Progress in MFDn: Mathematics and Computer Science

Philip Sternberg

in collaboration with

Esmond Ng
Chao Yang
James Vary
Pieter Maris
Masha Sosonkina
Hung Le

Progress Since Last Year

- Parallel speedup of basis space generation
- Hierarchical scheme for computing sparsity of Hamiltonian
- Use OpenMP to take advantage of multicore architectures
- Hybrid storage scheme for Hamiltonian
- Integration of MFDn with optimization techniques
- Improved inner loop efficiency

Parallel Basis Generation

To balance load, basis states are assigned to processors cyclically mod n_{diag}

Testing a basis state for m_j, parity, and oscillator quanta conditions takes much more time than getting the next odometer-ordered state

New algorithm divides number of validity tests by n_{diag}

Time to generate basis states comparable to reading from disk

Finding Nonzeros

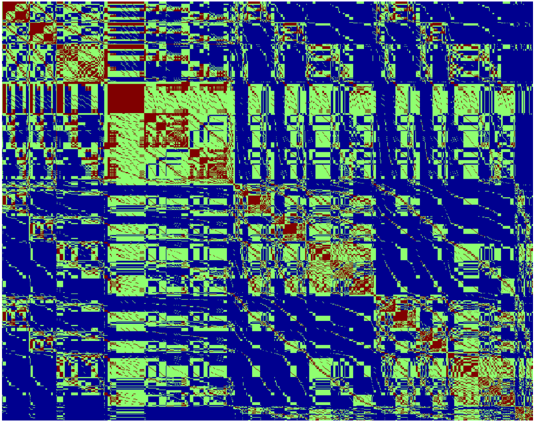

- nonzeros
- potentially nonzero blocks
- zero blocks

Tradeoff between many fine blocks and few coarse blocks

Multilevel Blocking (New Algorithm)

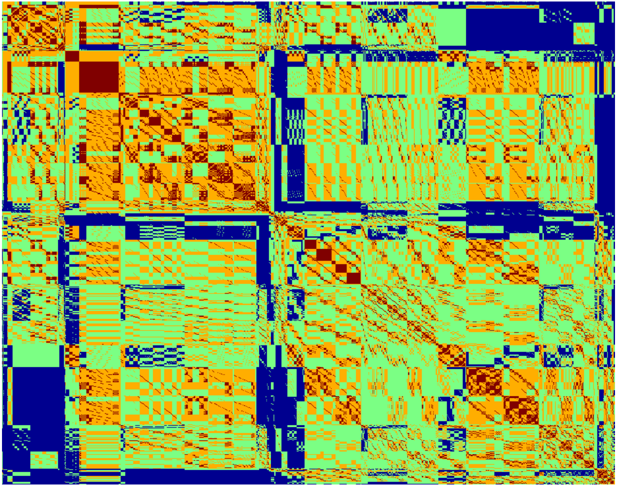

Any number of levels can be used

Performance Results

Time to compute sparsity:

	one level	multiple levels
^{6}He (small example)	180 seconds	90 seconds
^{12}C, Nmax = 8, 2-body	$\sim$ 1 hour	$\sim$ 13 minutes
^{16}O, Nmax = 8, 2-body	$\sim$ 2 hours	$\sim$ 20 minutes
^{14}F, Nmax = 8, 2-body	$\sim$ 8 hours ?	$\sim$ 35 minutes

Time to compute sparsity grows as $O(\text{nnz})$ ($\sim O(\text{n}^{1.5})$), rather than $O(\text{n}^2)$.

Thus, as problems get sparser, the advantage of multiple levels increases.

Performance Improvement in Separate Parts of MFDn

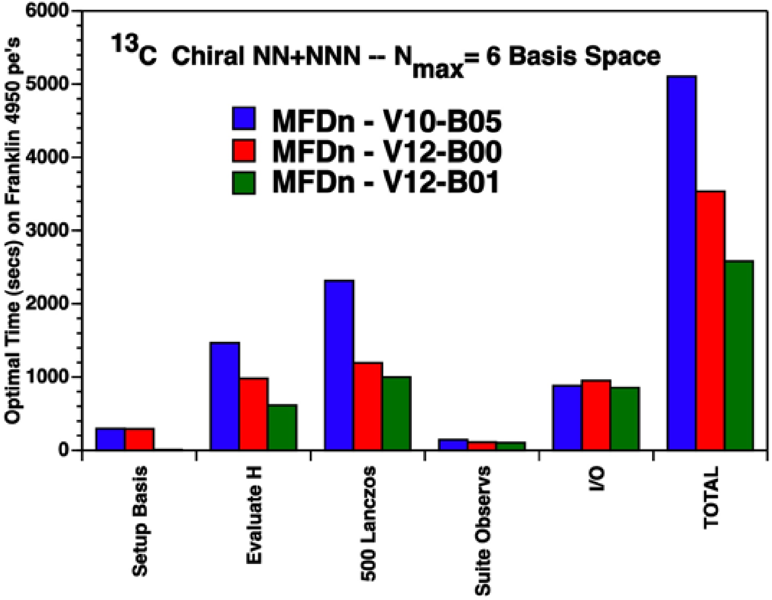

Using Threads for Multicore

Many current supercomputers are multicore (XT4, BG/P)

Cores/Node are increasing, so to achieve best performance, all codes will need to take advantage of multiple cores, using hybrid of distributed and shared memory model

Using OpenMP, have observed following initial speed-ups on quad-core system (Jaguar)

- $\sim 3\times$ speedup to evaluate matrix elements
- $\sim 2\times$ speedup to perform matrix-vector multiply
- $\sim 3\times$ speedup to evaluate suite of observables

Further work needed to integrate into production code

Hybrid Storage of Hamiltonian

We can store dense blocks as contiguous values (filling in with some zeros), with integers indicating location and size.

$A =$

1.0	2.0		3.0	
4.0	5.0			6.0
		7.0	8.0	
9.0		1.0	0.0	
				2.0

Many nonzeros are in dense blocks, using much less memory to store the Hamiltonian and taking less time to perform matvec.

Many Dense Blocks

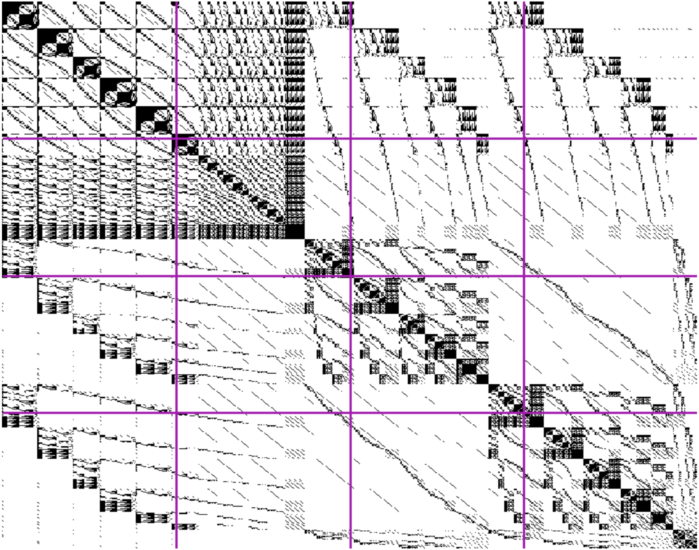

Dense Blocks Broken Up by Cyclic Assignment

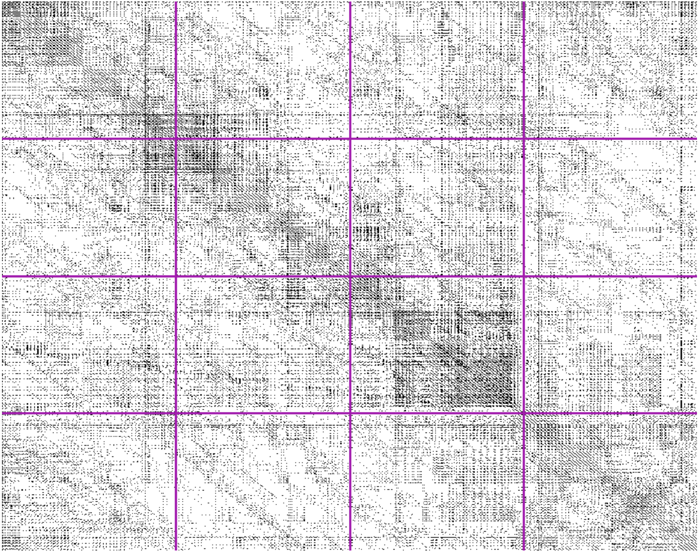

Future work

- Integrate load-balanced reorthogonalization with PARPACK
  - Current PARPACK-enabled MFDn does not use all processors to reorthogonalize
  - Will modify PARPACK to use MFDn's distribution scheme for Lanczos vectors so PARPACK will exhibit same scaling properties as current production code
- Investigate alternative eigensolvers (LOBPCG)
- Alternative partitioning strategy to keep dense blocks together using same hierarchical scheme as for computing sparsity
- Write highly tuned variant of MFDn for cases with dim $10^7 - 10^8$ as will be used for external fields
- Engineer modifications to code as production scale problem size increases

Papers

- M. Sosonkina, A. Sharda, A. Negoita, J.P. Vary, Integration of Ab Initio Nuclear Physics Calculations with Optimization Techniques. International Conference on Computational Science. In Proc. 8th International Conference, Kraków, Poland (June 23-25, 2008), Lecture Notes in Computer Science, v. 5101. pp.833-842.
- P. Sternberg, P. Maris, E.G. Ng, M. Sosonkina, H.V. Le, J.P. Vary, C. Yang, “Accelerating Full Configuration Interaction Calculations for Nuclear Structure”. Submitted to SC08, Austin (November 15-21, 2008).
- P. Sternberg, P. Maris, E.G. Ng, J.P. Vary, C. Yang, “High Performance Nuclear Structure Calculation”. SIAM Journal of Scientific Computing, In Preparation.

- J.P. Vary, A.M. Shirokov, P. Maris, “Light nuclei without a core,” to be published in the Proceedings of the XXX Reuniao de Fisica Nuclear no Brasil, August 2007. arXiv:0804.0836.
- S.K. Bogner, R.J. Furnstahl, P. Maris, R.J. Perry, A. Schwenk, J.P. Vary, “Convergence in the no-core shell model with low-momentum two-nucleon interactions,” Nuclear Physics A 801, 21(2008); arXiv:0708.3754.
- A. F. Lisetskiy, B. R. Barrett, M.K.G. Kruse, P. Navratil, I. Stetcu, and J. P. Vary, “Ab-initio shell model with a core” Submitted to Physical Review C, June 2008.
- A. M. Shirokov, A. I. Mazur, J. P. Vary, and E. A. Mazur, “Inverse scattering J-matrix approach to nucleon-nucleus scattering and the shell model”, To be submitted to the ArXiv this week.
- J.P. Vary, P. Maris, A. Negoita, P. Navratil, V.G. Gueorguiev, W. E. Ormand, A. Nogga, A. Shirokov and S. Stoica, “Elements of the ab initio No Core Shell Model,” in Exotic Nuclei and Nuclear/Particle Astrophysics (II), Proceedings of the Carpathian Summer School of Physics 2007, L. Trache and S. Stoica, Editors, AIP Conference Proceedings 972, 49(2008).