An overview of the R-matrix electron-impact scattering package at NERSC

Workshop on Leadership-class Machines, Petascale Applications and Performance Strategies

July 27-30, Granlibakken Resort

Connor Ballance

Auburn University

Collaborators : D C Griffin
: S D Loch
: N R Badnell
: M S Pindzola

Acknowledgments : Dept of Energy
: NERSC

Structure of Talk

1. What is the R-matrix method ?

2. What is the underlying computational formalism and structure of the code base ?

3. Scaling up from 1000 -> 10000 processors on franklin.

4. Current bottlenecks/possible future directions

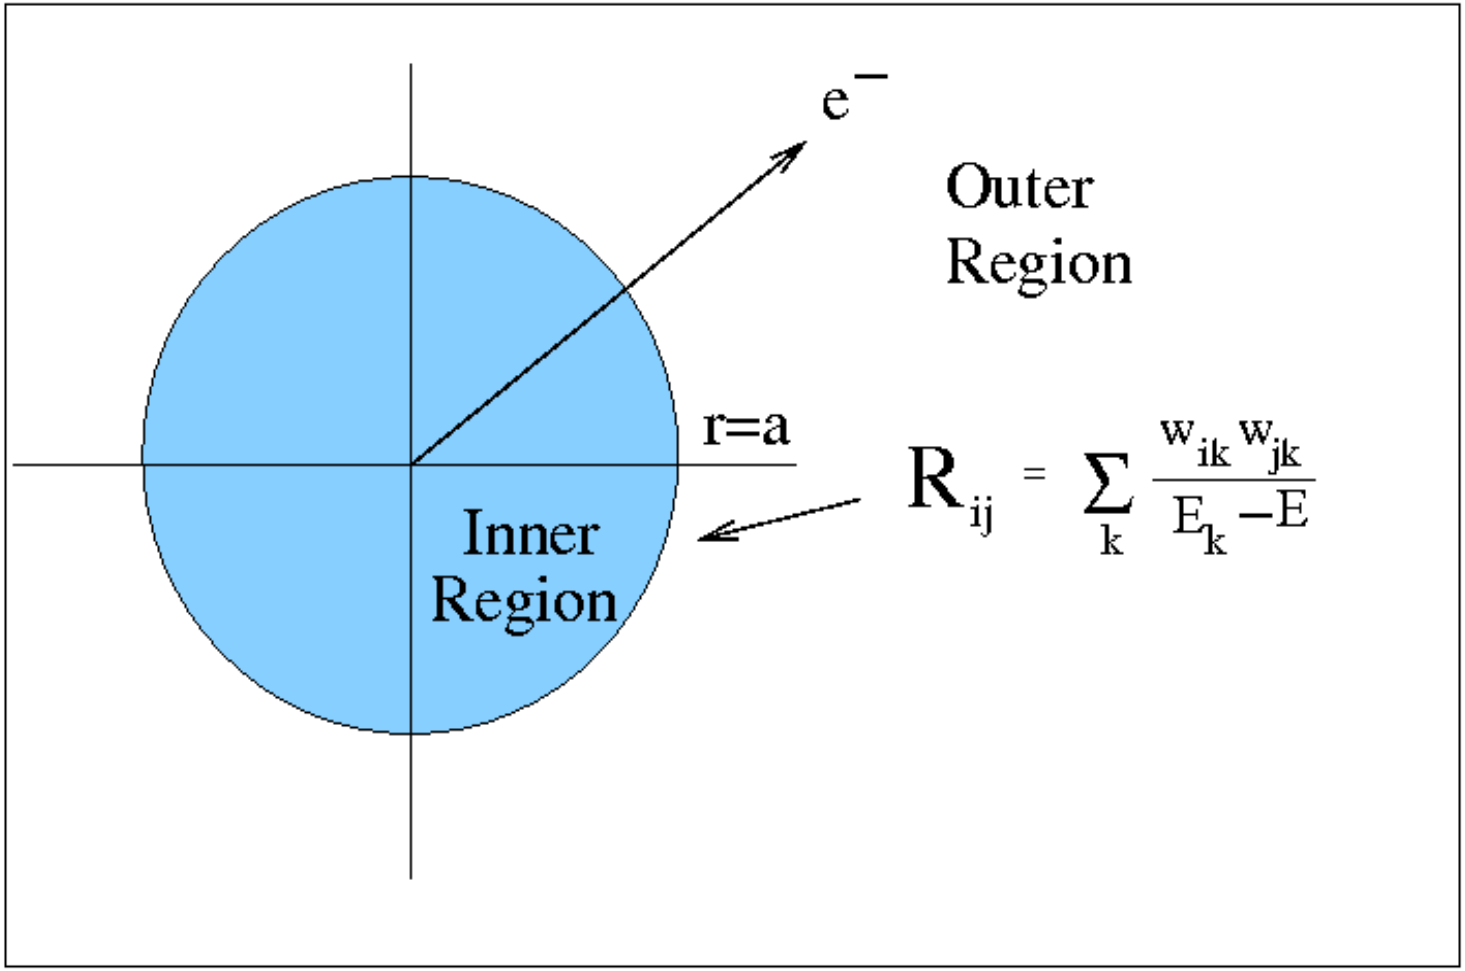
e⁻
Outer
Region
r=a
Inner
Region
$R_{ij} = \sum_k \frac{w_{ik} w_{jk}}{E_k - E}$

Computational Method

1. The formation of many real symmetric matrices (Hamiltonians), typically 60-100, requiring anywhere from 10-500 Gb of storage

2. The diagonalisation of each matrix, from which EVERY eigenvalue and EVERY eigenvector is required.

We achieve this through the ScaLapack package and in particular routines : pdsyevx and pdsyevd. Preference is given to the latter as it ensures orthogonality between all eigenvectors.

Matrices vary in size from 2K * 2K to 100K * 100 K , depending on the complexity of the atomic target

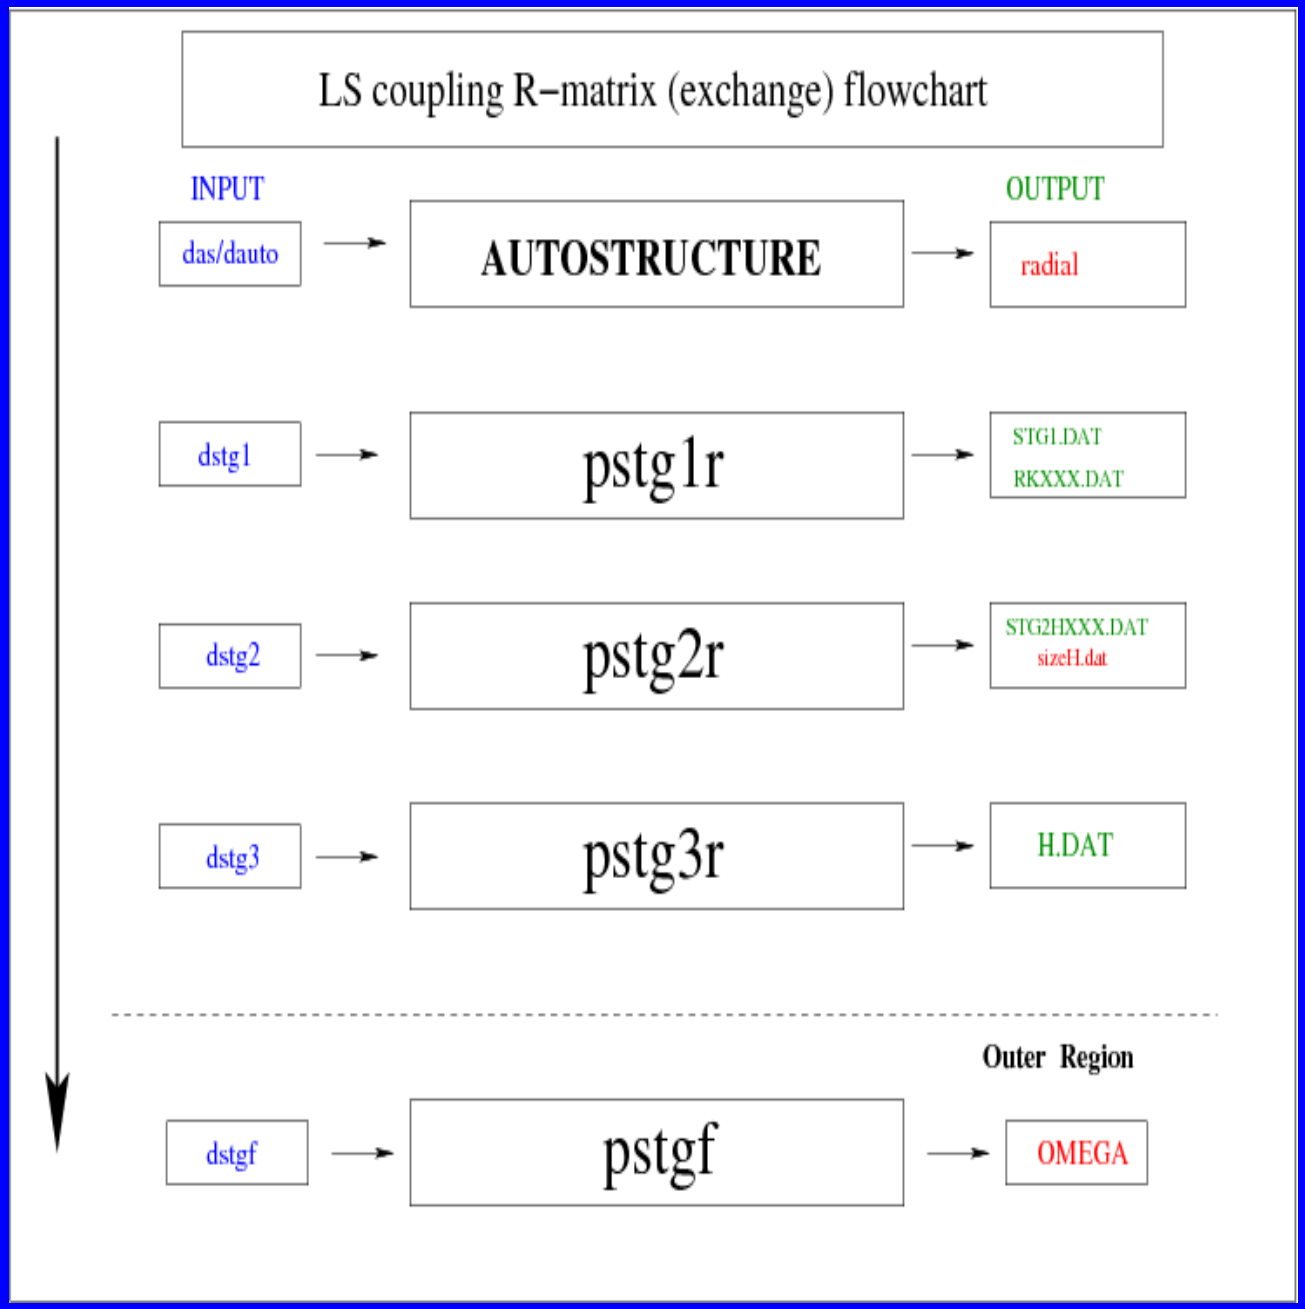

FORMATION

DIAGONALISATION

PHYSICAL RESULT

pstgf is well tuned ... DGEMM and loop-unrolling, make future impact on code optimisation difficult!

FRANKLIN (2500 procs) during the dual core period

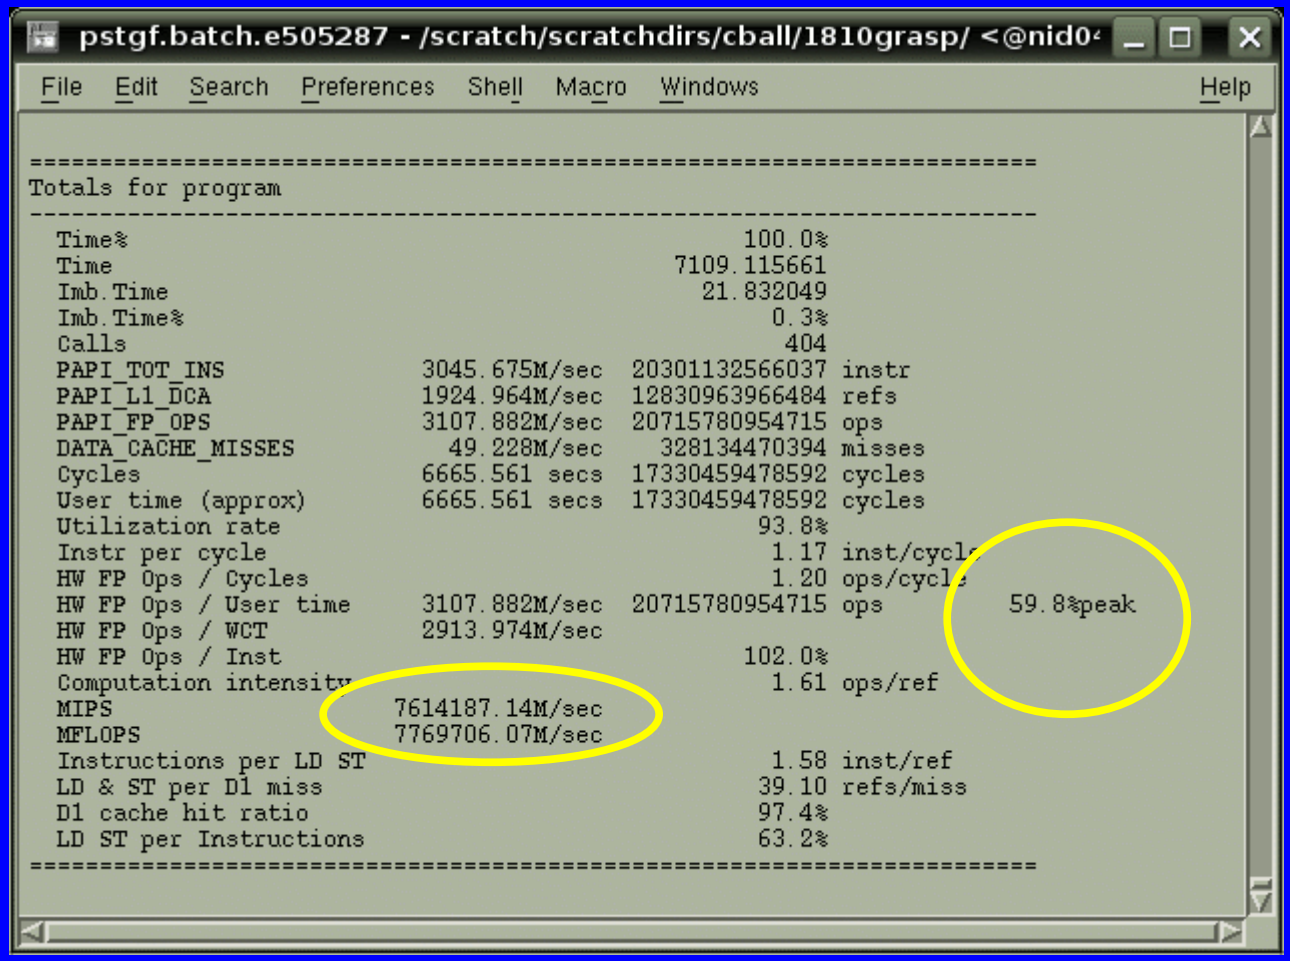

However plenty of progress to be made on the formation and diagonalisation

The formation of the scattering Hamiltonian

The formation of the continuum-continuum part of the N+1 electron Hamiltonian is the most time consuming ....

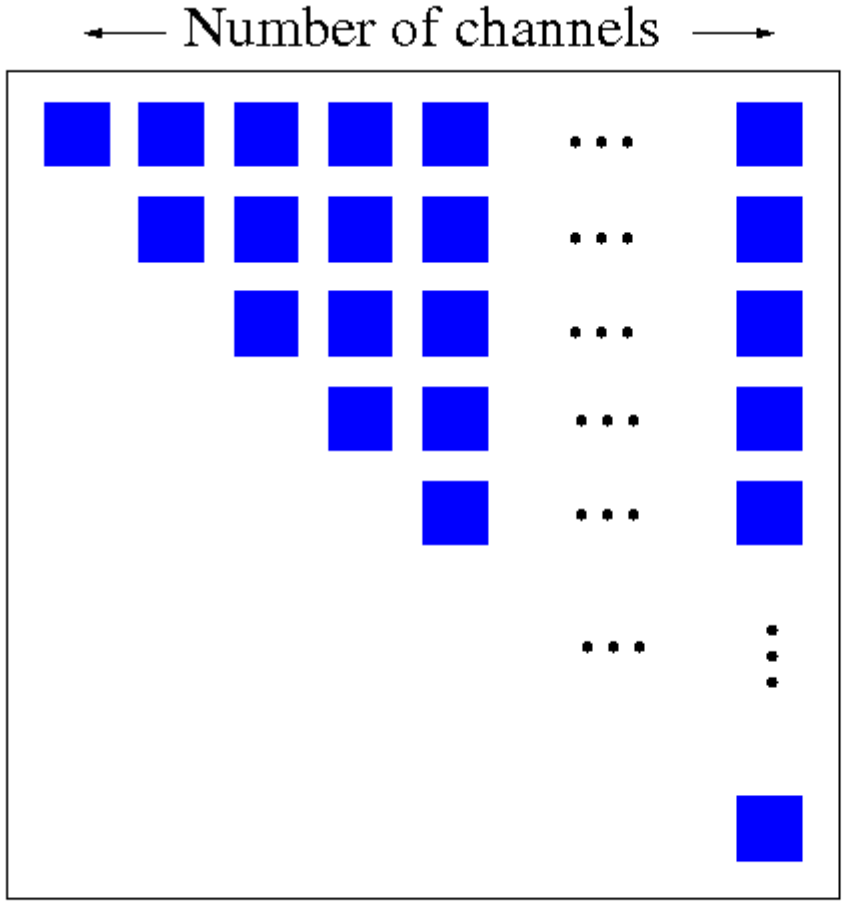

Therefore, if there are several thousand scattering channels

=

(NCHAN*(NCHAN+1))/2

matrix blocks !!!

i.e. several million

not very subtle, but effective .... parallel concurrent diagonalisation

(with help from NERSC consultants : Woo-Sun Yang ,Katie Antypas)

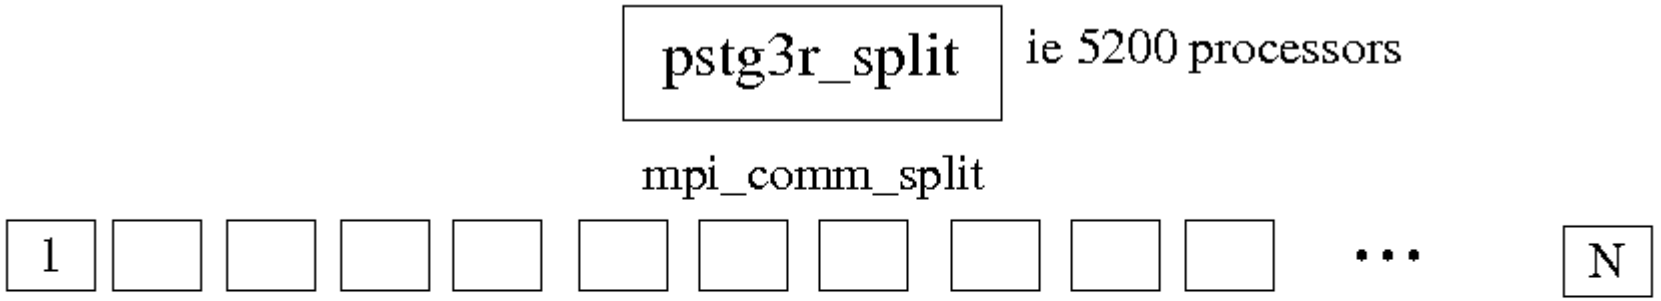

Each small box represents a partial wave

Each subgroup reads a single Hamiltonian and diagonalises it in parallel, concurrently with every other.

Almost endlessly scalable

Reduces all future R−matrix close−coupling calculations to the time required for a single partial wave.

Present situation when 'scaling up' on franklin .....

96 matrices , each 10721*10721 in dimension

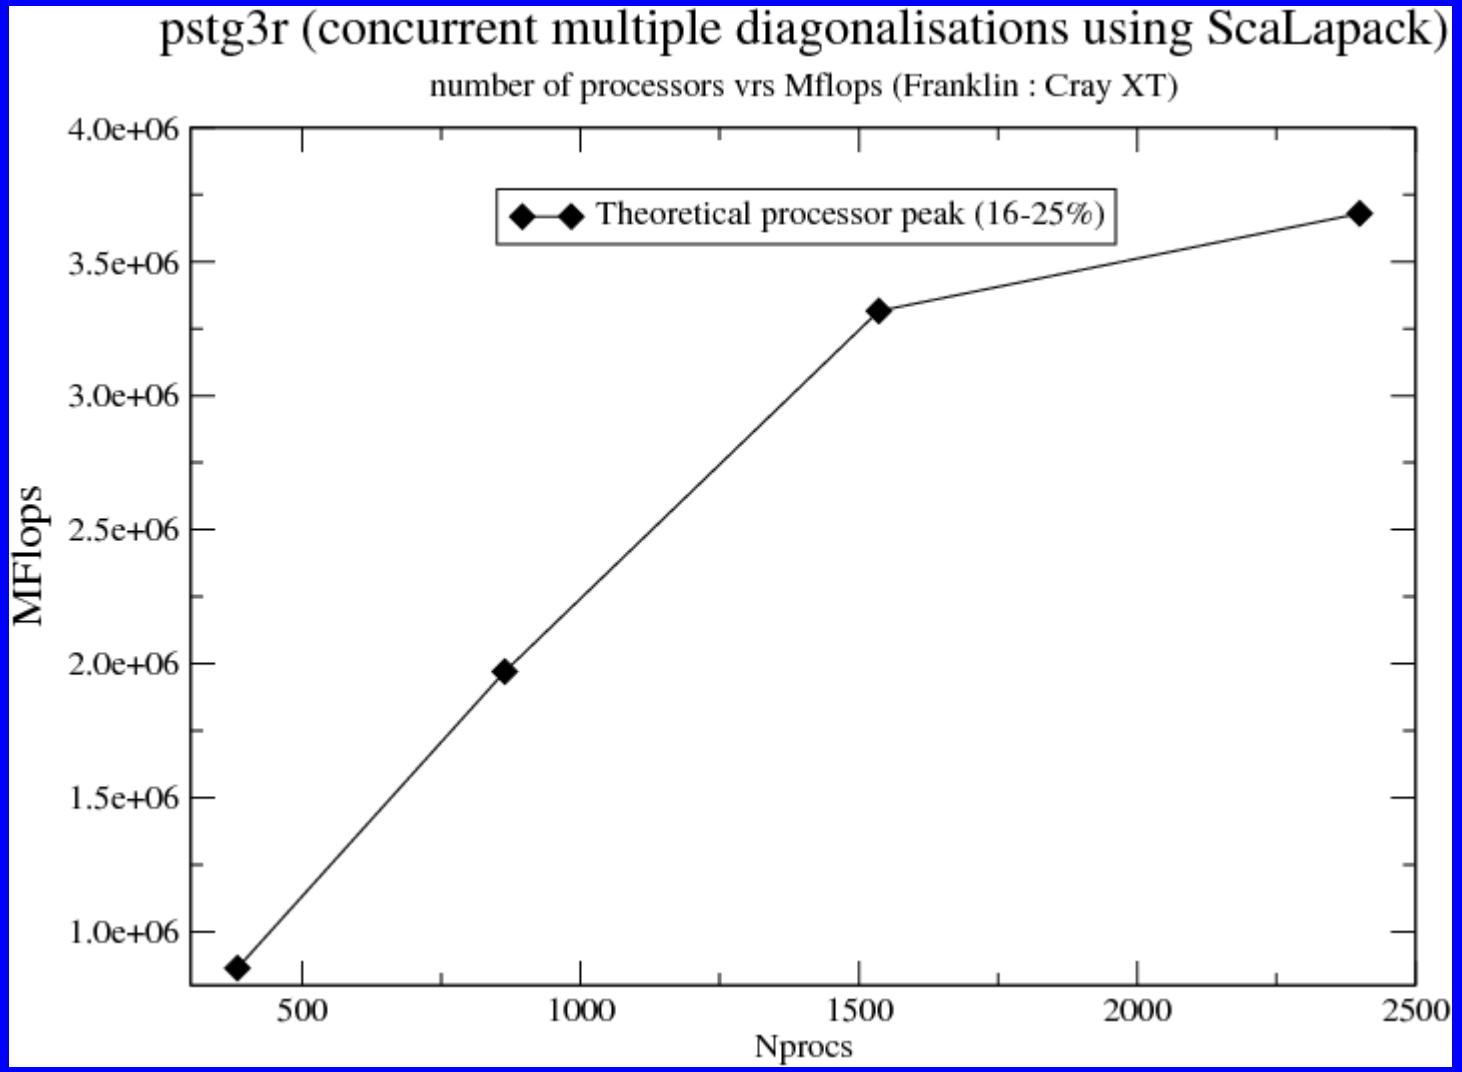

sample profiling for the same case ....

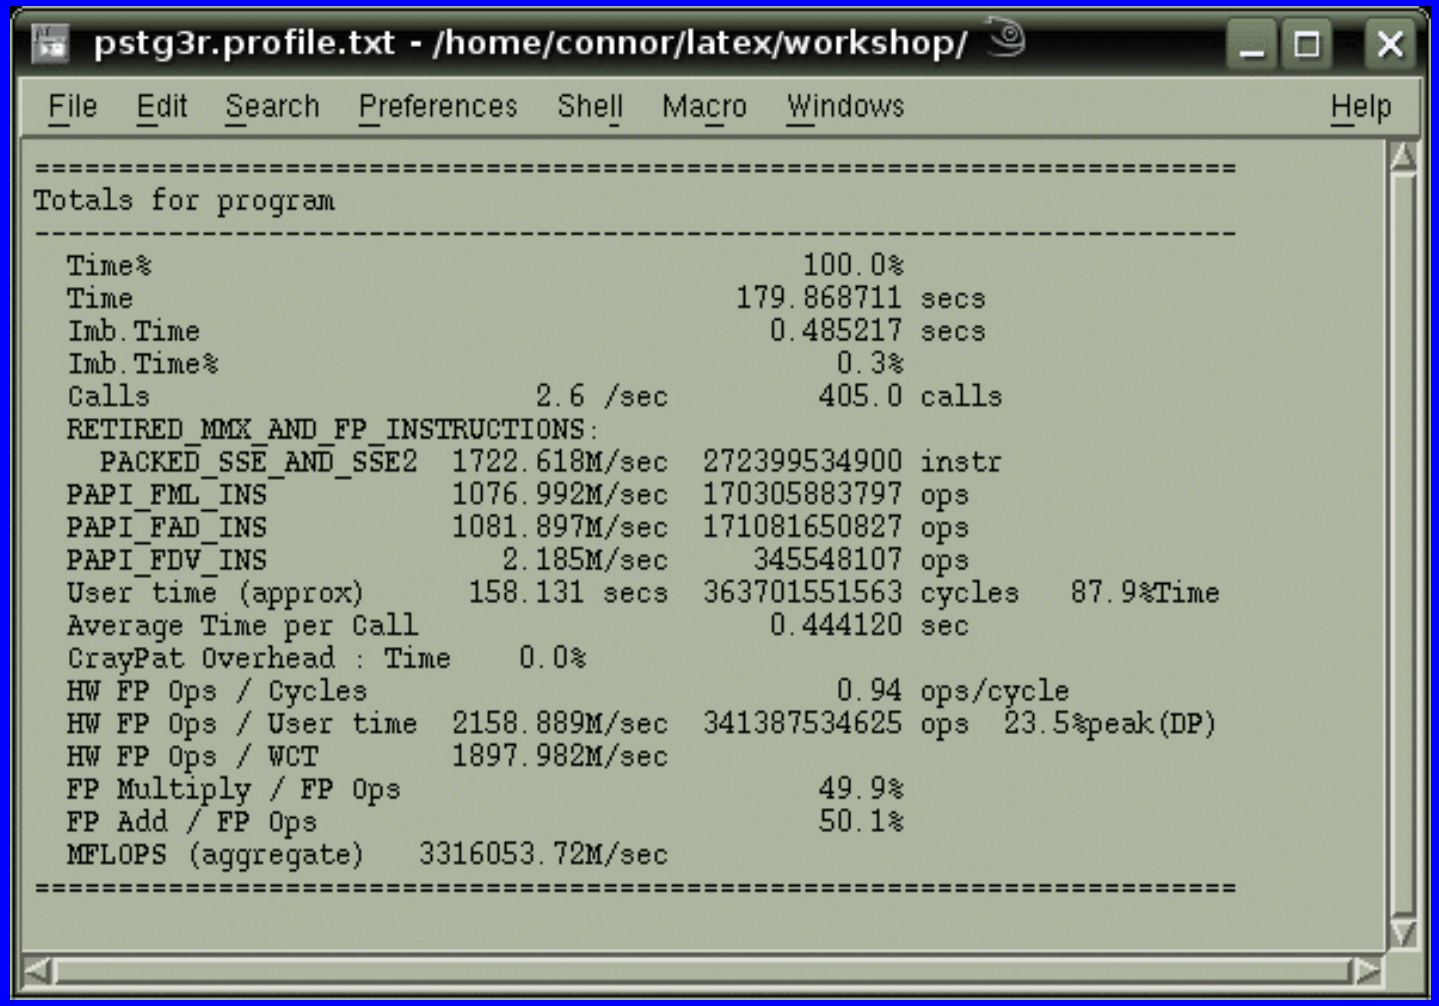

In a nutshell ,one of the biggest remaining challenges to continued scaling of the problem is:

the construction of the Hamiltonian,

then distributing the pertinent information to

a particular processor,
in a particular ScaLapack context,
before diagonalisation.