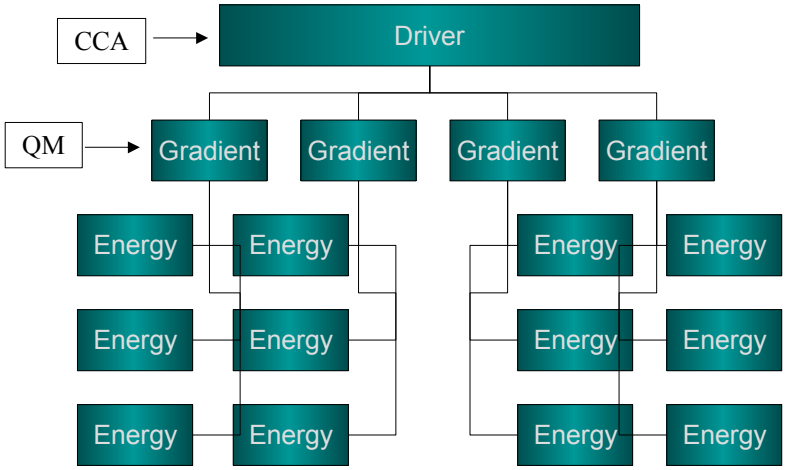

TL Windus

Capabilties:

- Chemically accurate thermochemistry
  - Many-body methods required
- Mixed QM/QM/MM dynamics
  - Accurate free-energy integration
  - Simulation of extended interfaces
- Families of relativistic methods

Participants:

- Harrison, UT/ORNL
- Sherrill, GATech
- Gordon, Windus, Iowa State / Ames
- Head-Gordon, U.C. Berkeley / LBL
- Crawford, Valeev, VTech.
- Bernholc, NCSU
- (Knowles, U. Cardiff, UK)
- (de Jong, PNNL)
- (Shepard, ANL)
- (Sherwood, Daresbury, UK)

Trends in Chemistry Codes

- Multi-level parallelism
  - Effective path for most applications to scale to 100K processors
  - Examples: coarse grain over vibrational degrees of freedom in numerical hessian, or geometries in a surface scan or parameter study
  - Conventional distributed memory within each subtask
  - Fine grain parallelism within a few processor SMP (multi-threads, OpenMP, parallel BLAS, …)
- Efficient exploitation of fine grain parallelism is a major concern on future architectures

NWChem Overview

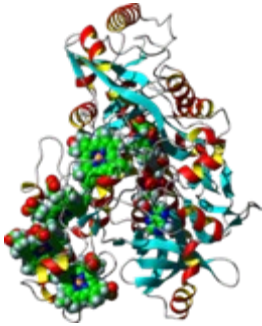

- Provides major modeling and simulation capability for molecular science
  - Broad range of molecules, including biomolecules, nanoparticles and heavy elements
  - Electronic structure of molecules (non-relativistic, relativistic, ECPs, first and second derivatives)
  - solid state capability (DFT plane-wave, CPMD)
  - Molecular dynamics, molecular mechanics

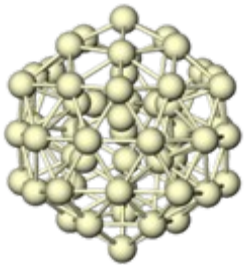

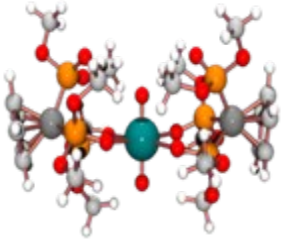

- Emphasis on modularity and portability.
- Freely distributed
- Performance characteristics – designed for MPP
- Portable – runs on a wide range of computers

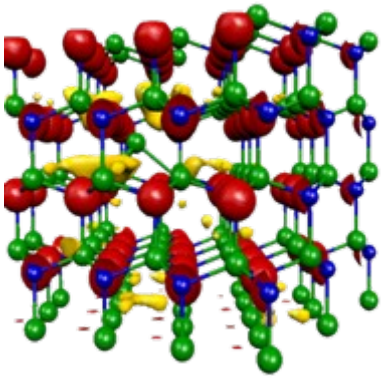

NWChem Molecular electronic structure

- The following quantum mechanical methods are available to calculate energies, analytic first derivatives and second derivatives with respect to atomic coordinates:
- Self Consistent Field (SCF) or Hartree Fock (RHF, UHF).
- Gaussian Density Functional Theory (DFT), using various (spin-restricted or unrestricted) exchange-correlation potentials
  - local and non-local (gradient-corrected),
  - meta-GGA,
  - hybrid (previous functionals plus HF Exchange)
- with formal N^4 and N^3 scaling (when auxiliary basis sets are used – RI-SCF or DFT-CDFit).
- Additional DFT capabilities:
  - Time-Dependent DFT.
  - Spin-orbit DFT (SODFT), using exchange-correlation potentials (spin-unrestricted).

Correlated Methods

- The following quantum mechanical methods are available to calculate energies and analytic first derivatives:
- MP2 including semi-direct using frozen core and RHF and UHF reference.
- Multiconfiguration SCF (MCSCF).
- The following methods are available to compute energies only. First and second derivatives are computed by finite difference of the energies.
- CCSD, CCSD(T), CCSD+T(CCSD), with RHF reference.
- Selected-CI with second-order perturbation correction.
- MP2 fully-direct with RHF reference.
- Resolution of the identity integral approximation MP2 (RI-MP2), with RHF and UHF reference.
- CIS, TDHF, TDDFT, and Tamm-Dancoff TDDFT for excited states with RHF, UHF, RDFT, or UDFT reference.
- TCE: Tensor Contraction Engine module, that can generate
  - CCSD(T) and CCSD[T] for closed- and open-shell systems (TCE module)
  - UCCD, ULCCD, UCCSD, ULCCSD, UQCISD, UCCSDT, and UCCSDTQ with RHF, UHF, or ROHF reference.
  - UCISD, UCISDT, and UCISDTQ with RHF, UHF, or ROHF reference.
  - Non-canonical UMP2, UMP3, and UMP4 with RHF or UHF reference.
  - EOM-CCSD, EOM-CCSDT, EOM-CCSDTQ for excitation energies, transition moments, and excited-state dipole moments of closed- and open-shell systems
  - CCSD, CCSDT, CCSDTQ for dipole moments of closed- and open-shell systems, plus … later in the presentation

Molecular electronic structure

- For all methods, the following operations may be performed:
- Single point energy
- Geometry optimization (minimization and transition state)
- Molecular dynamics on the fully *ab initio* potential energy surface
- Numerical first and second derivatives automatically computed if analytic derivatives are not available
- Normal mode vibrational analysis in cartesian coordinates
- ONIOM hybrid method of Morokuma and co-workers
- Evaluation of static, one-electron properties. (e.g. Spin-spin coupling , NMR shielding)
- For closed and open shell SCF and DFT:
- COSMO energies - the continuum solvation `COnductor-like Screening MOdel' of A. Klamt and G. Schürmann to describe dielectric screening effects in solvents.
- Generation of the electron density file for graphical display
- Electrostatic potential fit of atomic partial charges (CHELPG method with optional RESP restraints or charge constraints)
- In addition, interfaces are provided to
- The natural bond orbital (NBO) package
- Python: The Python programming language has been embedded within NWChem and many of the high level capabilities of NWChem can be easily combined and controlled by the user to perform complex operations.

Linear/Reduced Scaling Methods

- $O(N^2)$ and higher scaling of
  - storage or
  - computational requirements
- ... form a “wall” that parallel software must surmount

MADNESS

Multiresolution
Adaptive
Numerical
Scientific
Simulation

Multiresolution Adaptive Numerical Scientific Simulation

Ariana Beste[1], George I. Fann[1], Robert J. Harrison[1,2], Rebecca Hartman-Baker[1], Jun Jia[1], Shinichiro Sugiki[1]

[1]Oak Ridge National Laboratory

[2]University of Tennessee, Knoxville

In collaboration with

Gregory Beylkin[4], Fernando Perez[4], Lucas Monzon[4], Martin Mohlenkamp[5] and others

[4]University of Colorado

[5]Ohio University

harrisonrj@ornl.gov

The DOE funding

- This work is funded by the U.S. Department of Energy, the divisions of Advanced Scientific Computing Research and Basic Energy Science, Office of Science, under contract DE-AC05-00OR22725 with Oak Ridge National Laboratory. This research was performed in part using
  - resources of the National Energy Scientific Computing Center which is supported by the Office of Energy Research of the U.S. Department of Energy under contract DE-AC03-76SF0098,
  - and the Center for Computational Sciences at Oak Ridge National Laboratory under contract DE-AC05-00OR22725 .

SciDAC

Scientific Discovery through Advanced Computing

Multiresolution chemistry objectives

- Scaling to 1+M processors ASAP
- Complete elimination of the basis error
  - One-electron models (e.g., HF, DFT)
  - Pair models (e.g., MP2, CCSD, …)
- Correct scaling of cost with system size
- General approach
  - Readily accessible by students and researchers
  - Higher level of composition
  - Direct computation of chemical energy differences
- New computational approaches
  - *Fast algorithms with guaranteed precision*

Molecular orbitals of water

2-d contour plot

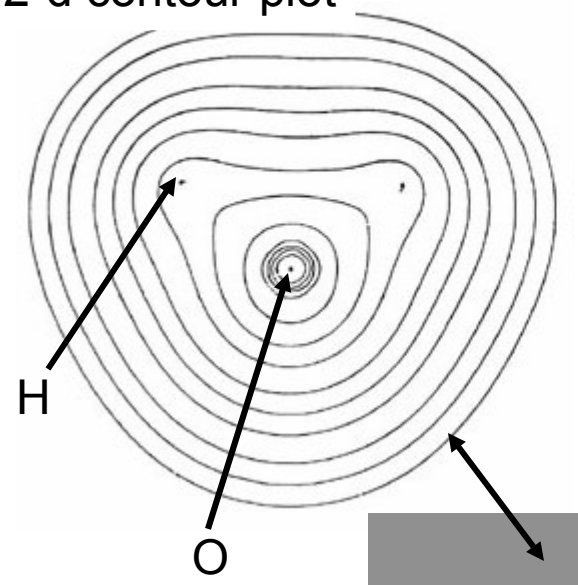

Iso-surfaces are 3-d contour plots – they show the surface upon which the function has a particular value

Water has 10 electrons (8 from oxygen, 1 from each hydrogen).

It is closed-shell, so it has 5 molecular orbitals each occupied with two electrons.

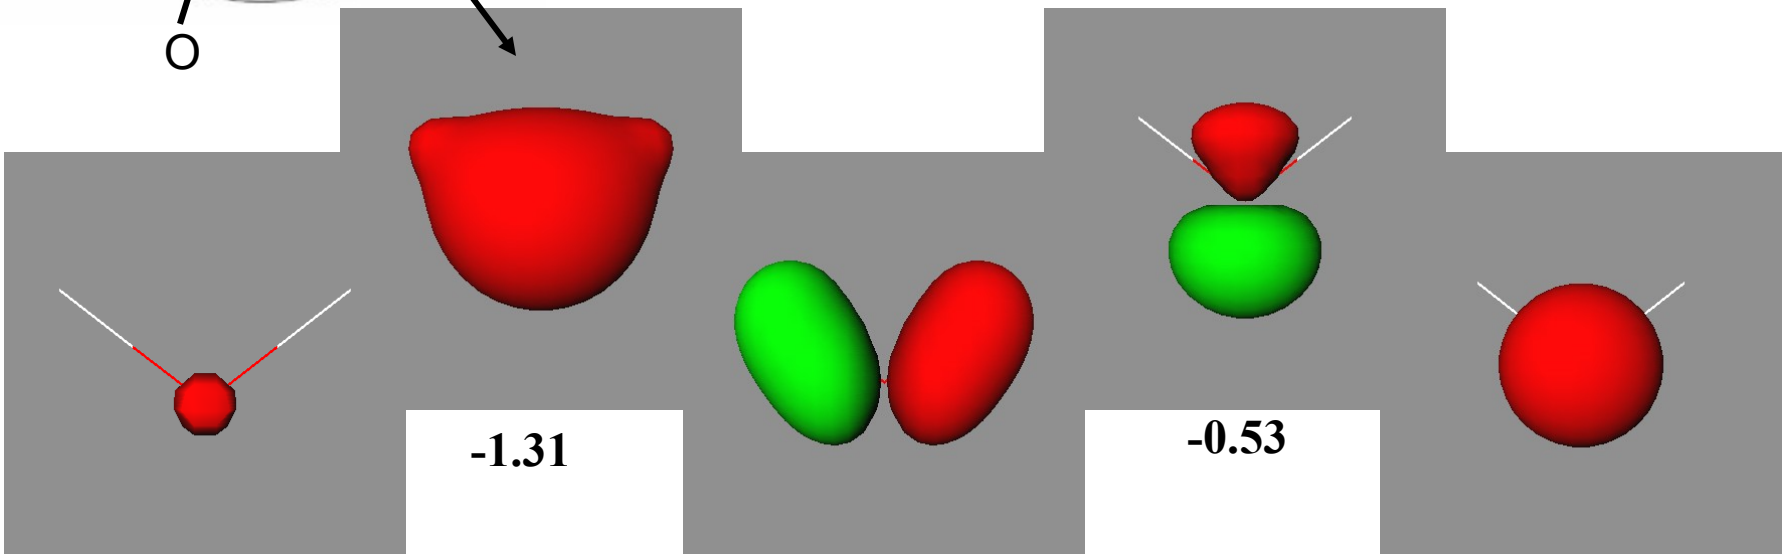

The energy of each orbital in atomic units

Linear Combination of Atomic Orbitals (LCAO)

- Molecules are composed of (weakly) perturbed atoms
  - Use finite set of atomic wave functions as the basis
  - Hydrogen-like wave functions are exponentials
- E.g., hydrogen molecul

$$1s(r) = \mathrm{e}^{-|r|}$$

$$\phi(r) = \mathrm{e}^{-|r-a|} + \mathrm{e}^{-|r-b|}$$

- Smooth function of molecular geometry
- MOs: cusp at nucleus with exponential decay

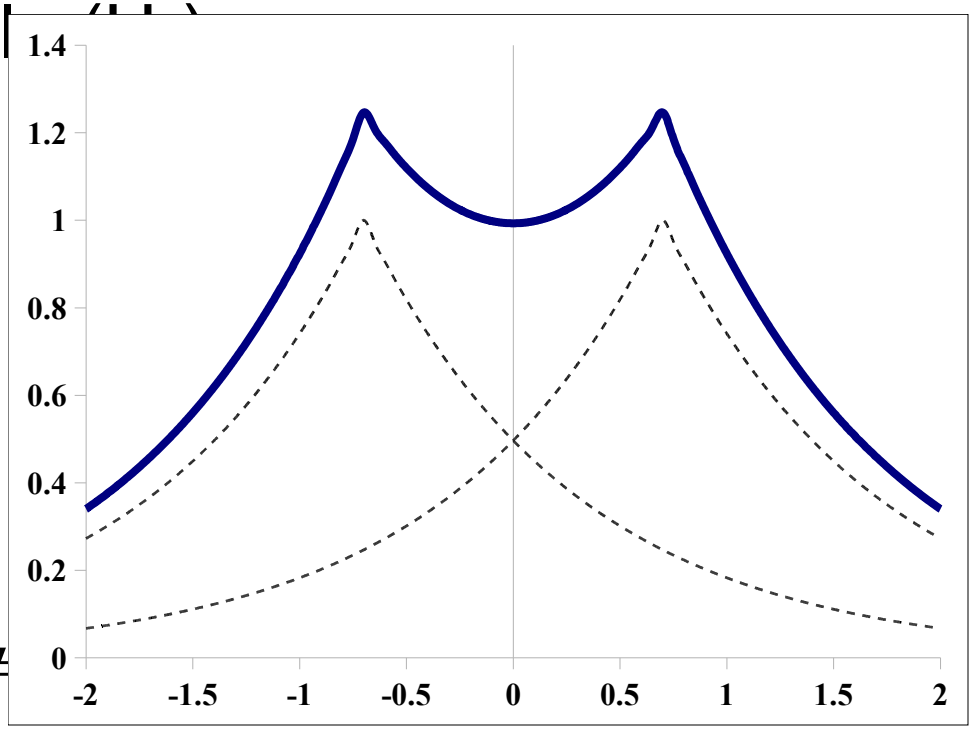

LCAO with Gaussian Functions

- Cannot compute integrals over exponential orbitals
- Boys (1950) noted that Gaussians are feasible
  - 6D integral reduced to 1D integrals which are tabulated once and stored (related to error function)
- Gaussian functions form a complete basis
  - With enough terms any radial function can be approximated to any precision using a linear combination of Gaussian functions

$$f(r) = \sum_{i=1}^{N} c_i e^{-a_i r^2} + O(\epsilon)$$

LCAO

- A fantastic success, but …
- Basis functions have extended support
  - causes great inefficiency in high accuracy calculations (functions on different centers overlap)
  - origin of non-physical density matrix
- Basis set superposition error (BSSE)
  - incomplete basis on each center leads to over-binding as atoms are brought together
- Linear dependence problems
  - accurate calculations require balanced approach to a complete basis on every atom
  - molecular basis can have severe linear dependence
- Must extrapolate to complete basis limit
  - unsatisfactory and not feasible for large systems

Essential techniques for fast computation

- Multiresolution

$$V_0 \subset V_1 \subset \cdots \subset V_n$$

$$V_n = V_0 + \left(V_1 - V_0\right) + \cdots + \left(V_n - V_{n-1}\right)$$

- Low-separation rank

$$f(x_1, \ldots, x_n) = \sum_{l=1}^{M} \sigma_l \prod_{i=1}^{d} f_i^{(l)}(x_i) + O(\epsilon)$$

$$\left\| f_i^{(l)} \right\|_2 = 1 \qquad \sigma_l > 0$$

- Low-operator rank

$$A = \sum_{\mu=1}^{r} u_\mu \sigma_\mu v_\mu^T + O(\epsilon)$$

$$\sigma_\mu > 0 \qquad v_\mu^T v_\lambda = u_\mu^T u_\lambda = \delta_{\mu\nu}$$

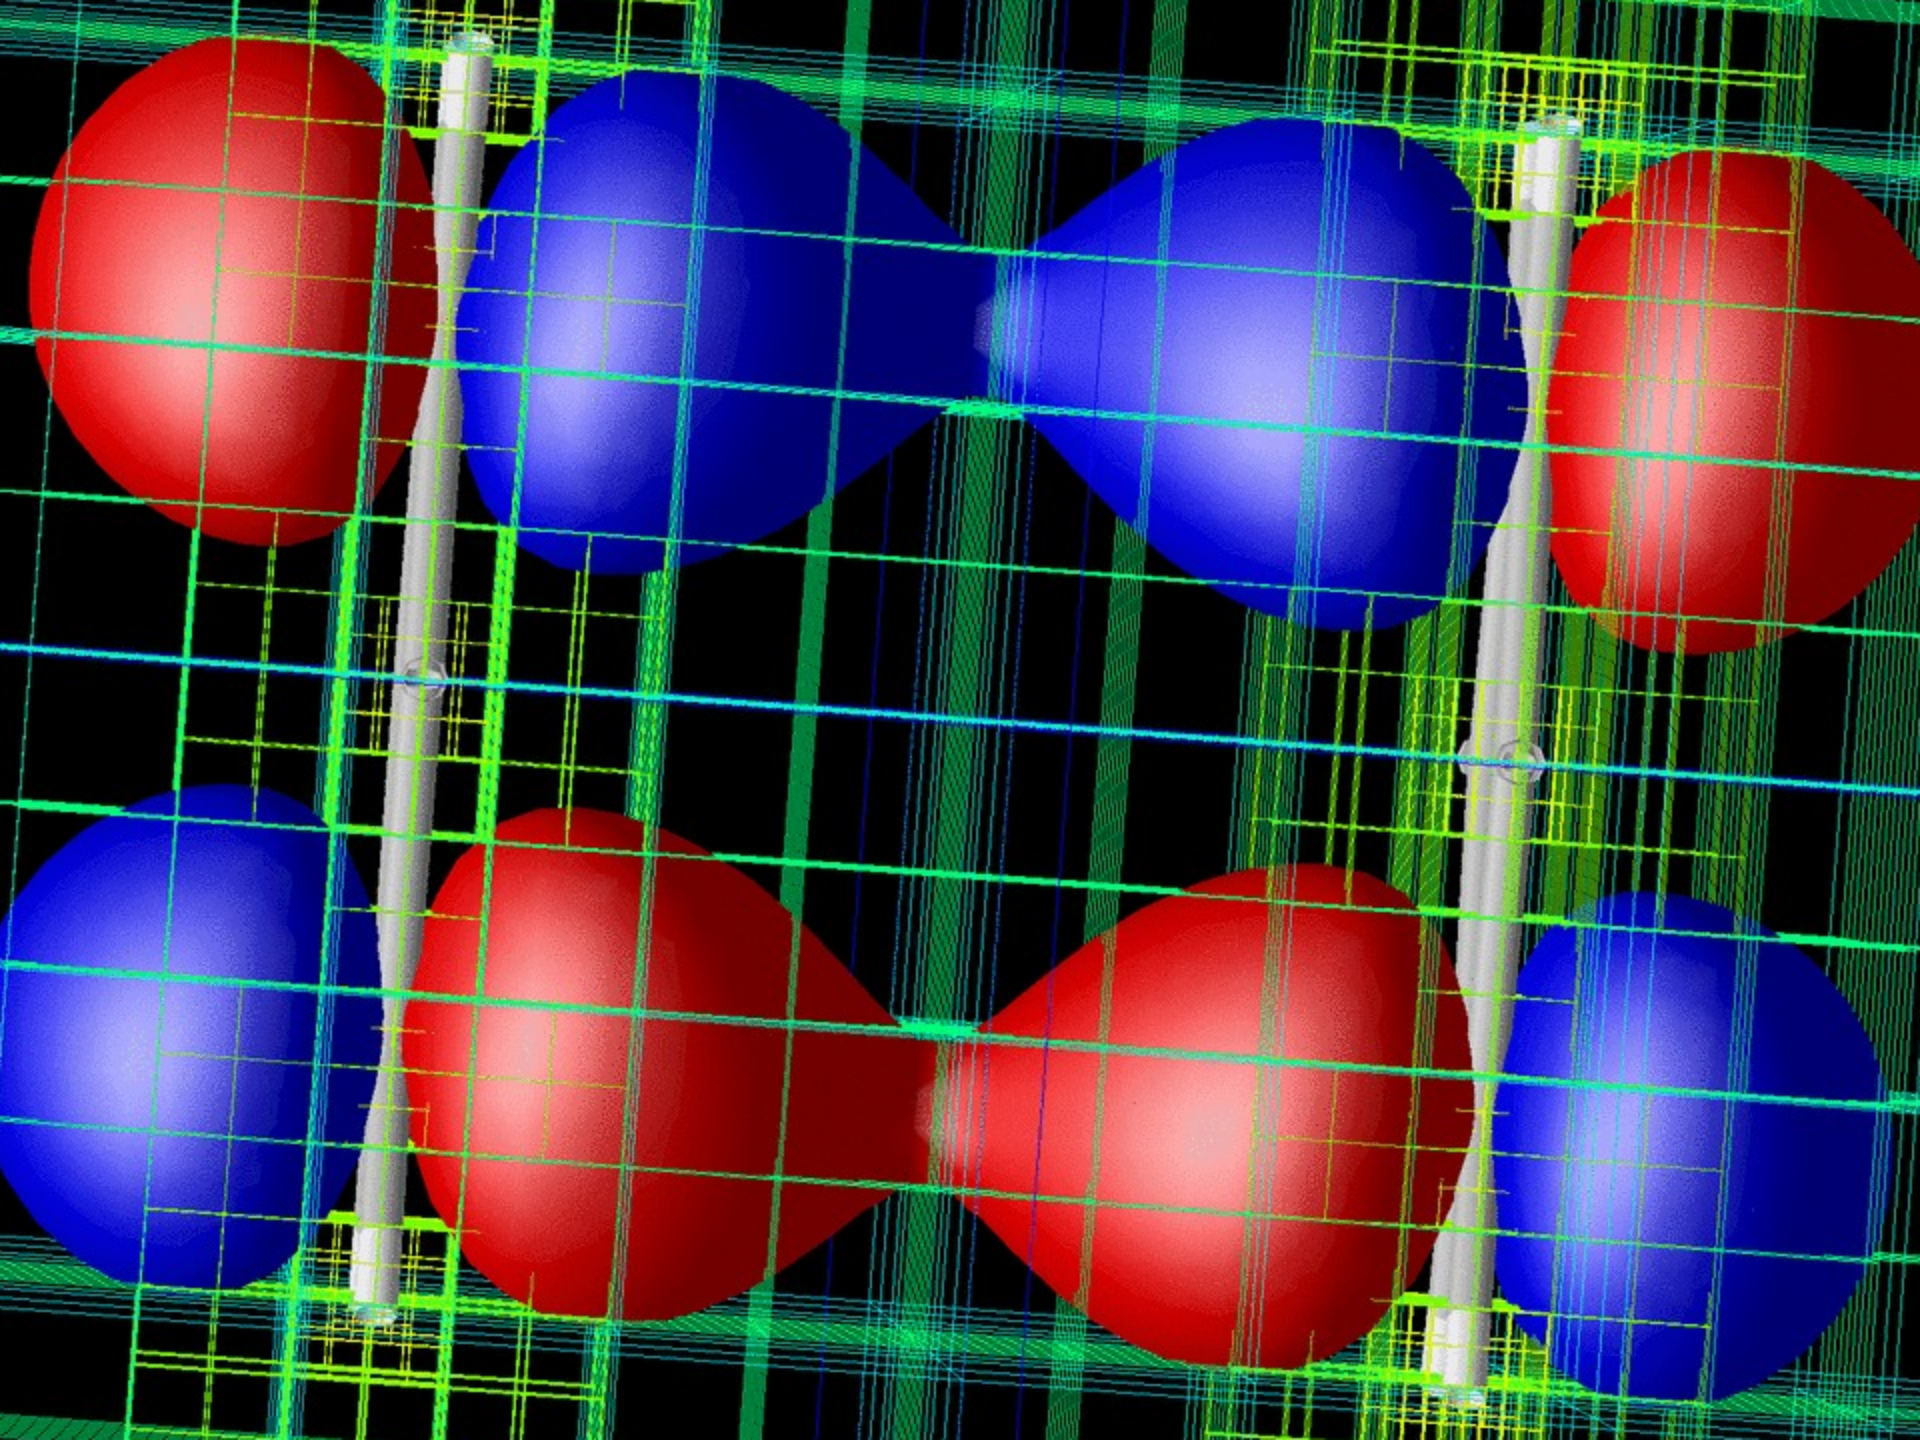

Please forget about wavelets

- They are not central
- Wavelets are a convenient basis for spanning V_n-V_{n-1} and understanding its properties
- But you don't actually need to use them
  - MADNESS does still compute wavelet coefficients, but *Beylkin's new code does not*
- Please remember this …
  - Discontinuous spectral element with multi-resolution and separated representations for fast computation with guaranteed precision in many dimensions.

Integral Formulation

- Solving the integral equation
  - Eliminates the derivative operator and related “issues”
  - Converges as fixed point iteration *with no preconditioner*

$$\left(-\tfrac{1}{2}\nabla^2 + V\right)\Psi = E\Psi$$

$$\Psi = -2\left(-\nabla^2 - 2E\right)^{-1} V\Psi$$

$$= -2G*\left(V\Psi\right)$$

$$\left(G*f\right)(r) = \int ds \frac{e^{-k|r-s|}}{4\pi|r-s|} f(s) \quad \text{in 3D ; } k^2 = -2E$$

Such Green’s Functions (bound state Helmholtz, Poisson) can be rapidly and accurately applied with a single, sparse matrix vector product.

Separated form for integral operators

$$T * f = \int ds\, K(r-s) f(s)$$

- Approach
  - Represent the kernel over a finite range as a sum of products of 1-D operators (often, not always, Gaussian)

$$r^{n,l-l'}_{ii',jj',kk'} = \sum_{\mu=0}^{M} X^{n,l_x-l'_x}_{ii'} Y^{n,l_y-l'_y}_{jj'} Z^{n,l_z-l'_z}_{kk'} + O(\epsilon)$$

  - Only need compute 1D transition matrices (X,Y,Z)
  - SVD the 1-D operators (low rank away from singularity)
  - Apply most efficient choice of low/full rank 1-D operator
  - Even better algorithms not yet implemented

Accurate Quadratures

$$\frac{e^{-\mu r}}{r} = \frac{2}{\sqrt{\pi}} \int_0^\infty e^{-x^2 t^2 - \mu^2/4t^2}\, dt$$

$$= \frac{2}{\sqrt{\pi}} \int_{-\infty}^\infty e^{-x^2 e^{2s} - \mu^2 e^{-2s}/4 + s}\, ds$$

- Trapezoidal quadrature
  - Geometric precision for periodic functions with sufficient smoothness
- Beylkin & Monzon
  - Further reductions, but not automatic

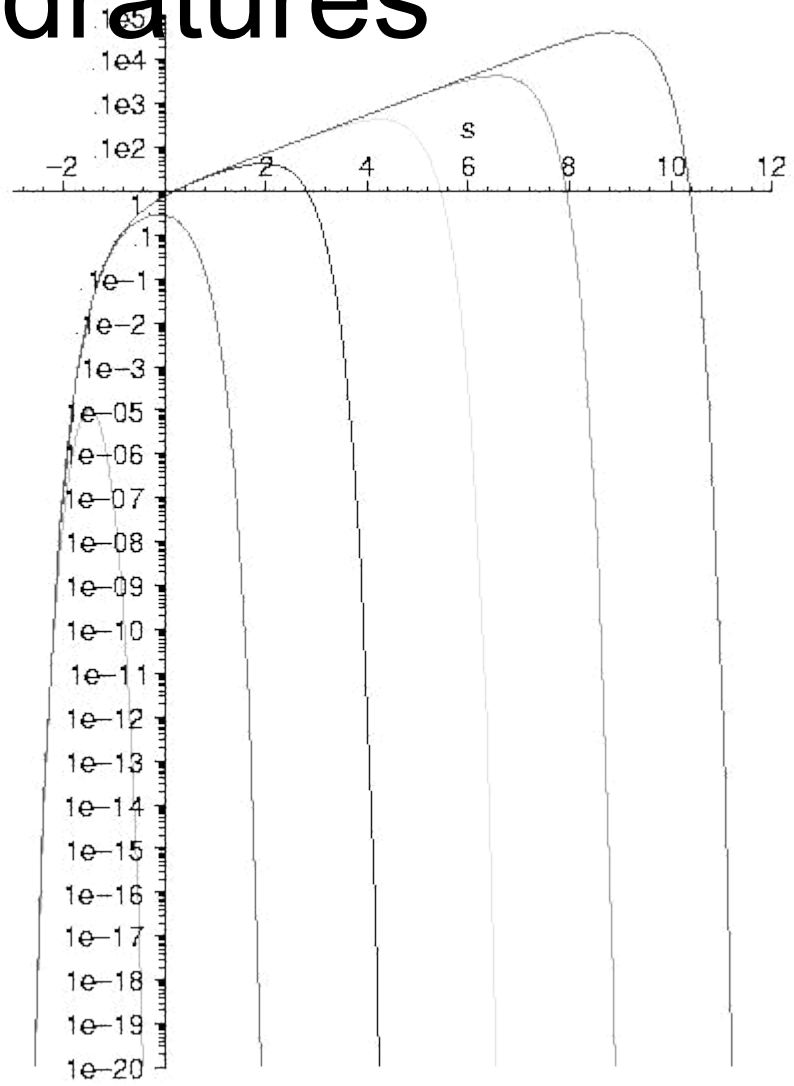

The kernel for x=1e-4,1e-3,1e-2,1e-,1e0.

The curve for x=1e-4 is the rightmost

Applications under active development

- DFT & HF for electrons
  - Energies, gradients, spectra, non-linear optical properties, Raman intensities (Harrison, Sekino, Yanai)
  - Molecules & periodic systems (Eguilez and Thornton)
- Atomic and molecular physics
  - Exact dynamics of few electron systems in strong fields (Krstic and Vence), MCSCF for larger systems
- Nuclear structure
  - G. Fann, et al.
- Preliminary studies in fusion and climate

TDDFT and CIS
T. Yanai with N.C. Handy

- Solve directly for the orbital response

$$\left(1-\hat{\rho}^{0}\right)\left[\left(\hat{F}^{0}-\epsilon_{p}^{0}\right)x_{p}\left(r\right)+\left\{\frac{\partial\hat{g}}{\partial\rho}\left[\rho^{0}\right]*\left(\sum_{i}^{occ}x_{i}\left(r\right)\phi_{i}^{\dagger}\left(r'\right)+\sum_{i}^{occ}\phi_{i}\left(r\right)y_{i}^{\dagger}\left(r'\right)\right)\right\}\phi_{p}\left(r\right)\right]$$
$$=\omega\, x_{p}\left(r\right),$$

$$\left(1-\hat{\rho}^{0}\right)^{\dagger}\left[\left(\hat{F}^{0}-\epsilon_{p}^{0}\right)^{\dagger}y_{p}\left(r\right)+\left\{\frac{\partial\hat{g}}{\partial\rho}\left[\rho^{0}\right]*\left(\sum_{i}^{occ}x_{i}\left(r\right)\phi_{i}^{\dagger}\left(r'\right)+\sum_{i}^{occ}\phi_{i}\left(r\right)y_{i}^{\dagger}\left(r'\right)\right)\right\}^{\dagger}\phi_{p}\left(r\right)\right]$$
$$=-\omega\, y_{p}\left(r\right),$$

  - Neglect *y* for CIS or Tamm-Dancoff

H_2 HOMO and CIS excited states

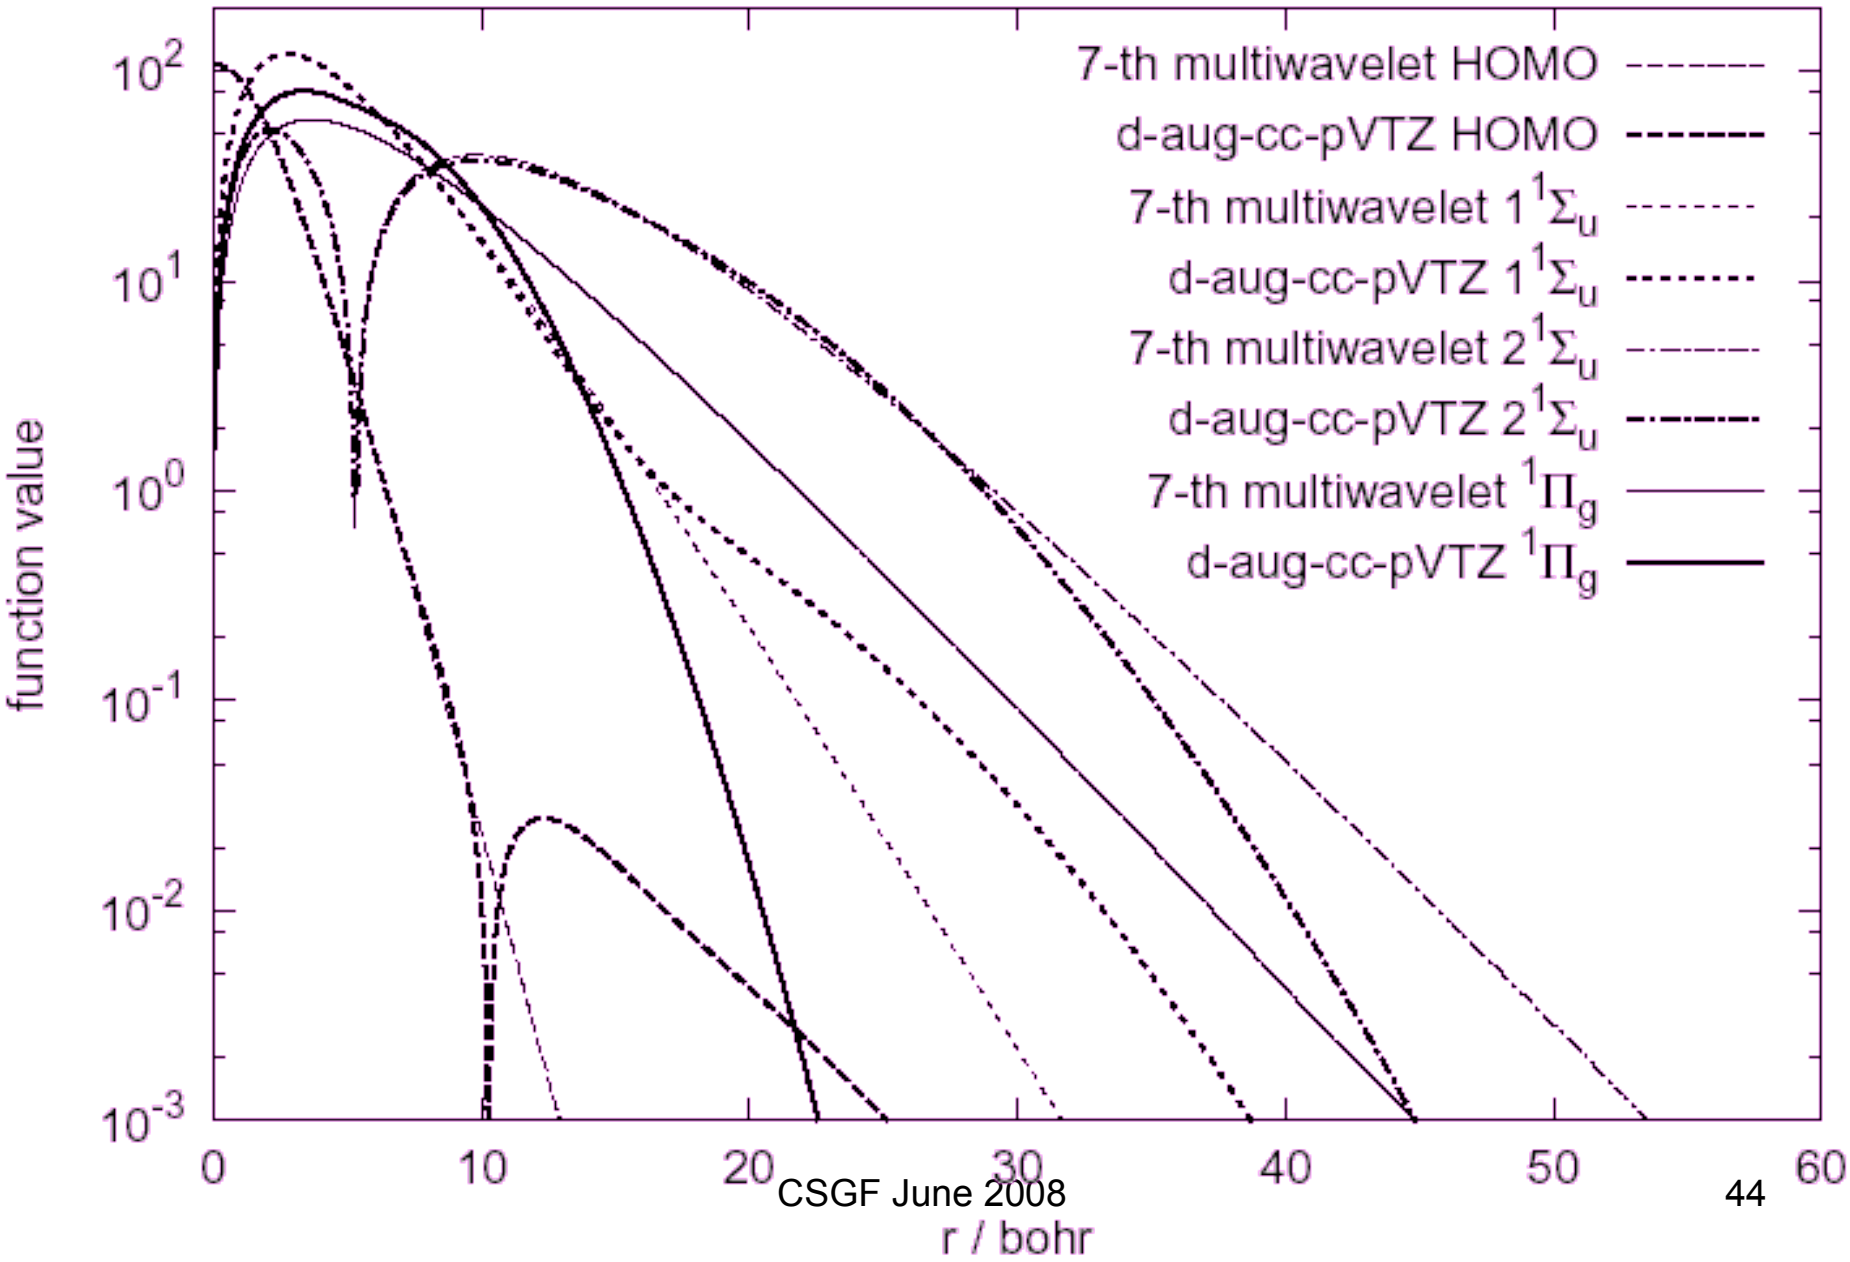

Time evolution

- Multiwavelet basis not optimal
  - Not strongly band limited
  - Explicit methods very unstable
    (DG introduces flux limiters, we use filters)
- Semi-group approach
  - Split into linear and non-linear parts

$$\dot{u}(x,t) = \hat{L}u + N(u,t)$$

$$u(x,t) = e^{\hat{L}t}u(x,0) + \int_0^t e^{\hat{L}(t-\tau)}N(u,\tau)\,d\tau$$

- Trotter-Suzuki methods
  - Time-ordered exponentials
  - Chin-Chen gradient correction (JCP 114, 7338, 2001)

$$e^{A+B} = e^{A/2}e^{B}e^{A/2} + O(\|[[A,B],A]\ldots\|)$$

Exponential propagator

- Imaginary time Schrodinger equation
  - Propagator is just the heat kernel

$$\left(-\frac{1}{2}\nabla^2 + V(x)\right)\psi(x,t) = \dot{\psi}(x,t)$$

$$\psi(x,t) \simeq e^{\nabla^2 t/4} e^{-Vt} e^{\nabla^2 t/4} \psi(x,0)$$

$$e^{\nabla^2 t/2} f(x) = \frac{1}{\sqrt{2\pi t}} \int_{-\infty}^{\infty} e^{-\frac{(x-y)^2}{2t}} f(y)\, dy$$

$$\lim_{t \to \infty} \psi(x,t) = \psi_0(x)$$

  - Wrap in solver to accelerate convergence

Exponential propagator

- Free-particle propagator in real time

$$\psi(x,t) = e^{i\nabla^2 t/2}\psi(x,0) = \frac{1}{\sqrt{2\pi i t}}\int_{-\infty}^{\infty} e^{-\frac{(x-y)^2}{2it}}\psi(y,0)\,dy$$

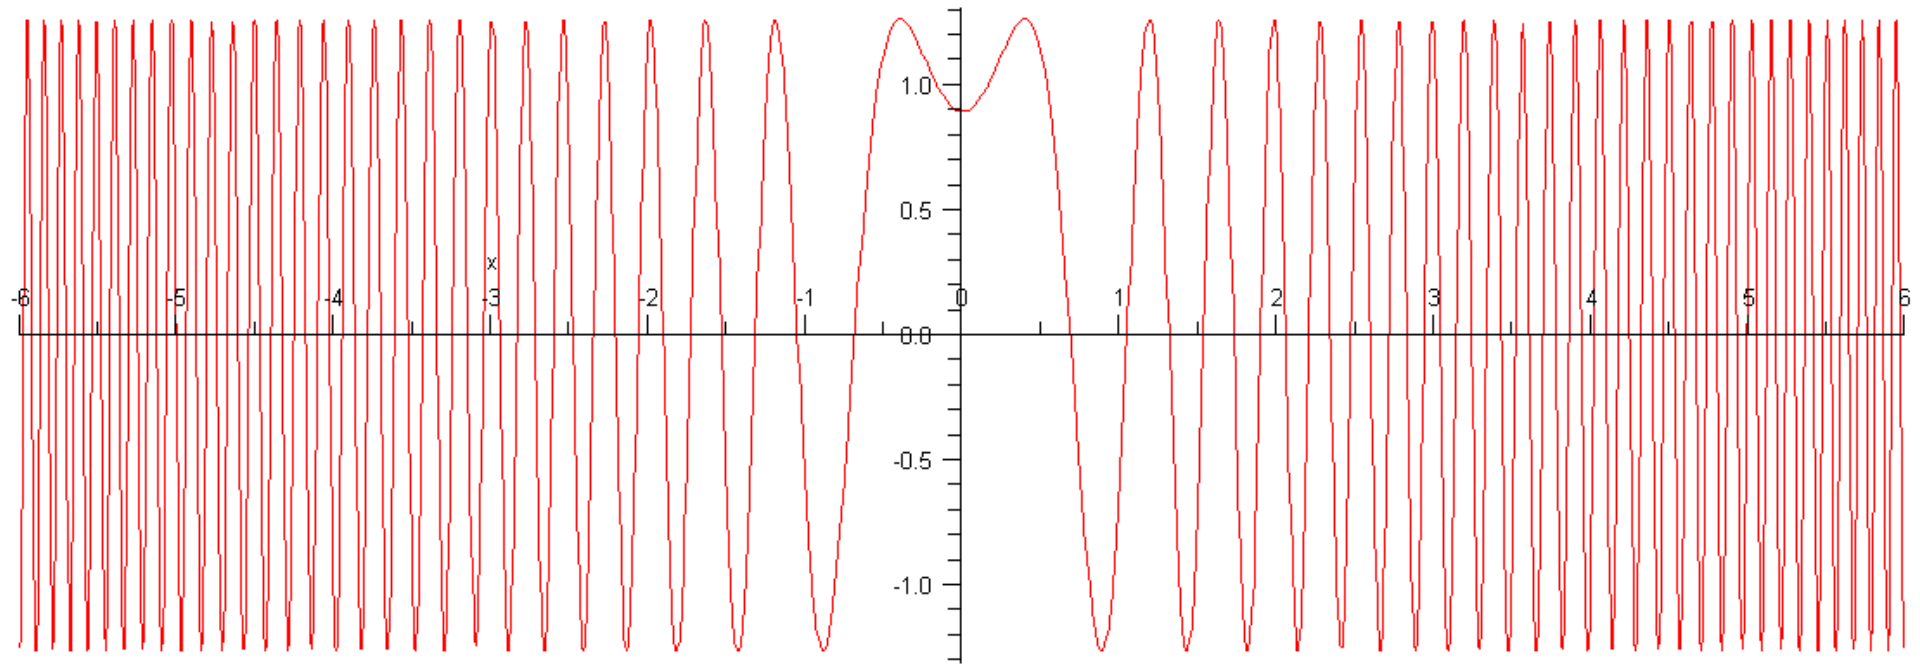

Exponential propagator

- Combine with projector onto band limit

$$\hat{G}_0(k,t,c) = e^{-i\frac{k^2 t}{2}}\left(1+(k/c)^{30}\right)^{-1}$$

$$h = \frac{\pi}{c} \qquad t_{crit} = \frac{2h^2}{pi}$$

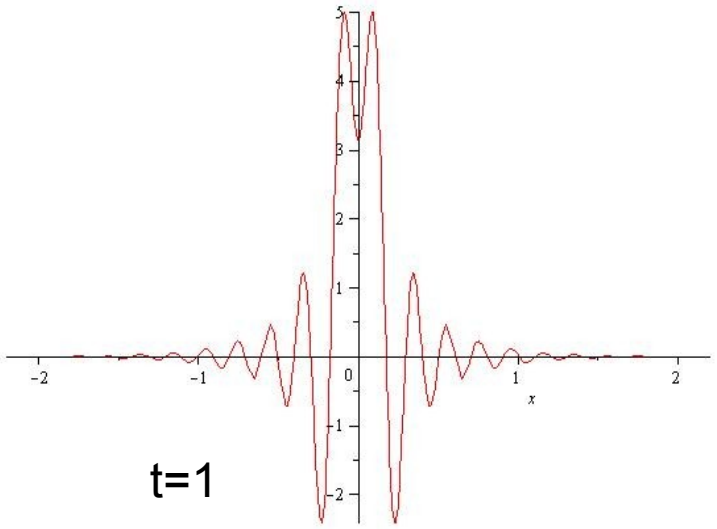

t=1

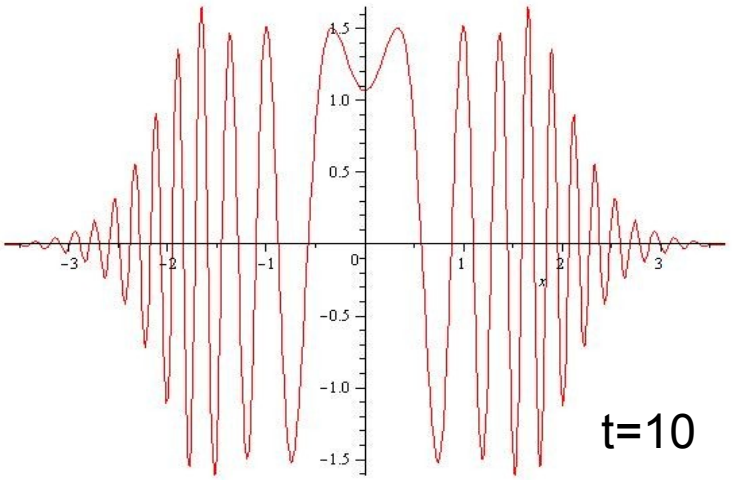

t=10

H-atom in laser field

- One electron – A still interesting test case
  - E.g., high-harmonic generation
  - With P. Krstic and N.E. Vence
- Preparing for T2O runs
  - Lie propagator much faster and stable

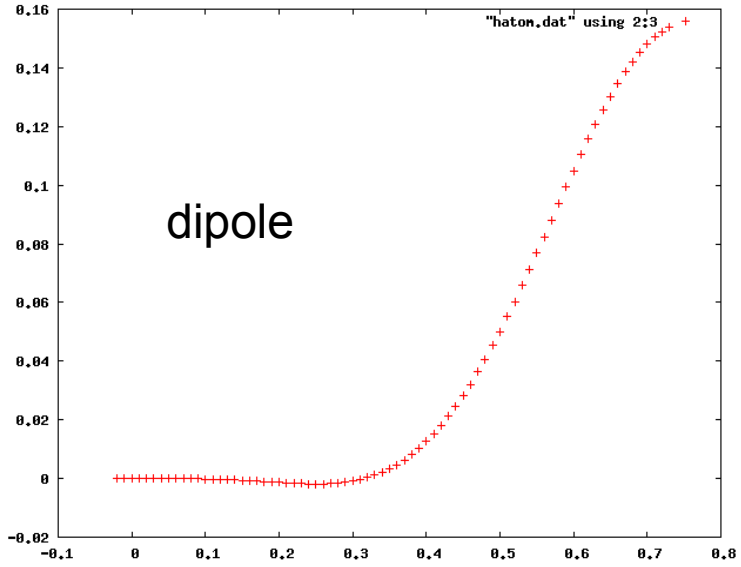

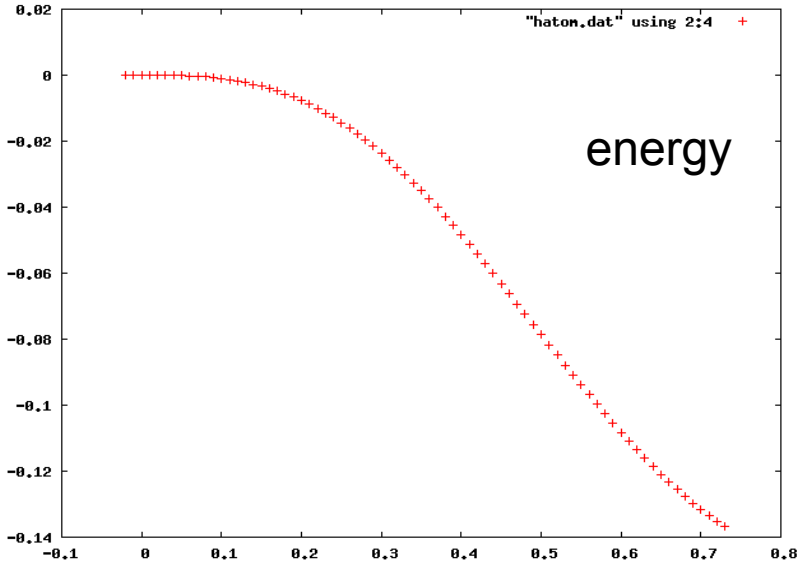

Path to linear scaling HF & DFT

- Need speed and precision
  - Absolute error cost $O\left(N \ln N / \epsilon\right)$
  - Relative error cost $O\left(N \ln 1 / \epsilon\right)$
- Coulomb potential
- HF exchange potential
- Orbital update
- Orthogonalization and or diagonalization
- Linear response properties

Electron correlation

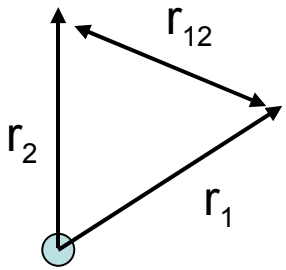

- All defects in the mean-field model are ascribed to electron correlation
- Consideration of singularities in the Hamiltonian imply that for a two-electron singlet atom (e.g., He)

$$\Psi(r_1, r_2, r_{12}) = 1 + \tfrac{1}{2} r_{12} + O(r_{12}^2) \quad \text{as} \quad r_{12} \to 0$$

- Include the inter-electron distance in the wavefunction
  - E.g., Hylleraas 1938 wavefunction for He

$$\Psi(r_1, r_2, r_{12}) = e^{-\zeta(r_1 + r_2)}(1 + a r_{12} + \cdots)$$

  - Potentially very accurate, but not systematically improvable, and (until recently) not computationally feasible for many-electron systems

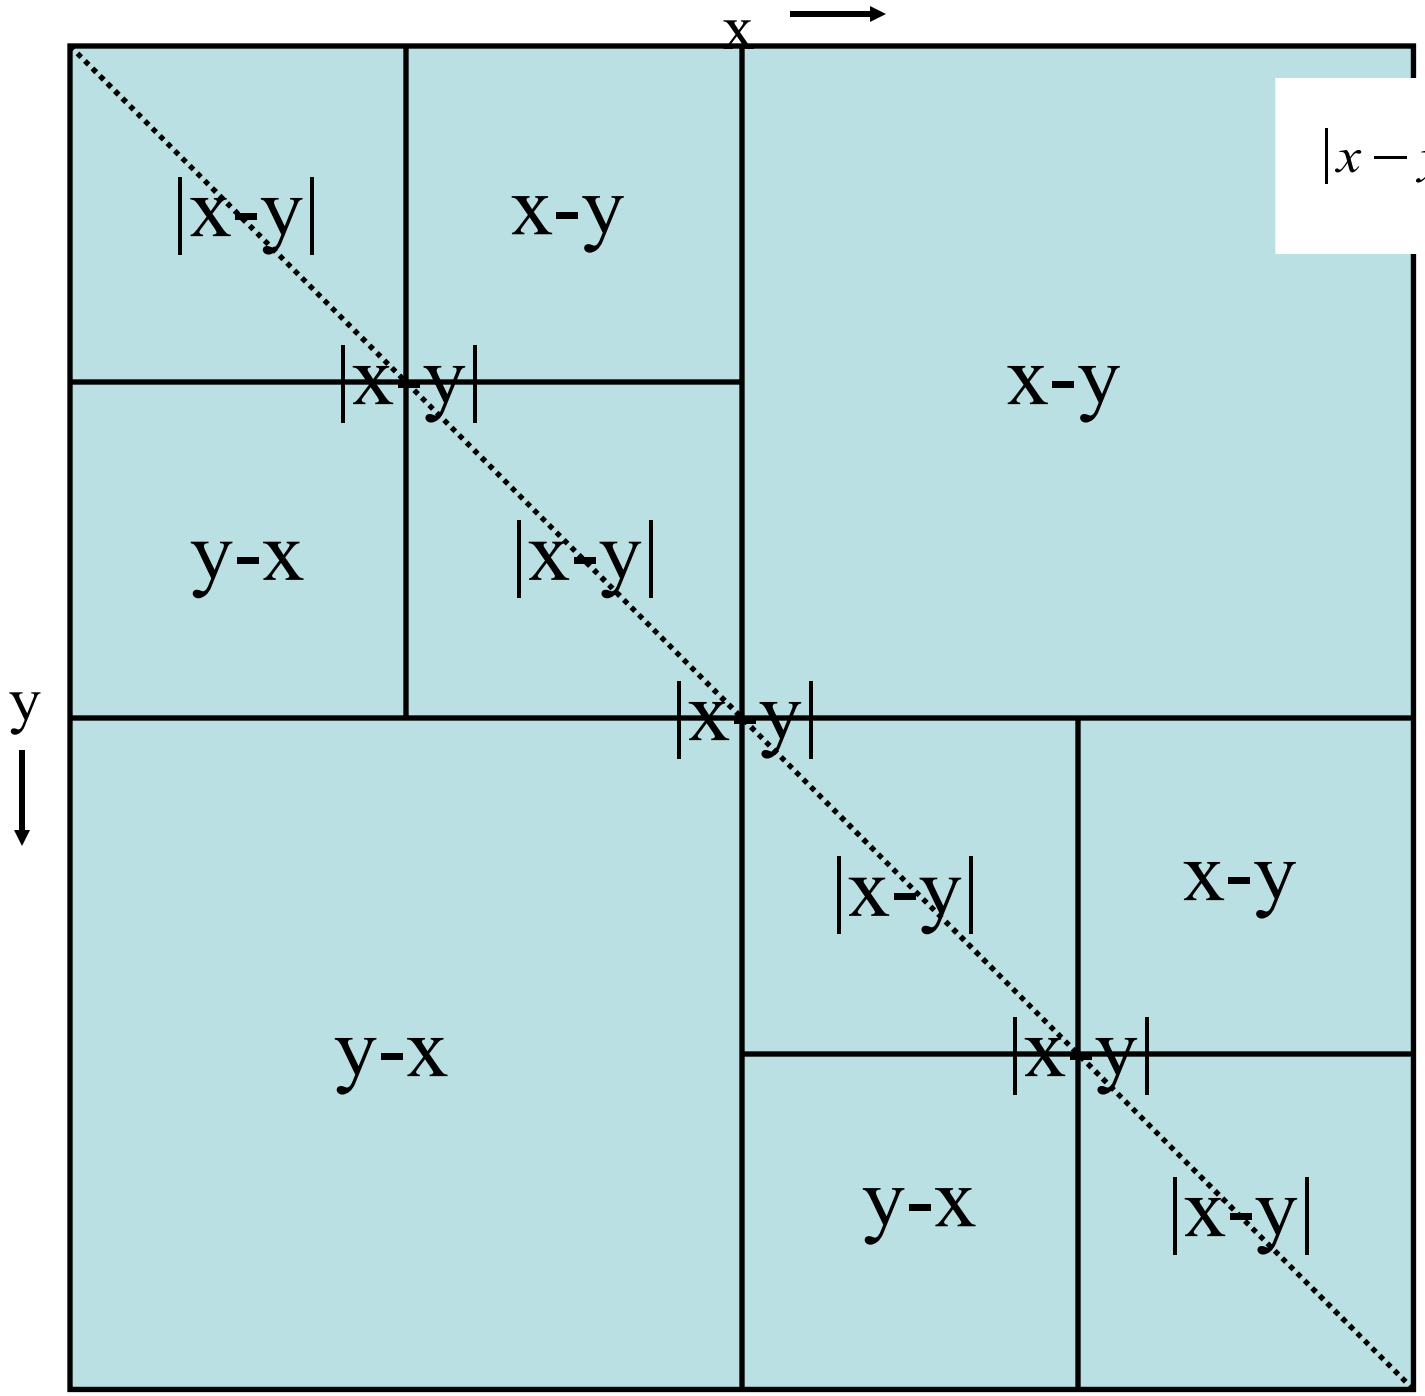

$$|x-y| = \sum_{\mu=1}^{r} f_\mu(x)\, g_\mu(y)$$

r = separation rank

In 3D, ideally must be one box removed from the diagonal

Diagonal box has full rank

Boxes touching diagonal (face, edge, or corner) have increasingly low rank

Away from diagonal $r = O(-\log \varepsilon)$

High-level composition

- Close to the physics

$$E = \langle \psi | -\frac{1}{2}\nabla^2 + V | \psi \rangle + \int \psi^2(x) \frac{1}{|x-y|} \psi^2(y)\, dx\, dy$$

```
operatorT op = CoulombOperator(k, rlo, thresh);
functionT rho = psi*psi;
double twoe = inner(apply(op,rho),rho);
double pe = 2.0*inner(Vnuc*psi,psi);
double ke = 0.0;
for (int axis=0; axis<3; axis++) {
    functionT dpsi = diff(psi,axis);
    ke += inner(dpsi,dpsi);
}
double energy = ke + pe + twoe;
```

High-level composition

- Express ***ALL*** available parallelism without burdening programmer
  - Internally, MADNESS is looking after data and placement and scheduling of operations on individual functions
  - Programmer must express parallelism over multiple functions and operators
    - But is *not* responsible for scheduling or placement

High-level composition

- E.g., make the matrix of KE operator
  - All scalar operations include optional fence
    - E.g., functionT scale(const functionT& f, T scale, bool fence=true)
  - Internally, operations on vectors schedule all tasks with only one fence

```
Tensor<double>
kinetic_energy_matrix(World& world,
                      const vector<functionT>& v) {
  int n = v.size();
  Tensor<double> r(n,n);
  for (int axis=0; axis<3; axis++) {
     vector<functionT> dv = diff(world,v,axis);
     r += inner(world, dv, dv);
  }
  return r.scale(0.5);
}
```

MADNESS architecture

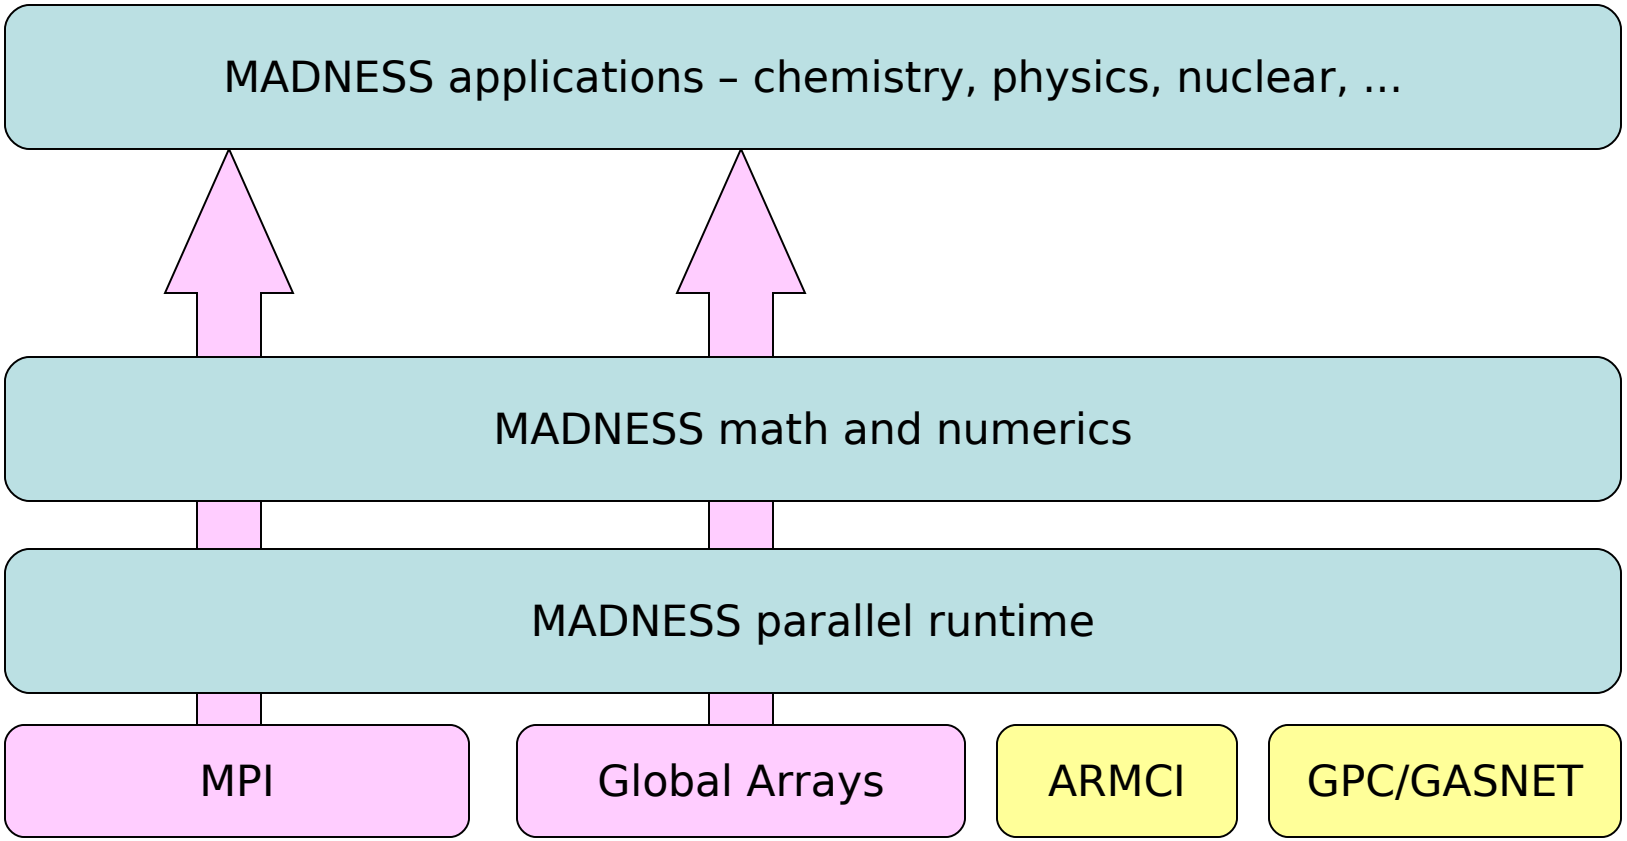

Runtime Objectives

- Scalability to 1+M processors ASAP
- Runtime responsible for
  - scheduling and placement,
  - managing data dependencies,
  - hiding latency, and
  - Medium to coarse grain concurrency
- Compatible with existing models
  - MPI, Global Arrays
- Borrow successful concepts from Cilk, Charm++, Python
- Anticipating next gen. languages

Key elements

- Futures for hiding latency and automating dependency management
- Global names and name spaces
- Non-process centric computing
  - One-sided messaging between objects
  - Retain place=process for MPI/GA legacy
- Dynamic load balancing
  - Data redistribution, work stealing, randomization

Futures

- Result of an asynchronous computation
  - Cilk, Java, HPCLs
- Hide latency due to communication or computation
- Management of dependencies
  - Via callbacks

```
int f(int arg);
ProcessId me, p;

Future<int> r0=task(p, f, 0);
Future<int> r1=task(me, f, r0);

// Work until need result

cout << r0 << r1 << endl;
```

Process “me” spawns a new task in process “p” to execute `f(0)` with the result eventually returned as the value of future `r0`. This is used as the argument of a second task whose execution is deferred until its argument is assigned. Tasks and futures can register multiple local or remote callbacks to express complex and dynamic dependencies.

Global Namespaces

- Specialize global names to containers
  - Hash table done
  - Arrays, etc., planned
- Replace global pointer (process+local pointer) with more powerful concept
-
- User definable map from keys to “owner” process

```
class Index;  // Hashable
class Value {
   double f(int);
};

WorldContainer<Index,Value> c;
Index i,j;  Value v;
c.insert(i,v);
Future<double> r =
   c.task(j,&Value::f,666);
```

A container is created mapping indices to values.

A value is inserted into the container.

A task is spawned in the process owning key `j` to invoke `c[j].f(666)`.

Namespaces are a large part of the elegance of Python and success of Charm++ (chares+arrays)

Summary

- Huge computational resources are rushing towards us
  - Tremendous scientific potential
  - Tremendous challenges
    - Research
    - Education
    - Community
- UT and ORNL are at the very center
  - Think of us when you have/want something fun and challenging to do

Extra Slides

HF Exchange (T. Yanai)

- HF or exact exchange
  - Features in the most successful XC functionals

$$\hat{K}\, f(x) = \sum_{i}^{\text{occupied}} n_i \phi_i(x) \int dy \frac{\phi_i(y) f(y)}{|x-y|}$$

  - Invariant to unitary rotation of occupied states with same occupation number
  - Localize the orbitals – only O(1) products but potential is still global
  - Compute potential only where orbital non-zero
  - Cost to apply to all orbitals circa O(N)

Orbital update

- Directly solve for localized orbitals that span space of occupied eigenfunctions
  - Rigorous error control from MRA refinement
  - Never construct the eigenfunctions
  - Update only diagonal multipliers
    - Off diagonal from localization process

$$\phi_i(x) = -(\hat{T} - \zeta)^{-1}\left((V + \zeta)\phi_i - \sum_j^{\text{occupied}} \phi_j(x)\epsilon_{ji}\right)$$

Inner products

- The most expensive term for plane wave codes leading to cost O(N^2 M)
- Inexpensive in MRA basis

$$\langle f|g\rangle = s_f^{00}\,.\,s_g^{00} + \sum_{n=0}^{\text{۲}^n-\text{۱}} \sum_{l=\text{۰}} d_f^{nl}\,.\,d_g^{nl}$$

  - Orthogonal basis from local adaptive refinement implies zero/reduced work if
    - Functions do not overlap
    - Functions locally live at different length scales

Advancing Scientific Discovery

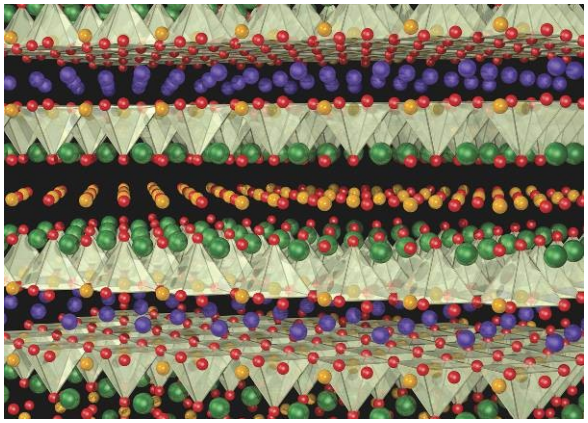

Resolved decades-long controversy about validity of 2D Hubbard model in predicting behavior of high-temperature superconducting cuprate planes

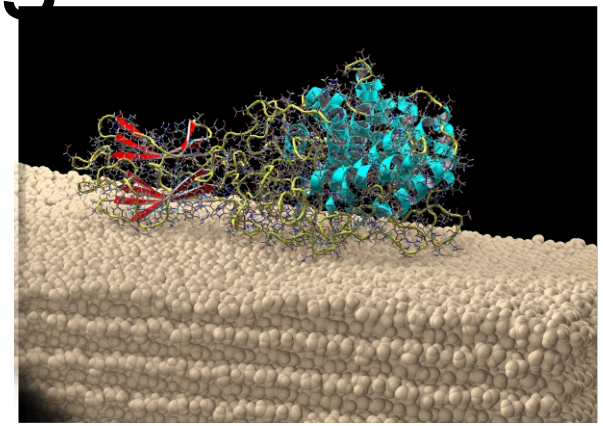

300K-atom models of cellulase enzyme on cellulose substrate reveal interior enzyme vibrations that influence reaction rates converting cellulose to ethanol

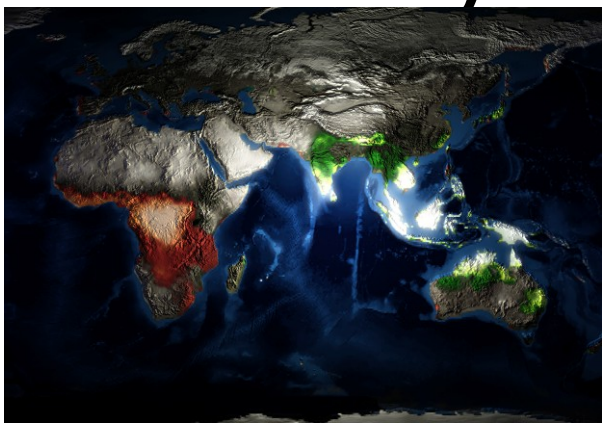

Addition and intercomparison of carbon-land models in new climate model is resolving key processes for carbon sources & sinks

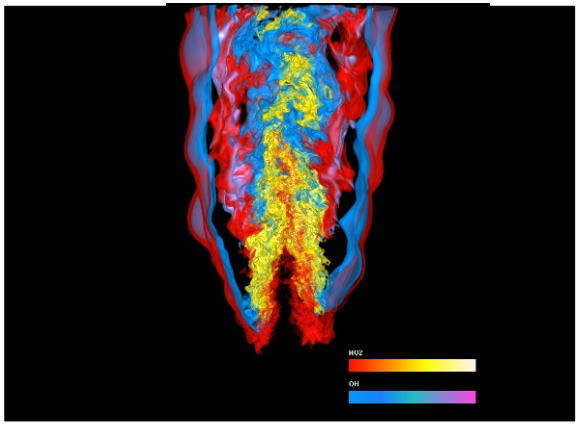

Turbulence chemistry revealed in study of lifted turbulent H_2/air jet flames in ignitive coflow relevant to diesel engines and gas turbines

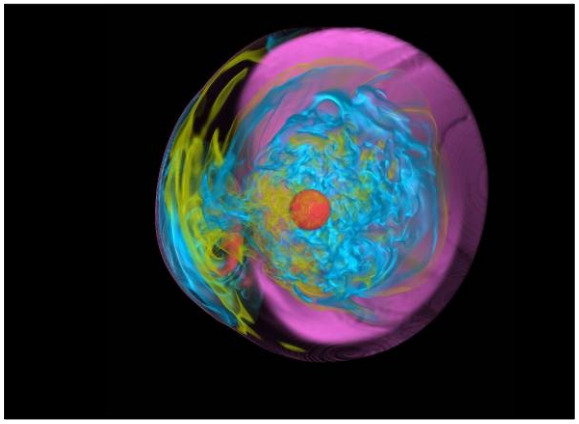

Instability of supernova shocks was discovered directly through simulation and core collapse pulsar mechanism was explained

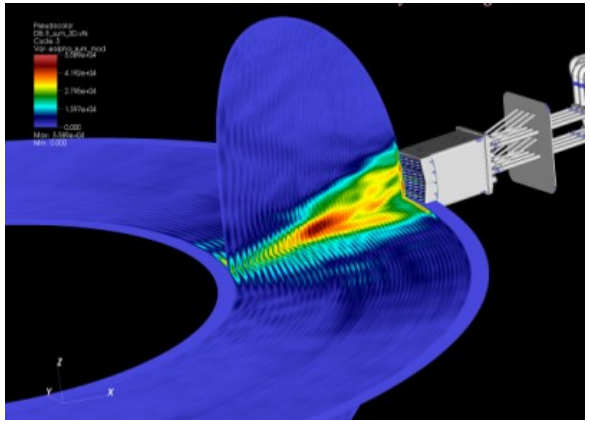

Providing increasing assurance that RF power will effectively heat ITER