Parallel Eigensolver Performance on High Performance Computers

Andrew Sunderland
Advanced Research Computing Group
STFC Daresbury Laboratory

Summary

- (Briefly) Introduce parallel diagonalization methods
  - Application Background
  - Numerical Library Routines
- Parallel (MPI) Performance Analysis
  - For
    - Applications
    - Eigensolver methods
    - High-end Architectures
- Petascaling strategies

Interest in Eigensolver Performance

- Demand for efficient eigensolver performance
  - Traditionally a huge computational bottleneck
    - Especially Electronic Structure codes, Atomic Molecular Codes, large-scale problems.
      - At STFC this includes PRMAT, CRYSTAL, GAMESS-UK, KPPW
    - Large amounts of memory, cpu time and communication required
  - Researchers are not always using most efficient methods
    - Several methods to choose from
      - New developments with Divide and Conquer and Multiple Relatively Robust Representations
    - Performance is application dependent (differs from linear systems)
      - $\alpha O(n^3)$, where α can vary

Standard Eigenvalue problem

Standard Eigenvalue Problem:

$$\mathbf{A}\,\mathbf{x} = \lambda\,\mathbf{x}$$

… where **A** is a dense real matrix and λ is an eigenvalue corresponding to eigenvector **x**.

Diagonalization:

$$\mathbf{A} = \mathbf{S}\,\Lambda\,\mathbf{S}^{-1}$$

… where the columns of matrix S are represented by the eigenvectors of **A** and the diagonal matrix Λ represents the associated eigenvalues

If A is symmetric, then:

$$\mathbf{A} = \mathbf{Q}\,\Lambda\,\mathbf{Q}^{T}$$

Standard Eigenvalue problem

Standard Eigenvalue Problem:

$$\mathbf{A}\,\mathbf{x} = \lambda\,\mathbf{x}$$

… where **A** is a dense real matrix and λ is an eigenvalue corresponding to eigenvector **x**.

Diagonalization:

$$\mathbf{A} = \mathbf{S}\,\Lambda\,\mathbf{S}^{-1}$$

… where the columns of matrix S are represented by the eigenvectors of **A** and the diagonal matrix Λ represents the associated eigenvalues

If A is symmetric, then:

$$\mathbf{A} = \mathbf{Q}\,\Lambda\,\mathbf{Q}^{T}$$

Λ captures key characteristics of A

Underlying Equations

Solve the non-relativistic Schrodinger equation:

$$H\,\Psi = E\,\Psi$$

PRMAT - Atomic & Molecular Physics Application Code

- R-matrix based atomic scattering calculations
- Baluja-Burke-Morgan approach - Diagonalize Representative of the Green's function $(H + L - EI)^{-1}$ within a basis (Hamiltonian Matrix)

CRYSTAL - Electronic Structure of Periodic Systems Code

- Diagonalization of the effective one electron Hamiltonian gives the eigenstates and associated energies that the electrons may occupy
- Periodic Hartree-Fock or density functional Kohn-Sham Hamiltonian matrix
- Eigensolver often part of iterative process

Underlying Equations

Solve the non-relativistic Schrodinger equation:

$$H\,\Psi = E\,\Psi$$

Dense, Real, Symmetric

PRMAT - Atomic & Molecular Physics Application Code

- R-matrix based atomic scattering calculations
- Baluja-Burke-Morgan approach - Diagonalize Representative of the Green's function $(H + L - EI)^{-1}$ within a basis (Hamiltonian Matrix)

CRYSTAL - Electronic Structure of Periodic Systems Code

- Diagonalization of the effective one electron Hamiltonian gives the eigenstates and associated energies that the electrons may occupy
- Periodic Hartree-Fock or density functional Kohn-Sham Hamiltonian matrix
- Eigensolver often part of iterative process

Solving the Symmetric Standard Eigenvalue Problem

The solution to the real symmetric Eigensolver problem usually takes place via three main steps

1. Reduction of the matrix to tri-diagonal form, typically using the Householder Reduction. $O(n^3)$
2. Solution of the real symmetric tri-diagonal Eigenproblem via one of the following methods:
   - Bisection for the Eigenvalues and inverse iteration for the Eigenvectors, up to $O(n^3)$
   - QR algorithm, up to $O(n^3)$
   - Divide & Conquer method (D&C), up to $O(n^3)$
   - Multiple Relatively Robust Representations (MRRR algorithm). O(nk)
3. Back transformation to find Eigenvectors for the full problem. $O(n^3)$

Other Methods

- Jacobi Method $O(n^3)$
- Symmetric Subspace Decomposition Algorithm

Parallel Real Symmetric Eigensolver Routines

Scalapack

- drivers for solving standard and generalized dense symmetric or dense Hermitian Eigenproblems.
- PDSYEV (QR Method) (Scalapack 1.5) (all eigenpairs only)
- PDSYEVX (Bisection & Inverse Iteration) (Scalapack 1.5) (all/subset of eigenpairs)
- PDSYEVD (D&C Method) (Scalapack (1.7) (all eigenpairs)
- PDSYEVR (MRRR Method) Scalapack (1.9?) - current version **under development** (all/subset of eigenpairs) *(Voemel et al.)*
- Memory overheads all ~O(n^2) except PDSYEVR O(n) (in theory)

Peigs

- General symmetric and standard symmetric eigenproblems
  - PDSPEV (Bisection & Inverse Iteration) (all/subset of eigenpairs)

Plapack

- QR method (all eigenpairs only)
- MRRR 'Multiple Relatively Robust Representations' (all/subset of eigenpairs)

Test Matrices

PRMAT

- Sector Hamiltonian matrices from external region Fe+ calculations

Matrix Characteristics

- Real, Symmetric
- Relatively well distributed eigenvalues

Matrix Sizes

- Dimension 10032, 20064, 35904, 63504

CRYSTAL

- Electronic Structure and Related properties of periodic systems
- All electron, local Gaussian basis set, DFT and Hartree-Fock

Matrix Characteristics

- Real, symmetric
- Many degeneracies - closely coupled eigenvalues.
- Varying degrees of sparsity

Matrix Sizes

- Dimension - 7194, 12354, 20480 (Crystal, SCF-calc)

Test Matrices

PRMAT

- Sector Hamiltonian matrices from external region Fe+ calculation

Matrix Characteristics

- Real, Symmetric
- Relatively well distributed eigenvalues

Matrix Sizes

- Dimension 10032, 20064, 35904, 63504

CRYSTAL

- Electronic Structure and Related properties of periodic systems
- All electron, local Gaussian basis set, DFT and Hartree-Fock

Matrix Characteristics

- Real, symmetric
- Many degeneracies - closely coupled eigenvalues.
- Varying degrees of sparsity

Matrix Sizes

- Dimension - 7194, 12354, 20480 (Crystal, SCF-calc)

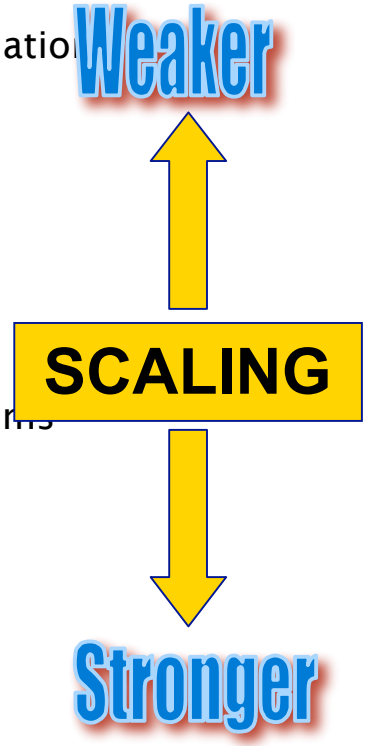

Machines

Cray XT4 *HECToR*
- DC 2.8 GHz Opterons 11328 cores

Cray XT3 *palu CSCS*
- DC 2.6 GHz Opterons 3328 cores

IBM p5-575 *HPCx*
- DC 1.7 GHz POWER5, HPS, 2560 cores

IBM BlueGene/P *Legion, STFC Daresbury*
- QC PowerPC 850 MHz, 4096 cores

Science & Technology Facilities Council

Machines

Libraries

Cray XT4 *HECToR*
- DC 2.8 GHz Opterons 11328 cores

Cray XT3 *palu CSCS*
- DC 2.6 GHz Opterons 3328 cores

Libsci, acml libraries

IBM p5-575 *HPCx*
- DC 1.7 GHz POWER5, HPS, 2560 cores

IBM BlueGene/P *Legion, STFC Daresbury*
- QC PowerPC 850 MHz, 4096 cores

Scalapack, pessl, essl

Parallel Performance of methods (CRYSTAL n=3888)

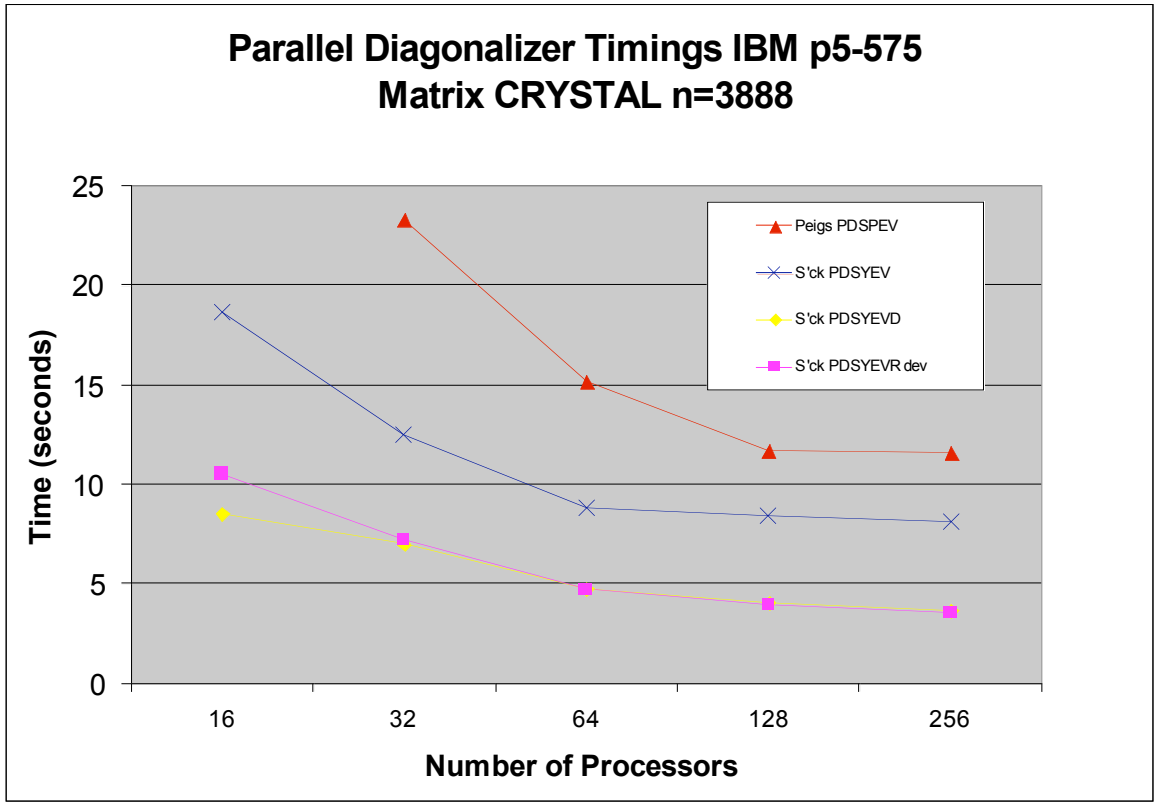

Relative Scaling with Matrix Size (PRMAT)

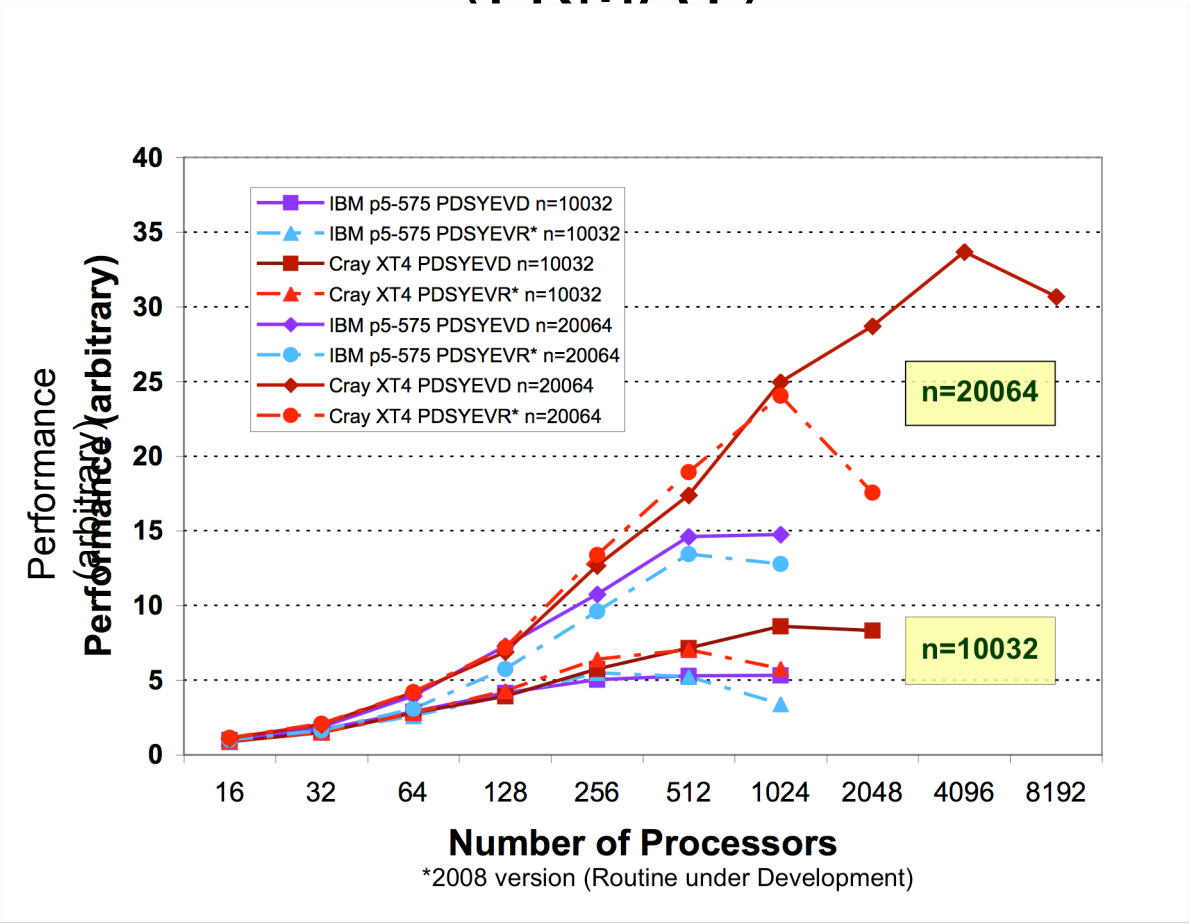

Breakdown of Timings within the eigensolver (CRYSTAL n=12354)

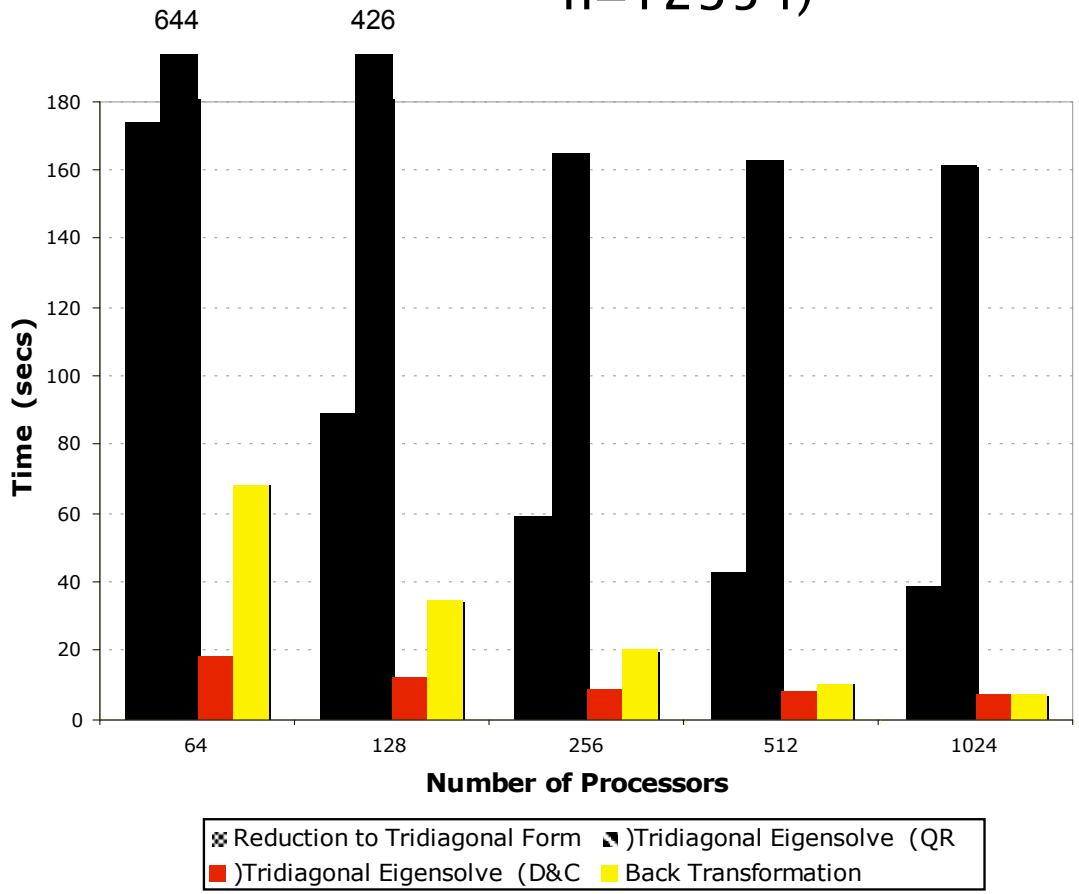

Relative Scaling with Matrix Size (PRMAT, large cases)

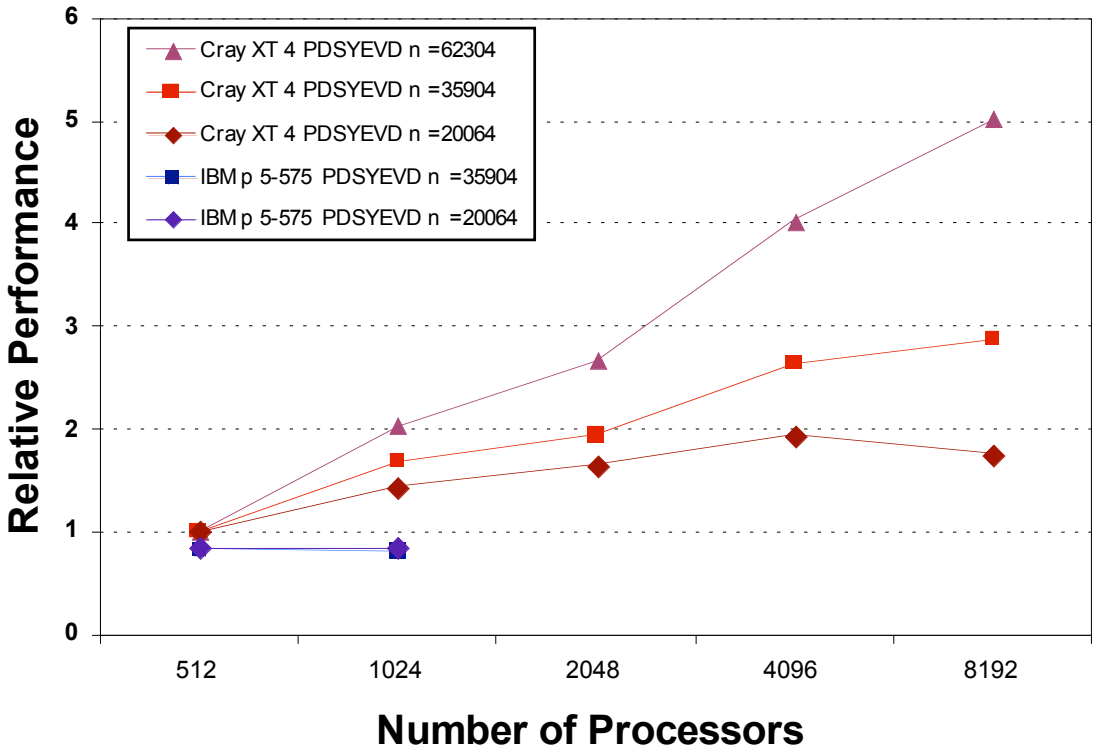

n=~62K

n=~35K

n=~20K

Science & Technology Facilities Council

Relative PDSYEVD Performance on HPC platforms

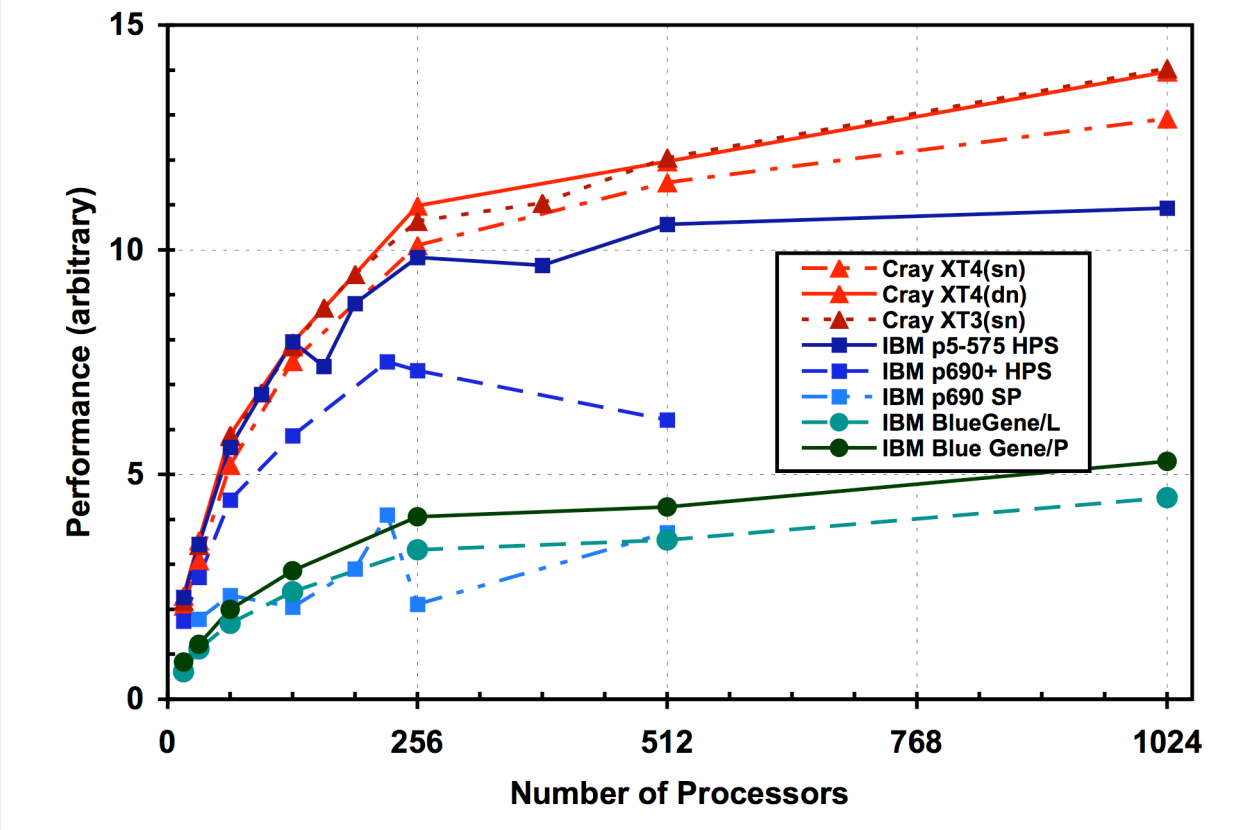

CRYSTAL
matrix n=7194

Model	Ncpus	CPU	Interconnect	Site
Cray XT4	11328	Opteron 2.8 GHz	SeaStar	UoE, UK
IBM	2560	POWER5 1.5 GHz	HPS	STFC, UK
IBM BG/P	4096	POWERPC 850 MHz	3D Torus	STFC, UK

Relative PDSYEVD Performance on HPC platforms

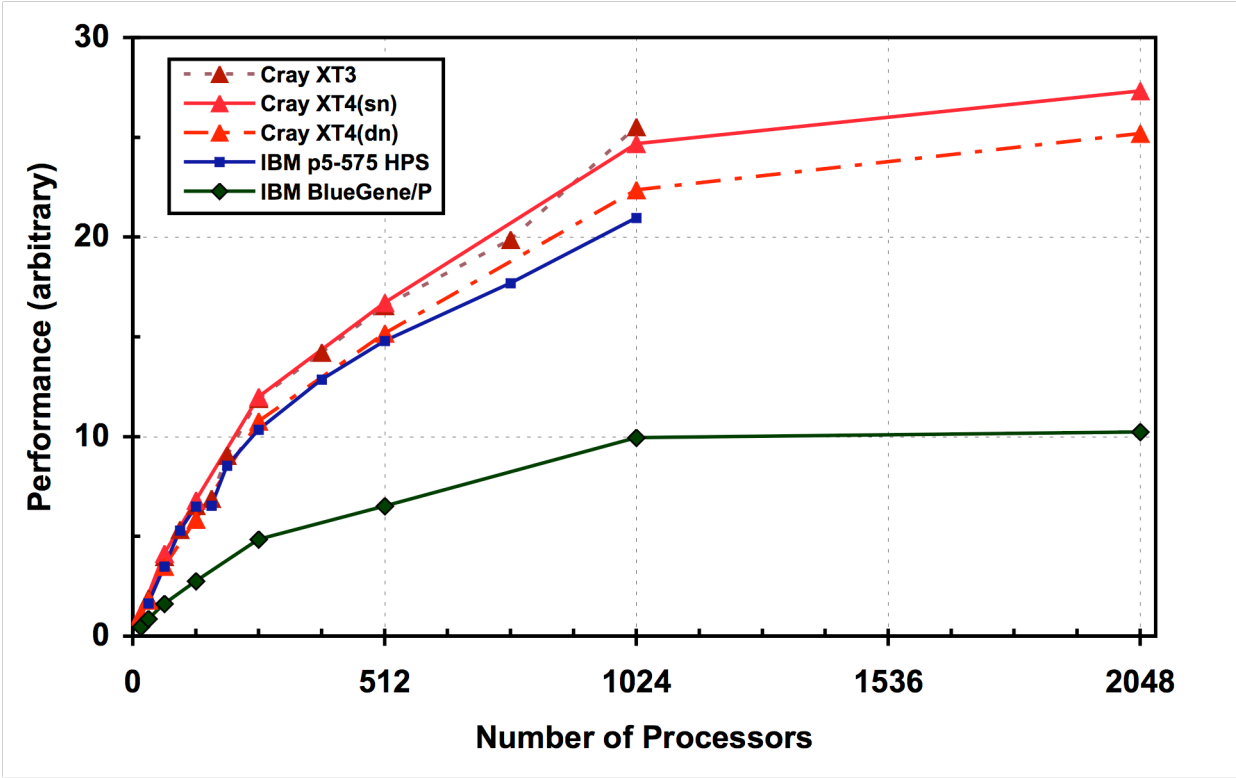

CRYSTAL
matrix n=20480

Model	Ncpus	CPU	Interconnect	Site
Cray XT4	11328	Opteron 2.8 GHz	SeaStar	UoE, UK
IBM	2560	POWER5 1.5 GHz	HPS	STFC, UK
IBM BG/P	4096	POWERPC 850 MHz	3D Torus	STFC, UK

Green Top 500 Supercomputer List

Machine	Mflops/watt	Green Top 500 Ranking
IBM BG/P (STFC)	357.23	1
IBM BG/L (STFC)	215	6
Cray XT3 (CSCS)	56.58	145
IBM p5-575 (HPCx)	30.96	403
Cray XT4 (HECToR)	21.01	460

http://www.green500.org/lists/2007/11/green500_200711.xls

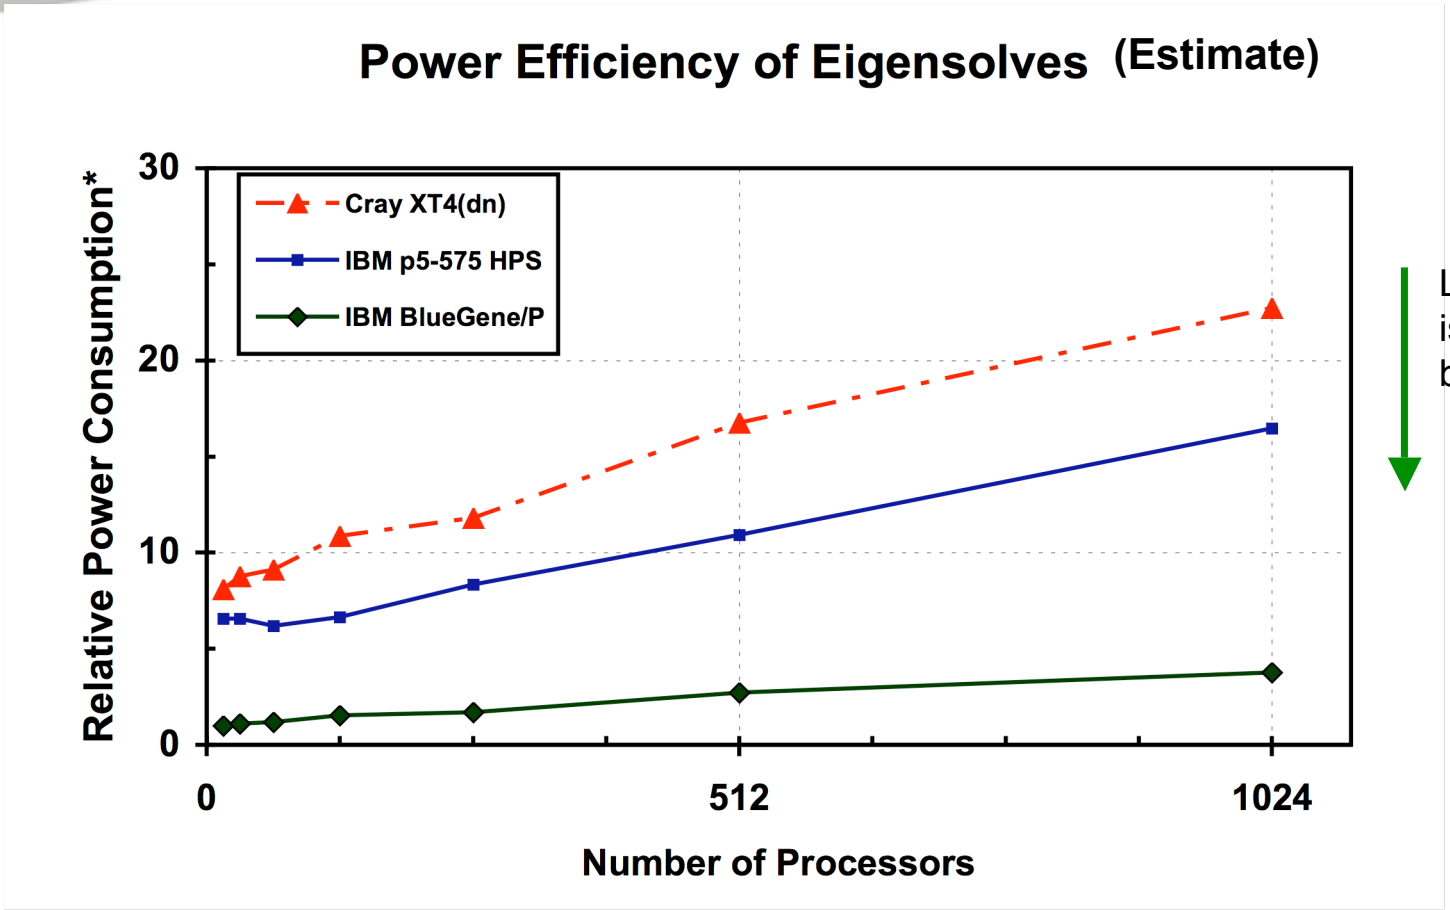

*Arbitrary, relative to to BG/P 16 procs

Petascaling PRMAT: Partition of Configuration Space

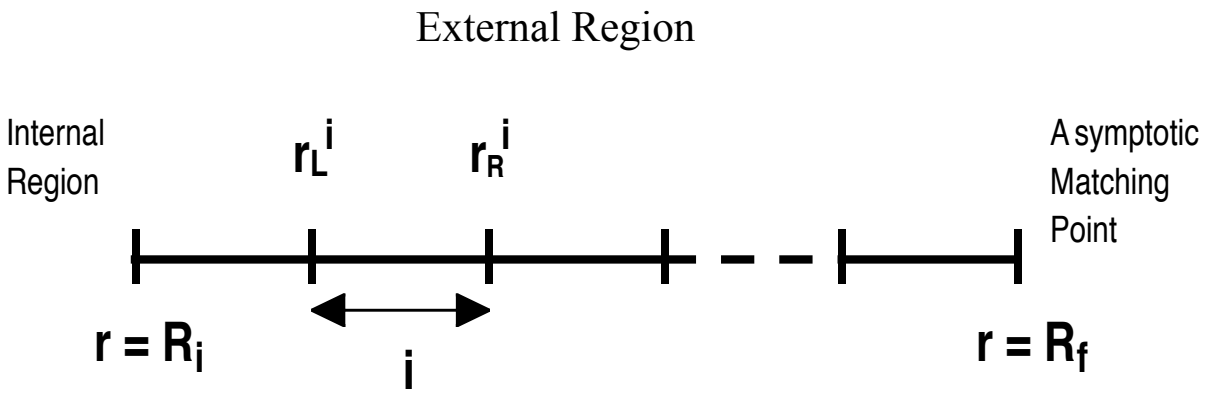

Parallel Diagonalization of large Hamiltonian matrix required for each sector

- Sector diagonalizations can be undertaken simultaneously on sub-groups of processors.

Petascaling Diagonalization Calculations in PRMAT

Sector Hamiltonian matrices, n=20064

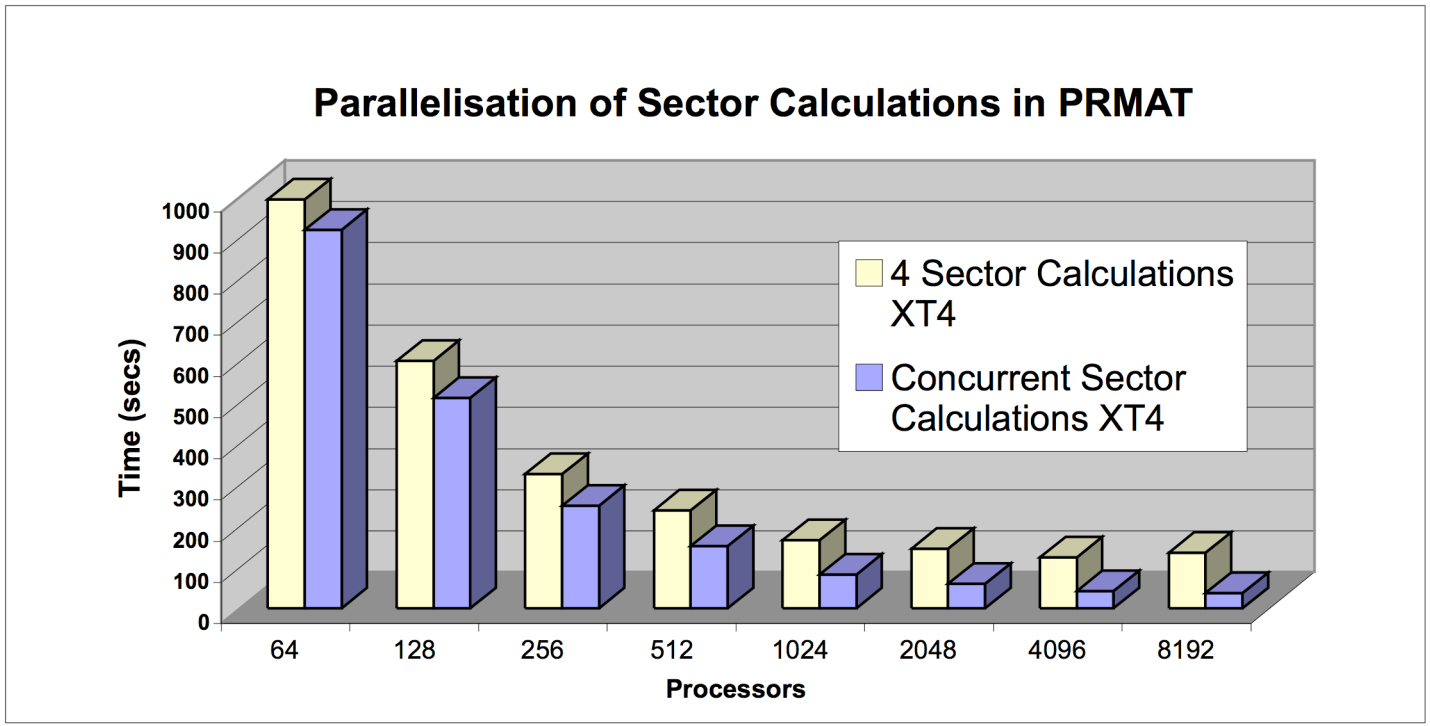

Lower is better

Petascaling CRYSTAL

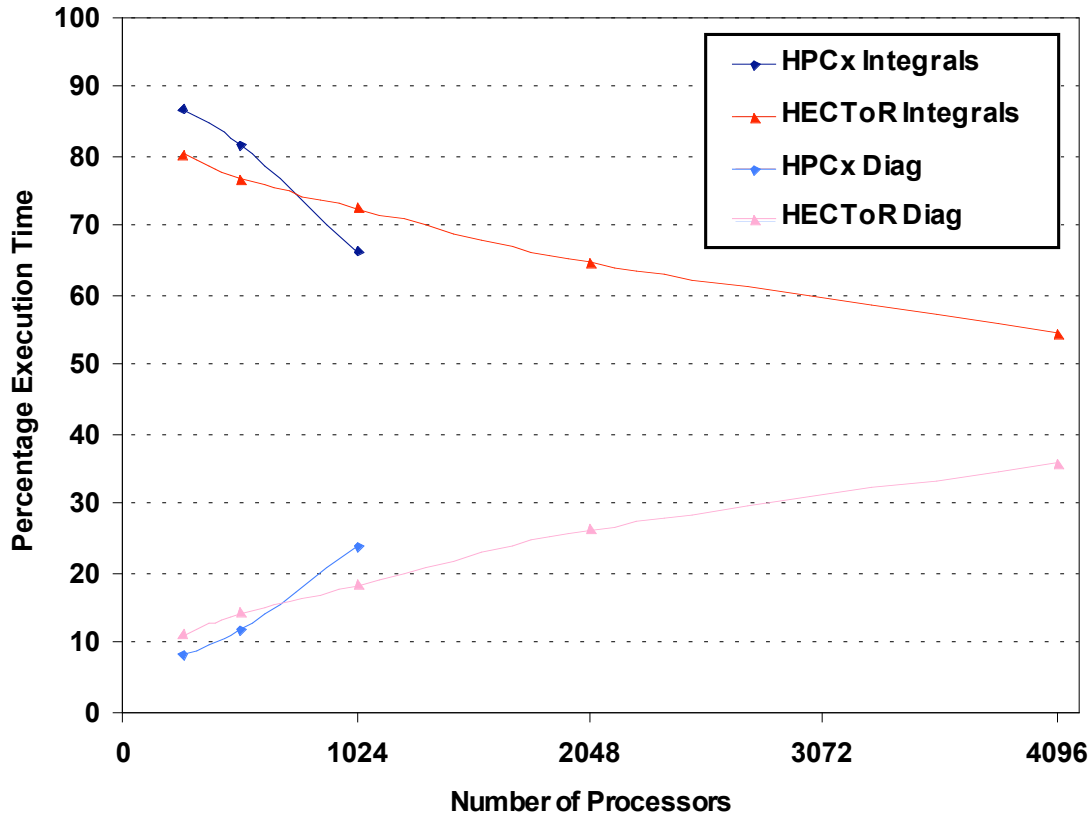

Petascaling CRYSTAL

Possible Approaches:

•New parallel block D&C methods under development

- Block tridiagonalization (Bai & Ward)
- Block D&C method (Bai & Ward)
  - Possibility of trading accuracy for performance
- Could also be in future releases of Scalapack

•Diagonalization-free methods

Conclusions (methods)

- PDSYEVD consistently the fastest parallel eigensolver (for all eigenpairs) for application matrices tested
  - particularly for degenerate eigenvalues
  - high memory overheads
- PDSYEVR performance close to matching PDSYEVD
  - final release may be even faster
  - should have lower memory overheads than PDSYEVD
  - ability to select subsets of eigenpairs
- Reduction to tridiagonal form dominates parallel performance
  - More efficient methods may be available in Scalapack 1.9 ?

Conclusions (machines)

- Cray XT4 / IBM Power 5 comparisons
  - Similar performance upto ~ 256 processors
  - Eigensolvers scale better on Cray XT4 on larger processor counts
    - Especially if the matrices are large
- May need new strategies for petascale architectures
  - How to scale performance upto 10000+ processors?
- Future Priorities
  - Mflops/s or Mflops/watt ?

Acknowledgements

Christian Voemel, ETH Zurich

Ian Bush, Mike Ashworth, STFC Daresbury Laboratory

Technical Reports on Parallel Eigensolver Performance Analysis:

http://www.hpcx.ac.uk/research/hpc

Lapack Working Notes (LAWNS):

- Updates on parallel diagonalizer implementations in Scalapack

  http://www.netlib.org/lapack/lawns