

# NWChem and MADNESS

## A brief exploration for exascale

Robert J. Harrison

Institute for Advanced Computational Science  
Stony Brook University

and

Center for Scientific Computing  
Brookhaven National Laboratory

[robert.harrison@stonybrook.edu](mailto:robert.harrison@stonybrook.edu)



**BROOKHAVEN**  
NATIONAL LABORATORY



# Molecular Science Software Project




PNNL

Yuri Alexeev,  
Eric Bylaska,  
Bert deJong,  
Mahin Hackler,  
Karol Kowalski,  
Lisa Pollack,  
Tjerk Straatsma,  
Marat Valiev,  
Edo Apra

ISU and Ames  
Theresa Windus

SBU & BNL  
Robert Harrison

<http://www.nwchem-sw.org>




## MS<sup>3</sup>

MOLECULAR SCIENCE  
SOFTWARE SUITE



## ECCE

EXTENSIBLE COMPUTATIONAL  
CHEMISTRY ENVIRONMENT



## NWChem

HIGH-PERFORMANCE COMPUTATIONAL  
CHEMISTRY SOFTWARE



## GA TOOLS

PARALLEL COMPUTING LIBRARIES  
AND SOFTWARE TOOLS

(Jarek Nieplocha), Manoj Krishnan,  
Bruce Palmer, Daniel Chavarría,  
Sriram Krishnamoorthy



Gary Black,  
Brett Didier,  
Todd Elsenthagen,  
Sue Havre,  
Carina Lansing,  
Bruce Palmer,  
Karen Schuchardt,  
Lisong Sun  
Erich Vorpapel

# Fock matrix in a Nutshell

$$F_{ij} = \sum_{kl} \left( 2(ij|kl) - (ik|jl) \right)$$

$$(\mu\nu|\sigma\lambda) = \int_{-\infty}^{\infty} g_{\mu}(r_1)g_{\nu}(r_1) \frac{1}{r_{12}} g_{\sigma}(r_2)g_{\lambda}(r_2) dr_1 dr_2$$



**1 integral contributes to 6 Fock Matrix elements**

$$(\mu\nu|\sigma\lambda) \otimes \begin{Bmatrix} D_{\mu\nu} \\ D_{\mu\sigma} \\ D_{\mu\lambda} \\ D_{\nu\sigma} \\ D_{\nu\lambda} \\ D_{\sigma\lambda} \end{Bmatrix} \Rightarrow \begin{Bmatrix} F_{\mu\nu} \\ F_{\mu\sigma} \\ F_{\mu\lambda} \\ F_{\nu\sigma} \\ F_{\nu\lambda} \\ F_{\sigma\lambda} \end{Bmatrix}$$

- Sparsity, variable integral costs, algorithm constraints, symmetry, shell blocking, ...

# Global Arrays (technologies)

Physically distributed data



Single, shared data structure

- Shared-memory-like model
  - Fast local access
  - NUMA aware and easy to use
  - MIMD and data-parallel modes
  - Inter-operates with MPI, ...
- BLAS and linear algebra interface
- Ported to major parallel machines
  - IBM, Cray, SGI, clusters,...
- Originated in an HPCC project
- Used by most major chemistry codes, financial futures forecasting, astrophysics, computer graphics
- Supported by DOE
- One of the legacies of Jarek Nieplocha, PNNL

# Global Arrays: A Portable “Shared-Memory” Programming Model for Distributed Memory Computers

Jaroslav Nieplocha, Robert J. Harrison and Richard J. Littlefield

Pacific Northwest Laboratory<sup>†</sup>, P.O. Box 999, Richland WA 99352

## Abstract

*Portability, efficiency, and ease of coding are all important considerations in choosing the programming model for a scalable parallel application. The message-passing programming model is widely used because of its portability, yet some applications are too complex to code in it while also trying to maintain a balanced computation load and avoid redundant computations. The shared-memory programming model simplifies coding, but it is not portable and often provides little control over interprocessor data transfer costs. This paper describes a new approach, called Global Arrays (GA), that combines the better features of both other models, leading to both simple coding and efficient execution. The key concept of GA is that it provides a portable interface through which each process in a MIMD parallel program can asynchronously access logical blocks of physically distributed matrices, with no need for explicit cooperation by other processes. We have implemented GA libraries on a variety of computer systems, including the Intel DELTA and Paragon, the IBM SP-1 (all message-passers), the Kendall Square KSR-2 (a nonuniform access shared-memory machine), and networks of Unix worksta-*

chemistry. At the same time, we and our colleagues at the Pacific Northwest Laboratory (PNL) have a short-term goal of developing, within the next three years, a suite of parallel chemistry application codes to be used in production mode for chemistry research at PNL's Environmental and Molecular Science Laboratory (EMSL) and elsewhere. The programming model and implementations described here have turned out to be useful for both purposes.

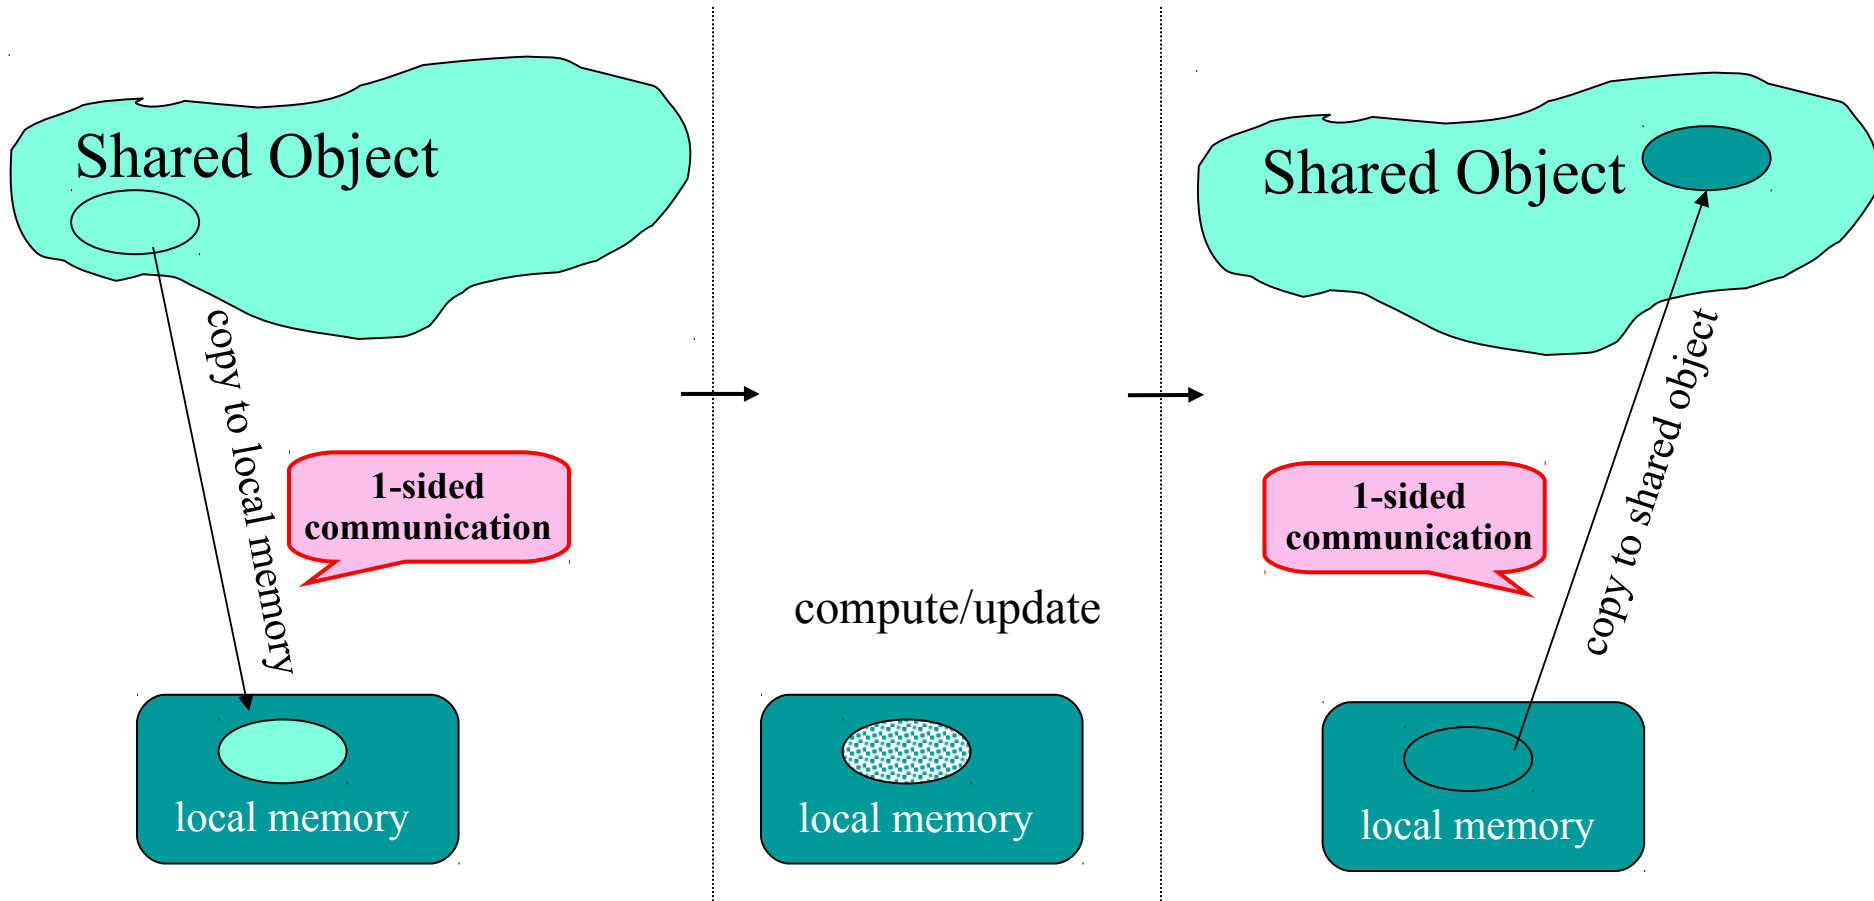
Two assumptions permeate our work. The first is that most high performance parallel computers currently and will continue to have physically distributed memories with Non-Uniform Memory Access (NUMA) timing characteristics, and will thus work best with application programs that have a high degree of locality in their memory reference patterns. The second assumption is that extra programming effort is and will continue to be required to construct such applications. Thus, a recurring theme in our work is to develop techniques and tools that allow applications with explicit control of locality to be developed with only a tolerable amount of extra effort.

There are significant tradeoffs between the important

# History and Design

- Prototyping at very start of NWChem project
  - Model full application not just kernel
  - 80-20 rule – more like 90-10 rule
- GA designed to solve a problem
  - Distributing large data structures while supporting irregular computation
  - Entire HF code
  - First 2 attempts (Linda-like) worked for kernel but not the rest of the code

# Non-uniform memory access model of computation



# Dynamic load balancing

```
my_next_task = SharedCounter(chunksize)
```

```
do i=1,max_i
```

```
  if(i.eq.my_next_task) then
```

```
    call ga_get(      )
```

(do work)



```
    call ga_acc(      )
```

```
    my_next_task = SharedCounter(chunksize)
```

```
  endif
```

```
enddo
```

```
Barrier()
```





# Distributed data SCF

- First success for NWChem and Global Arrays

```
do tiles of i
  do tiles of j
    do tiles of k
      do tiles of l
```

} Parallel loop nest

get patches ij, ik, il, jk, jl, kl

compute integrals

accumulate results back into patches

Mini-apps used to  
evaluate HPCS

languages Chapel,  
X10, Fortress

- just the data flow

$B = \text{block size}$

$$t_{\text{comm}} = O(B^2) \quad t_{\text{compute}} = O(B^4) \quad \frac{t_{\text{compute}}}{t_{\text{comm}}} = O(B^2)$$

# Higher-performance code

- Looks nothing like that!
- Sort shell pairs to evaluate in similar batches
  - Precomputation, vectorization – 10-fold speedup
  - Big increase in complexity and memory use
- Integral evaluation code – 100K lines!
- Careful screening with rigorous inequalities
  - Robustness, minimize overhead
-

# Highest-performance code

- Looks nothing like that!
- Strives for near linear scaling
- Coulomb interaction
  - Mix of FMM, FFT, and other fast methods
  - (near) linear scaling with system size
- Exchange interaction
  - Heavy screening and physical thresholding
- And this is just 1% of NWChem functionality

# The Tensor Contraction Engine: A Tool for Quantum Chemistry

## Oak Ridge National Laboratory

*David E. Bernholdt,*  
Venkatesh Choppella, *Robert*  
*Harrison*

## Pacific Northwest National Laboratory

*So Hirata*

## Louisiana State University

*J Ramanujam,*

## Ohio State University

*Gerald Baumgartner,* Alina  
Bibireata, Daniel Cociorva,  
Xiaoyang Gao, Sriram  
Krishnamoorthy, Sandhya  
Krishnan, Chi-Chung Lam,  
Quingda Lu, *Russell M.*  
*Pitzer, P Sadayappan,*  
Alexander Sibiriyakov

## University of Waterloo

*Marcel Nooijen,* Alexander  
Auer


<http://www.cis.ohio-state.edu/~gb/TCE/>

# Tensor Contraction Engine (TCE) (Kowalski, PNNL)



Highly parallel codes are needed in order to apply the CC theories to larger molecular systems

Symbolic algebra systems for coding complicated tensor expressions: Tensor Contraction Engine (TCE)



OCE

$$+\frac{1}{4}v_{ef}^{mn,ef}t_{ij}^{ab} - \frac{1}{2}v_{ef}^{mn,ef}t_{mj}^{ab} +$$

TCE

```
next = NXTASK(nprocs, 1)
DO p3b = noab+1, noab+nvab
DO p4b = p3b, noab+nvab
DO h1b = 1, noab
DO h2b = h1b, noab
IF (next.eq.count) THEN
CALL GET_HASH_BLOCK(d_a, dbl_mb(k_a), dim
- 1 + (noab+nvab) * (h1b_1 - 1 + (noab+
+nvab) * (p3b_1 - 1)))
CALL GET_HASH_BLOCK_I(d_a, dbl_mb(k_a), d
```

	Expression <sup>a</sup>
$D_{ij}^a t_i^a$	$f_i^a + t_{ij}^a t_j^a - t_n^a t_n^i + t_{ni}^a t_n^i + t_{ni}^a t_n^i - \frac{1}{2} t_{no}^a v_{fi} + \frac{1}{2} t_{fi}^a v_{no} + \frac{1}{4} t_{ino}^a v_{fg}$
$D_{ij}^{ab} t_{ij}^{ab}$	$v_{ij}^{ab} + P(a/b) I_{ij}^a t_{ij}^{fb} - P(i/j) I_{ij}^n t_{ij}^{ab} + \frac{1}{2} t_{ij}^a t_{ij}^{fb} + \frac{1}{2} t_{no}^a v_{ij}^{no}$ $+ P(a/b) P(i/j) t_{in}^a t_{ij}^{fb} - \frac{1}{2} P(a/b) I_{ij}^n t_{ij}^{ab} + \frac{1}{2} t_{no}^a v_{ij}^{no}$ $- \frac{1}{2} P(i/j) I_{ij}^n t_{ij}^{ab} + t_{nij}^a t_{ij}^{fb} + P(i/j) t_{ij}^a t_{ij}^{fb} - P(a/b) t_n^a t_{ij}^{nb} + \frac{1}{4} t_{ijno}^a v_{fg}$
$D_{ijk}^{abc} t_{ijk}^{abc}$	$P(a/bc) I_{ijk}^a t_{ijk}^{abc} - P(i/jk) I_{ijk}^n t_{ijk}^{abc} + \frac{1}{2} P(a/bc) t_{ijk}^a t_{ijk}^{abc} + \frac{1}{2} P(i/jk) t_{ino}^a t_{ijk}^{abc}$ $+ P(a/bc) P(i/jk) t_{ijn}^a t_{ijk}^{abc} + P(a/bc) P(i/jk) t_{ij}^a t_{ijk}^{abc} - P(a/bc) P(i/jk) t_{in}^a t_{ijk}^{abc}$ $+ t_{nij}^a t_{ijk}^{abc} + \frac{1}{2} P(a/bc) I_{ijk}^n t_{ijk}^{abc} - P(i/jk) I_{ijk}^n t_{nojk}^a + \frac{1}{4} t_{ijkno}^a v_{fg}$
$D_{ijkl}^{abcd} t_{ijkl}^{abcd}$	$P(abcd) I_{ijkl}^a t_{ijkl}^{abcd} - P(i/jkl) I_{ijkl}^n t_{ijkl}^{abcd} + \frac{1}{2} P(abcd) t_{ijkl}^a t_{ijkl}^{abcd} + \frac{1}{2} P(i/jkl) t_{ino}^a t_{ijkl}^{abcd}$ $+ P(abcd) P(i/jkl) t_{ijn}^a t_{ijkl}^{abcd} + P(abcd) P(i/jkl) t_{ij}^a t_{ijkl}^{abcd} - P(abcd) P(i/jkl) t_{in}^a t_{ijkl}^{abcd}$ $+ P(abcd) P(i/jkl) t_{ij}^a t_{ijkl}^{abcd} - P(abcd) P(i/jkl) t_{in}^a t_{ijkl}^{abcd} + P(abcd) P(i/jkl) t_{ij}^a t_{ijkl}^{abcd}$ $+ \frac{1}{2} P(abcd) P(i/jkl) t_{ino}^a t_{ijkl}^{abcd} + t_{nijkl}^a t_{ijkl}^{abcd} + \frac{1}{2} P(abcd) I_{ijkl}^n t_{ijkl}^{abcd}$ $- \frac{1}{2} P(i/jkl) I_{ijkl}^n t_{nojkl}^a$
$D_{ijklm}^{abcde} t_{ijklm}^{abcde}$	$P(abcde) I_{ijklm}^a t_{ijklm}^{abcde} - P(i/jklm) I_{ijklm}^n t_{ijklm}^{abcde} + \frac{1}{2} P(abcde) t_{ijklm}^a t_{ijklm}^{abcde} + \frac{1}{2} P(i/jklm) t_{ino}^a t_{ijklm}^{abcde}$ $+ P(abcde) P(i/jklm) t_{ijn}^a t_{ijklm}^{abcde} + P(abcde) P(i/jklm) t_{ij}^a t_{ijklm}^{abcde} - P(abcde) P(i/jklm) t_{in}^a t_{ijklm}^{abcde}$ $+ P(abcde) P(i/jklm) t_{ij}^a t_{ijklm}^{abcde} - P(abcde) P(i/jklm) t_{in}^a t_{ijklm}^{abcde} + P(abcde) P(i/jklm) t_{ij}^a t_{ijklm}^{abcde}$ $+ \frac{1}{2} P(abcde) P(i/jklm) t_{ino}^a t_{ijklm}^{abcde} + t_{ijnklm}^a t_{ijklm}^{abcde} + \frac{1}{2} P(abcde) I_{ijklm}^n t_{ijklm}^{abcde}$ $+ P(abcde) P(i/jklm) t_{ij}^a t_{ijklm}^{abcde} - P(abcde) P(i/jklm) t_{in}^a t_{ijklm}^{abcde}$ $+ P(abcde) P(i/jklm) t_{ij}^a t_{ijklm}^{abcde} - P(abcde) P(i/jklm) t_{in}^a t_{ijklm}^{abcde}$ $+ P(abcde) P(i/jklm) t_{ij}^a t_{ijklm}^{abcde} - P(abcde) P(i/jklm) t_{in}^a t_{ijklm}^{abcde}$




Pacific Northwest  
NATIONAL LABORATORY



Office of  
Science  
U.S. DEPARTMENT OF ENERGY

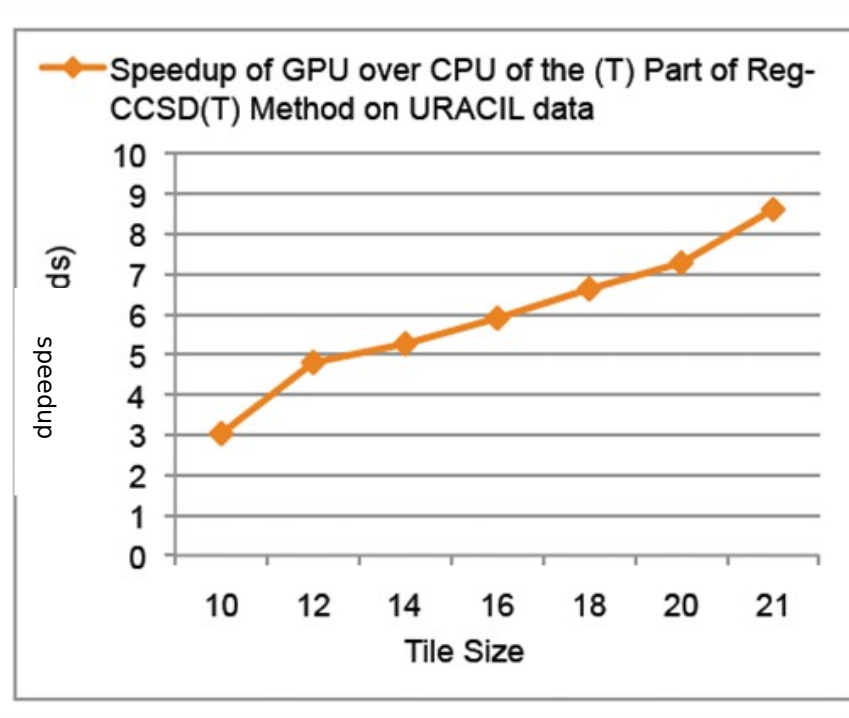

# Parallel performance (Karwolski et al., PNNL)



# Towards future computer architectures

## (Villa, Krishnamoorthy, Kowalski)

The CCSD(T)/Reg-CCSD(T) codes have been rewritten in order to take advantage of GPGPU accelerators  
Preliminary tests show very good scalability of the most expensive N7 part of the CCSD(T) approach




M

A

D

N



T

S



Multiresolution  
Adaptive  
Numerical  
Scientific  
Simulation

S



# Multiresolution Adaptive Numerical Scientific Simulation

*George I. Fann<sup>1</sup>, Diego Galindo<sup>1</sup>, Robert J. Harrison<sup>3</sup>,  
Scott Thornton<sup>2</sup>, Judy Hill<sup>1</sup>, and Jun Jia<sup>1</sup>*

*<sup>1</sup>Oak Ridge National Laboratory*

*<sup>2</sup>University of Tennessee, Knoxville*

*<sup>3</sup>Stony Brook University, Brookhaven National Laboratory*



*In collaboration with*

*Gregory Beylkin<sup>4</sup>, Lucas Monzon<sup>4</sup>, Hideo Sekino<sup>5</sup>  
and Edward Valeev<sup>6</sup>*

*<sup>4</sup>University of Colorado*

*<sup>5</sup>Toyohashi Technical University, Japan*

*<sup>6</sup>Virginia Tech*

**BROOKHAVEN**  
NATIONAL LABORATORY

*robert.harrison@gmail.com*



Stony Brook **University**

# Big picture


- Want robust algorithms that scale correctly with system size and are easy to write
- Robust, accurate, fast computation
  - Gaussian basis sets: high accuracy yields dense matrices and linear dependence –  $O(N^3)$
  - Plane waves: force pseudo-potentials –  $O(N^3)$
  - $O(N \log^m N \log^k \epsilon)$  is possible, guaranteed  $\epsilon$
- Semantic gap
  - Why are our equations just  $O(100)$  lines but programs  $O(1M)$  lines?
- Facile path from laptop to exaflop

# What is MADNESS?

- A general purpose numerical environment for reliable and fast scientific simulation
  - Chemistry, nuclear physics, atomic physics, material science, nanoscience, climate, fusion, ...
- A general purpose parallel programming environment designed for the peta/exa-scales
- Addresses many of the sources of complexity that constrain our HPC ambitions


<http://code.google.com/p/m-a-d-n-e-s-s>

<http://harrison2.chem.utk.edu/~rjh/madness/>



# IBM BGQ Team

- ANL
  - Alvaro Vasquez
  - Jeff Hammond
  - Nichols Romero
- OSU
  - Kevin Stocks
- SBU
  - Robert Harrison
- UTK
  - Scott Thornton



George Fann



Judy Hill



Gregory Beylkin




Rebecca  
Hartman-Baker


Jun Jia  
Tetsuya Kato  
Justus Calvin  
J. Pei



Ariana Beste



Eduard Valeyev



Alvaro Vasquez




Hideo Sekino




Robert Harrison



Nicholas Vence



Takahiro Ii



Scott Thornton



Matt Reuter




Paul Sutter

# What is MADNESS?

- A general purpose numerical environment for reliable and fast scientific simulation
  - Chemistry, nuclear physics, atomic physics, material science, nanoscience, climate, fusion, ...
- A general purpose parallel programming environment designed for the peta/exa-scales
- Addresses many of the sources of complexity that constrain our HPC ambitions

<http://code.google.com/p/m-a-d-n-e-s-s>

<http://harrison2.chem.utk.edu/~rjh/madness/>





# Why MADNESS?

- Reduces S/W complexity
  - MATLAB-like level of composition of scientific problems with guaranteed speed and precision
  - Programmer not responsible for managing dependencies, scheduling, or placement
- Reduces numerical complexity
  - Solution of integral not differential equations
  - Framework makes latest techniques in applied math and physics available to wide audience

E.g., with guaranteed precision of  $1e-6$  form a numerical representation of a Gaussian in the cube  $[-20,20]^3$ , solve Poisson's equation, and plot the resulting potential  
(all running in parallel with threads+MPI)

Let

$$\Omega = [-20, 20]^3$$

$$\epsilon = 1e - 6$$

$$g = x \rightarrow \exp(- (x_0^2 + x_1^2 + x_2^2)) * \pi^{-1.5}$$

In

$$f = \mathcal{F} g$$


$$u = \nabla^{-2} (-4 * \pi * f)$$

```
print "norm of f", <f>, "energy", <f|u> * 0.5
```

```
plot u
```

End

*output:* norm of f 1.00000000e+00 energy 3.98920526e-01



There are only two lines doing real work. First the Gaussian ( $g$ ) is projected into the adaptive basis to the default precision. Second, the Green's function is applied. The exact results are norm=1.0 and energy=0.3989422804.



Let

$$\Omega = [-20, 20]^3$$

$$r = x \rightarrow \sqrt{x_0^2 + x_1^2 + x_2^2}$$

$$g = x \rightarrow \exp(-2 * r(x))$$

$$v = x \rightarrow -\frac{2}{r(x)}$$

In

$$\nu = \mathcal{F} v$$

$$\phi = \mathcal{F} g$$

$$\lambda = -1.0$$

for  $i \in [0, 10]$

$$\phi = \phi * \|\phi\|^{-1}$$

$$V = \nu - \nabla^{-2} (4 * \pi * \phi^2)$$

$$\psi = -2 * (-2 * \lambda - \nabla^2)^{-1} (V * \phi)$$

$$\lambda = \lambda + \frac{\langle V * \phi | \psi - \phi \rangle}{\langle \psi | \psi \rangle}$$

$$\phi = \psi$$

print "iter", i, "norm",  $\|\phi\|$ , "eval",  $\lambda$

end

End

# He atom Hartree-Fock

Compose directly in terms of functions and operators

This is a Latex rendering of a program to solve the Hartree-Fock equations for the helium atom

The compiler also output a C++ code that can be compiled without modification and run in parallel

# “Fast” algorithms

- Fast in mathematical sense
  - Optimal scaling of cost with accuracy & size
- Multigrid method – Brandt (1977)
  - Iterative solution of differential equations
  - Analyzes solution/error at different length scales
- Fast multipole method – Greengard, Rokhlin (1987)
  - Fast application of dense operators
  - Exploits smoothness of operators
- Multiresolution analysis
  - Exploits smoothness of operators and functions

# The math behind the MADNESS

- Multiresolution


$$V_0 \subset V_1 \subset \dots \subset V_n$$
$$V_n = V_0 + (V_1 - V_0) + \dots + (V_n - V_{n-1})$$

- Low-separation rank

$$f(x_1, \dots, x_n) = \sum_{l=1}^M \sigma_l \prod_{i=1}^d f_i^{(l)}(x_i) + O(\epsilon)$$
$$\|f_i^{(l)}\|_2 = 1 \quad \sigma_l > 0$$

- Low-operator rank

$$A = \sum_{\mu=1}^r u_{\mu} \sigma_{\mu} v_{\mu}^T + O(\epsilon)$$
$$\sigma_{\mu} > 0 \quad v_{\mu}^T v_{\lambda} = u_{\mu}^T u_{\lambda} = \delta_{\mu\nu}$$



# Three equivalent representations

- Scaling function basis (reconstructed)

$$f^n(x) = \sum_{l=0}^{2^n-1} \sum_{i=0}^{k-1} s_{il}^n \phi_{il}^n(x)$$


- Multi-wavelet basis (compressed)

$$f^n(x) = \sum_{i=0}^{k-1} s_{i0}^0 \phi_{i0}^0(x) + \sum_{n'=0}^{n-1} \sum_{l=0}^{2^{n'}-1} \sum_{i=0}^{k-1} d_{il}^{n'} \psi_{il}^{n'}(x)$$

- Rapid compression/reconstruction
- Values at Gauss-Legendre points in each box
- Use appropriate basis for a given operation



# Molecular Electronic Structure



Energy and  
gradients

ECPs coming  
(Sekino, Kato)


Response  
properties  
(Vasquez and  
Sekino)

Still not as  
functional as  
previous  
Python version

*Spin density  
of solvated  
electron*

# Solid-state electronic structure

- Thornton, Egiluz and Harrison (UT/ORNL)
  - NSF OCI-0904972: Computational chemistry and physics beyond the petascale
- Full band structure with LDA and HF for periodic systems
- In development: hybrid functionals, response theory, post-DFT methods such as GW and model many-body Hamiltonians via Wannier functions



Coulomb potential isosurface in LiF

# Nuclear physics

J. Pei, G.I. Fann, W. Thornton

W. Nazarewicz

UT/ORNL

UNEDF Deformed SLDA in 3-D  
deformation in an external trap with aspect ratio of 1/16

- DOE UNDEF
- Nuclei & neutron matter
- ASLDA
- Hartree-Fock Bogliobulov
- Spinors
- Gamov states




Figure 4. Pairing density  $\kappa(x, y, z)$ , calculated by MADNESS-HFB for the elongated trap with  $\eta = 16$ . The box scale in view is  $x[-12, 12]$ ,  $y[-12, 12]$  and  $z[-32, 32]$ .




# BGQ Early Science Project Activities


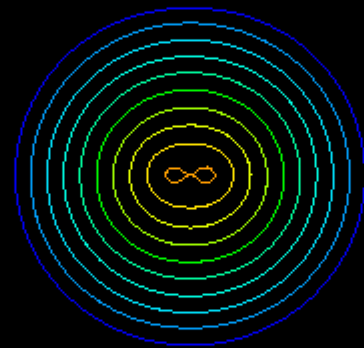
- We are testing a new linear-response module to solve TDDFT equations.
- Molecular properties (dipole polarizabilities, NMR chemical shifting, etc.)
- Support for well known hard pseudopotentials (i.e. Krack, Goedecker, etc.)
- Speedup of Hartree-Fock exchange evaluations via screening parameters.
- Implementation of new DFT functionals.
- Improving parallel scalability for current supercomputer architectures.

# UNEDF Solving HFB in Production

MADNESS-HFB Scalability Test in Jaguar XT5



Realizable scaling even for small problems in 3-D




Time  
dependent  
electronic  
structure

Vence,  
Krstic,  
Harrison  
UT/ORNL

H<sub>2</sub><sup>+</sup> molecule  
in laser field  
(fixed nuclei)

DB: wf3D-06543.vts  
Cycle: 6543

Volume  
Var: LogAbsRe  
-1.500  
-2.625  
-3.750  
-4.875  
-6.000  
Max: -0.6685  
Min: -8.000

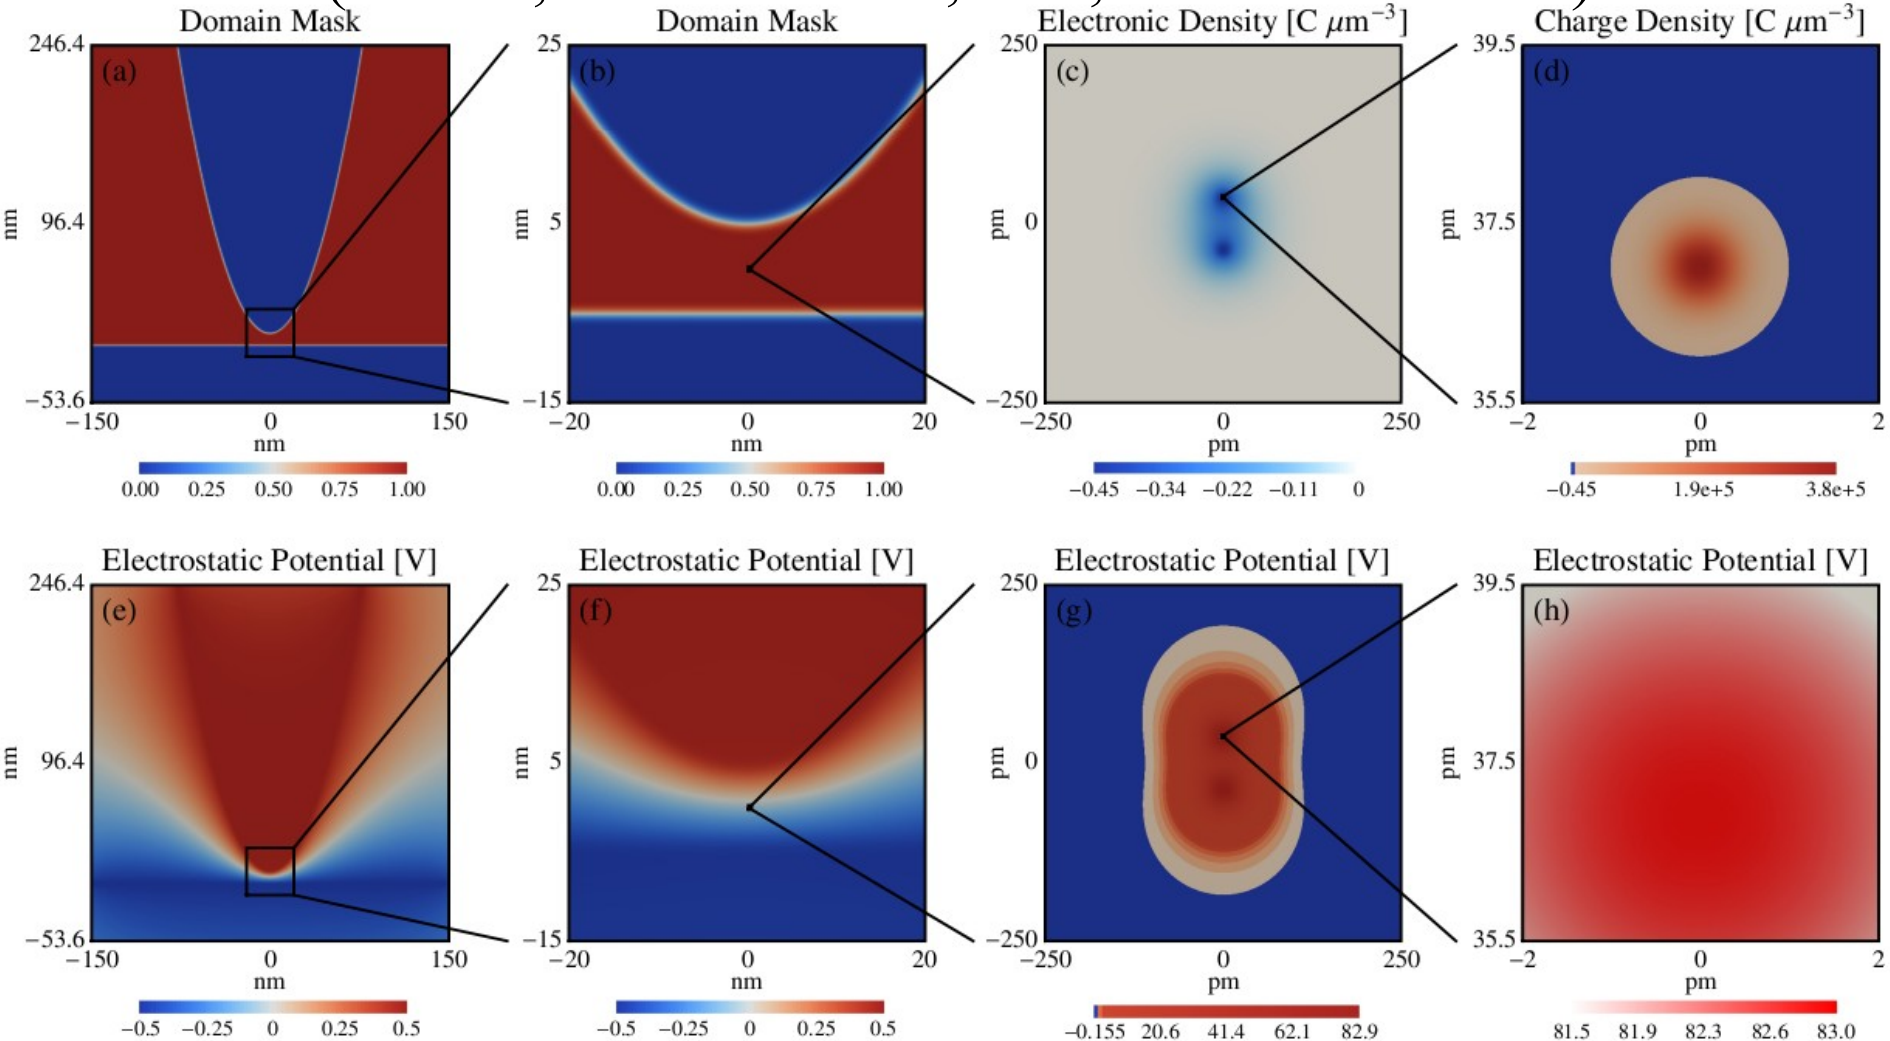


# H-atom IR pulse

The electric field repeatedly rips the electron out  $\sim 300$  Bohr and brings it back to rescatter off the nucleus.


# Nanoscale photonics

(Reuter, Northwestern; Hill, Harrison ORNL)



Diffuse domain approximation for interior boundary value problem; long-wavelength Maxwell equations; Poisson equation; Micron-scale Au tip 2 nm above Si surface with H<sub>2</sub> molecule in gap –  $10^7$  difference between shortest and longest length scales.

# Electron correlation (6D)



- All defects in mean-field model are ascribed to electron correlation
- Singularities in Hamiltonian imply for a two-electron atom

$$\Psi(r_1, r_2, r_{12}) = 1 + \frac{1}{2} r_{12} + \dots \quad \text{as } r_{12} \rightarrow 0$$

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

$$\Psi(r_1, r_2, r_{12}) = \exp(-\xi(r_1 + r_2)) (1 + a r_{12} + \dots)$$

- Potentially very accurate, but not systematically improvable, and (until recently) not computationally feasible for many-electron systems
- Configuration interaction expansion – slowly convergent

$$\Psi(r_1, r_2, \dots) = \sum_i c_i \left| \phi_1^{(i)}(r_1) \phi_2^{(i)}(r_2) \dots \right|$$

# Partitioned SVD representation

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

$r$  = separation rank

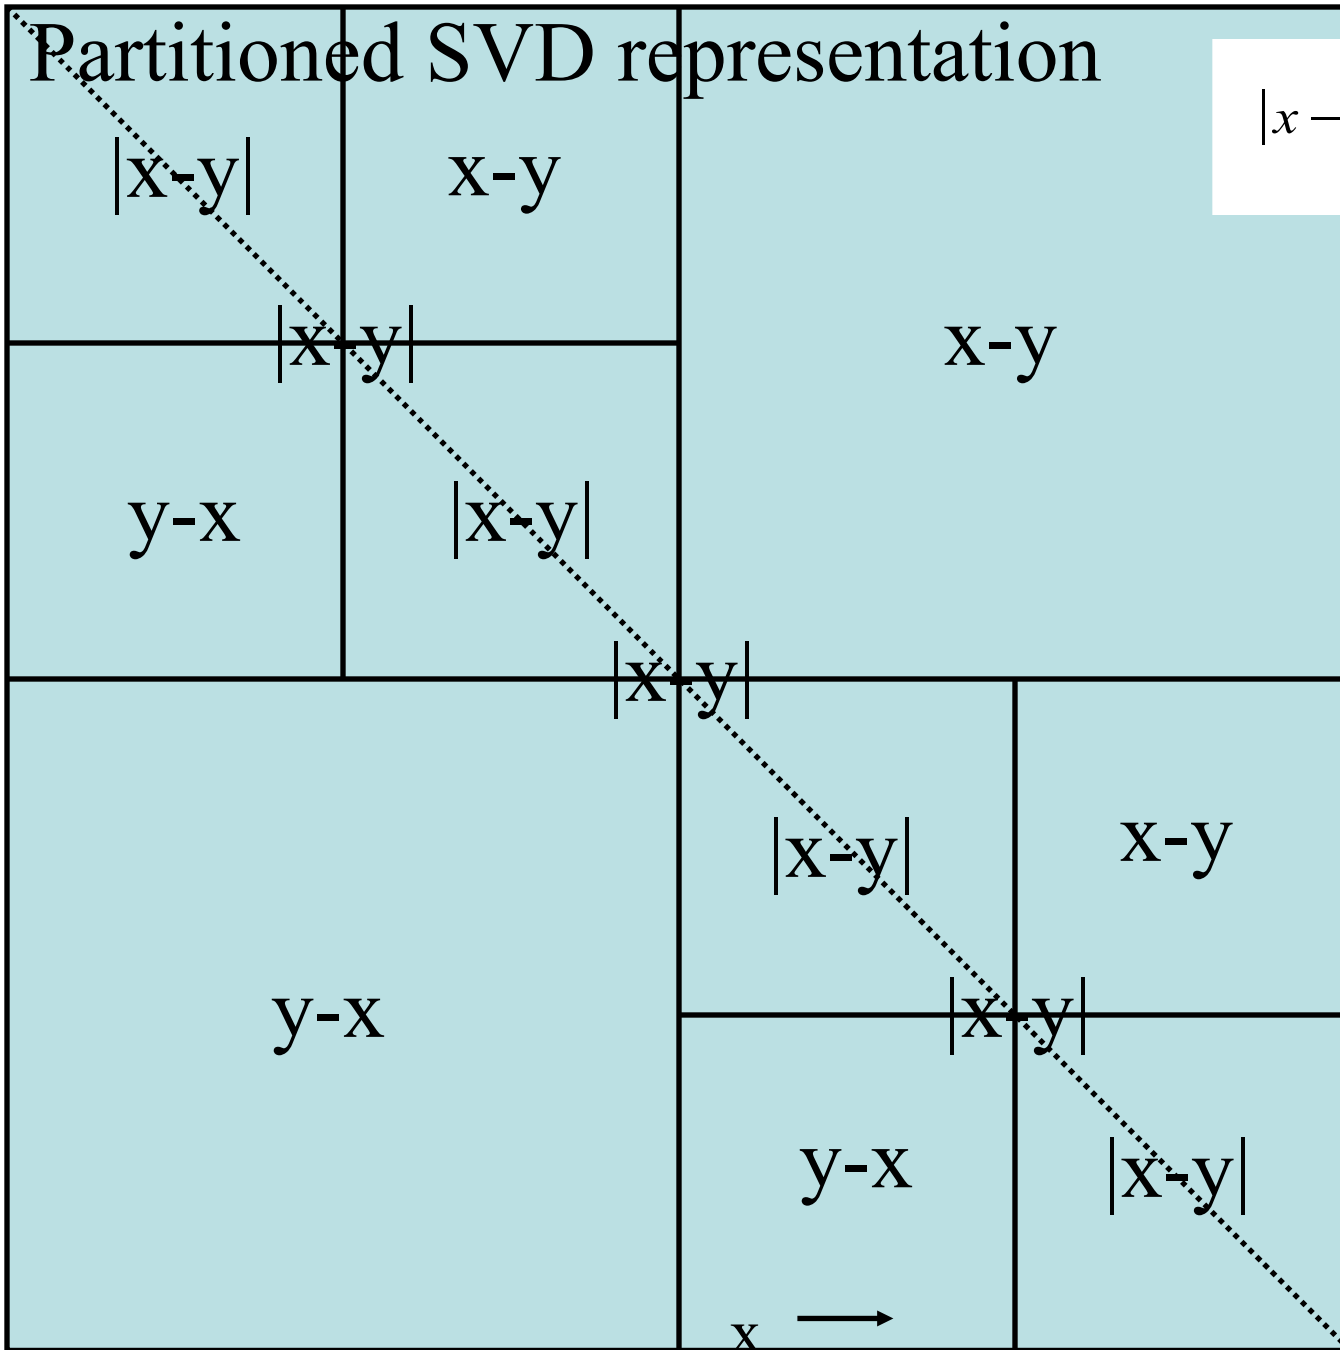
In 3D, ideally must be one box removed from the diagonal

Diagonal box has full rank

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

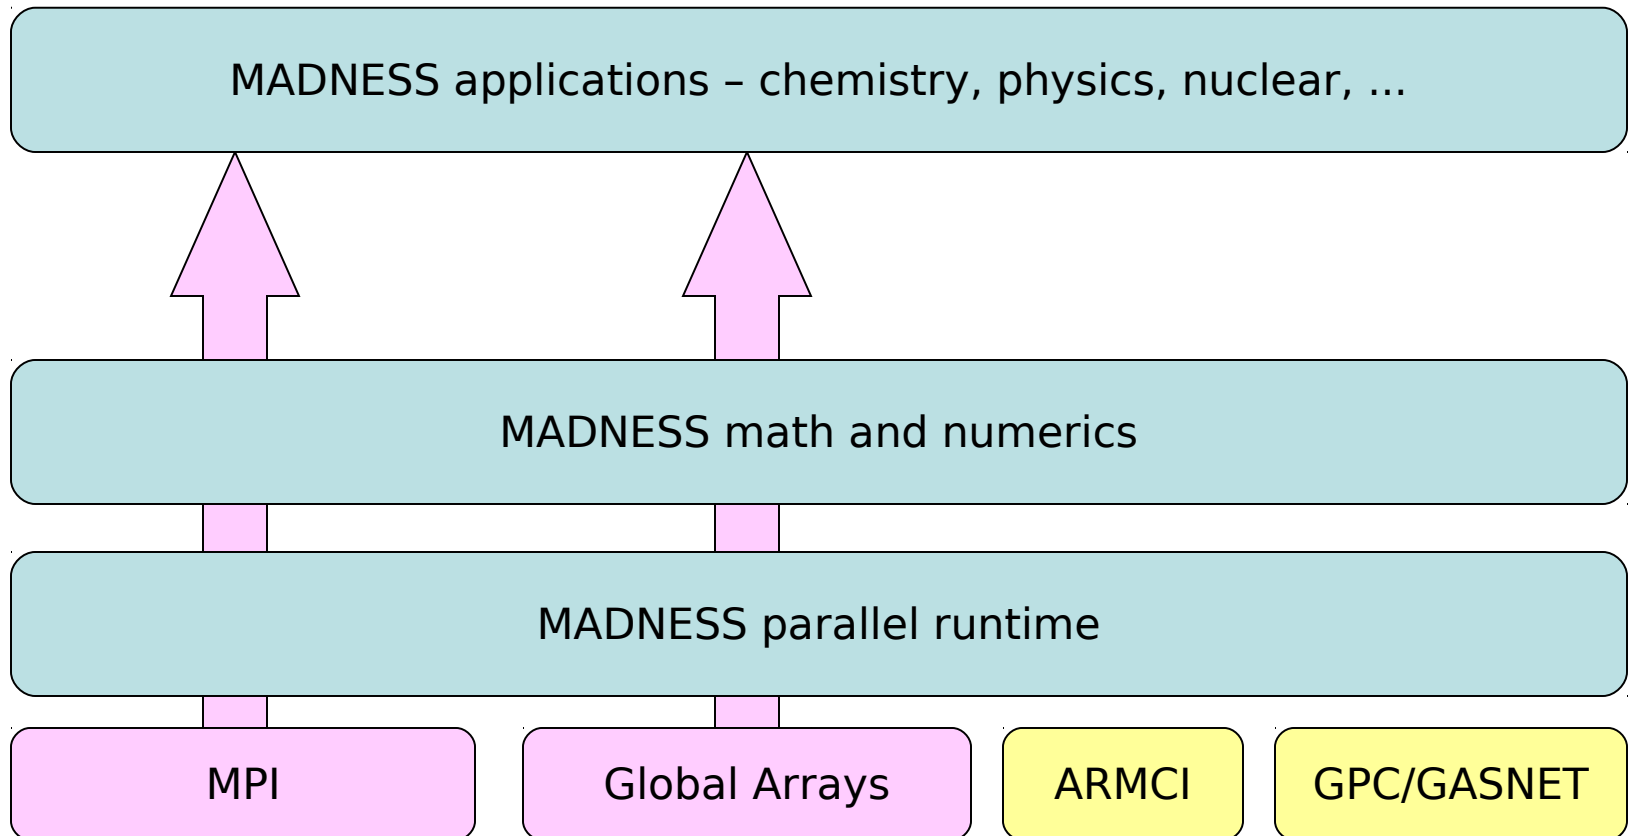
Away from diagonal  $r = O(-\log \epsilon)$

$y$   
↓



$x$  →

# MADNESS architecture



Intel Thread Building Blocks more scalable; also ported to BGQ



# Runtime Objectives

- Scalability to 1+M processors ASAP
- Runtime responsible for
  - scheduling and placement,
  - managing dependencies & hiding latency
- Compatible with existing models (MPI, GA)
- Borrow successful concepts from Cilk, Charm++, Python, HPCS languages

# Why a new runtime?

- MADNESS computation is irregular & dynamic
  - 1000s of dynamically-refined meshes changing frequently & independently (to guarantee precision)
- Because we wanted to make MADNESS itself easier to write not just the applications using it
  - We explored implementations with MPI, Global Arrays, and Charm++ and all were inadequate
- MADNESS is helping drive
  - One-sided operations in MPI-3, DOE projects in fault tolerance, ...

# Key runtime 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 compatibility
- Dynamic load balancing
  - Data redistribution, work stealing, randomization

# Futures

- Result of an asynchronous computation

- Cilk, Java, HPCLs, C++0x

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

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

- Hide latency due to communication or computation

```
// Work until need result
```

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

- Management of dependencies

- Via callbacks

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 Names

- Objects with global names with different state in each process
  - C.f. shared[threads] in UPC; co-Array
- Non-collective constructor; deferred destructor
  - Eliminates synchronization

```
class A : public WorldObject<A>
{
    int f(int);
};
ProcessID p;
A a;
Future<int> b =
    a.task(p, &A::f, 0);
```

A task is sent to the instance of a in process p. If this has not yet been constructed the message is stored in a pending queue. Destruction of a global object is deferred until the next user synchronization point.

# 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);
```

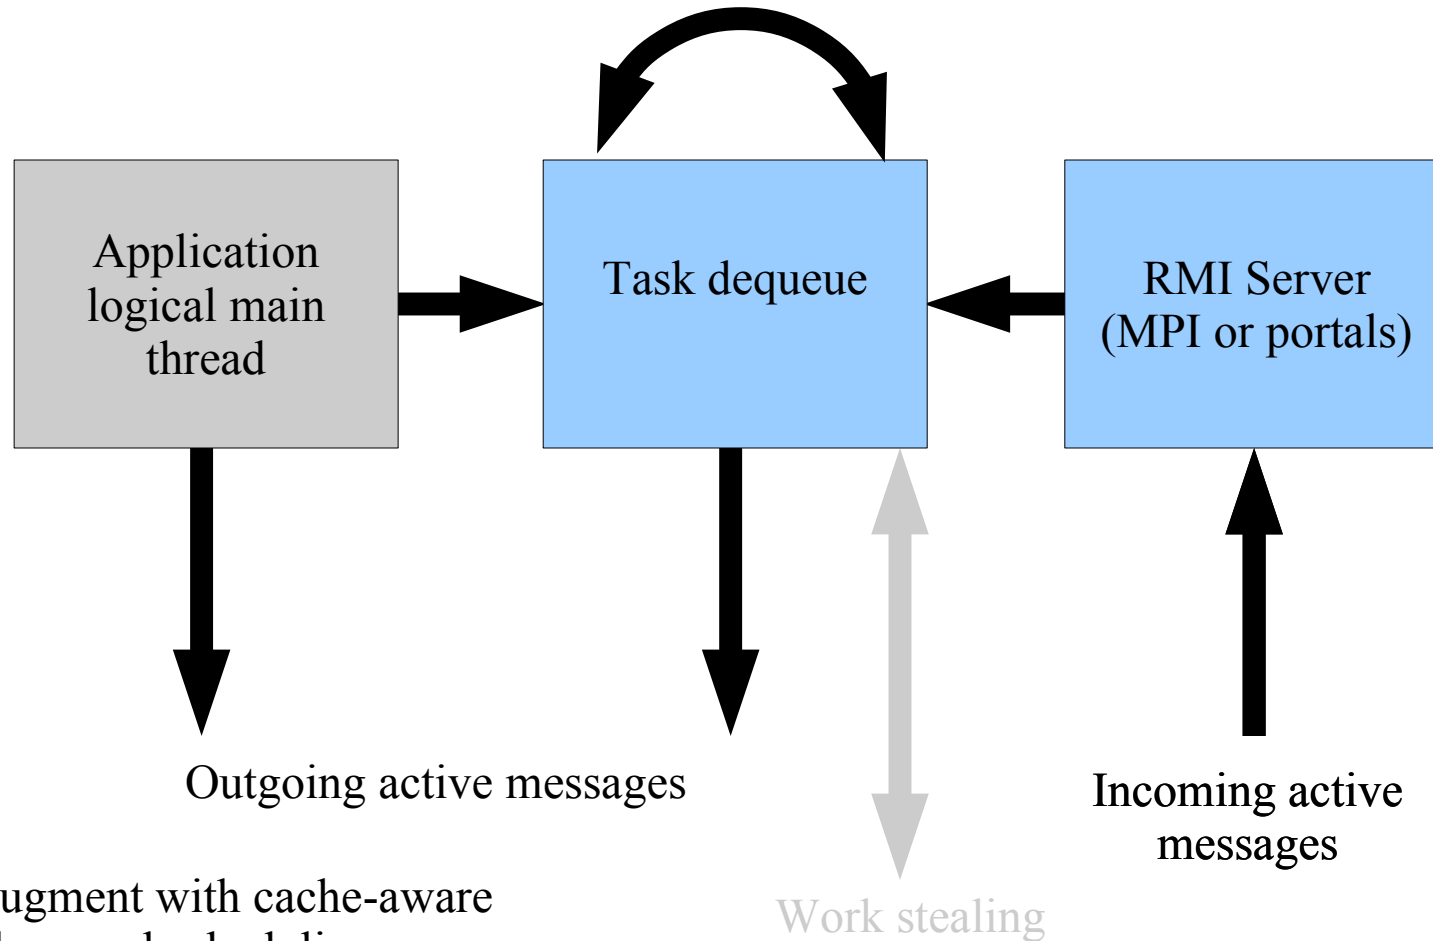
A container is created mapping indices to values.

A value is inserted into the container.

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



# Multi-threaded architecture



Must augment with cache-aware algorithms and scheduling


# Computational kernels

- Discontinuous spectral element
  - In each “box” a tensor product of coefficients
  - Most operations are small matrix-multiplication

$$r_{i'j'k'} = \sum_{ijk} s_{ijk} c_{ii'} c_{jj'} c_{kk'} = \sum_k \left( \sum_j \left( \sum_i s_{ijk} c_{ii'} \right) c_{jj'} \right) c_{kk'}$$
$$\Rightarrow r = ((s^T c)^T c)^T c$$


- Typical matrix dimensions are 2 to 30
- E.g.,  $(20,400)^T * (20,20)$

# Comparison with MKL, Goto, ATLAS on Intel Xeon 5355 for $(20,400)^T * (20,n)$ .




# MtXM performance on BGQ

(n,n)\*(n,n) small matrix multiply various thread counts



transform(n,n,n) various thread counts



## Kevin Stocks OSU

64 threads, best performance is  
139.7 GFLOPS (trans(400,20,20))

\* Theoretical peak is 204.8


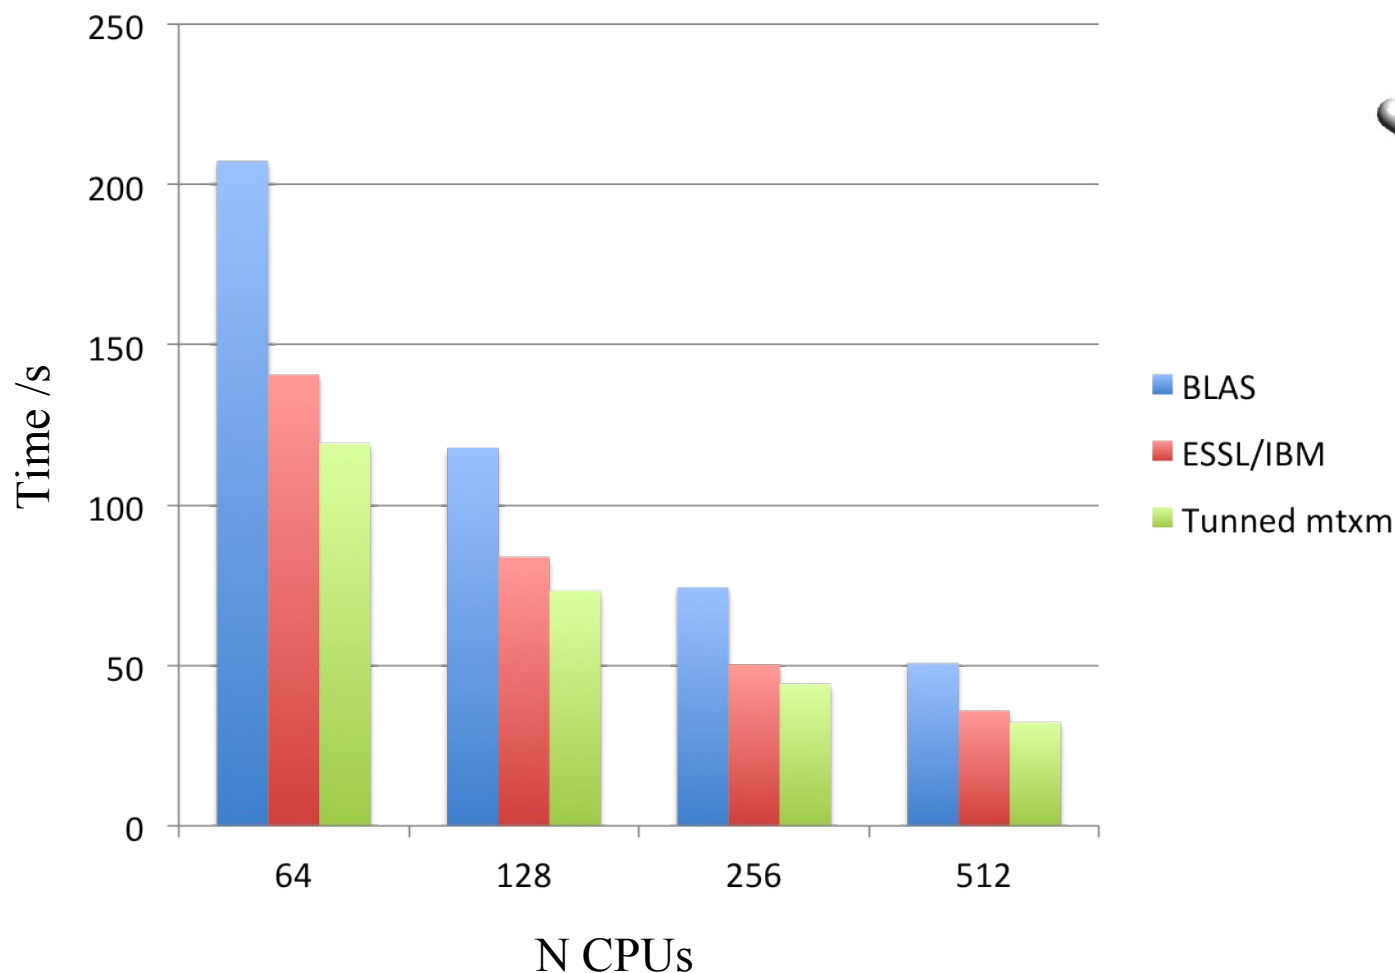
\* Linpack is approx. 166.3

(scaling top500 results to one node)

# Benefit of tuned mTxm in BG/P Performance

## Strong scaling

Molecular system with 13 heavy atoms, DFT, k=8, one iteration



# Summary

- Exascale programming models
  - Resilience, Power, Performance, Productivity
  - Productivity is arguably the most important
  - Enable innovation and discovery at scale
- MADNESS and NWChem
  - Frameworks – places for disciplines to meet to leverage investments and expertise
  - Face difference challenges in moving forward