Present and Future Computing needs in Atomic Physics

John Ludlow, Connor Ballance, Stuart Loch, Teck-Ghee Lee, Mitch Pindzola

Auburn University

Science Goals

- To calculate atomic and molecular collision processes of relevance to controlled fusion energy
- Processes include electron-impact excitation and ionization of atoms and their ions, dielectronic recombination of ions and heavy particle impact excitation, ionization and charge transfer with atoms and ions
- Ensure collisional data are interfaced with plasma modelling codes (ADAS, TRANSP)
- We are presently focused on light elements like H, He, Li, Be, B, C, Ne
- In the next 3-5 years we shall look at heavier fusion related elements such as Xe, Mo, W

ADAS

- The fundamental atomic data is processed through the ADAS suite of codes to give generalized collisional-radiative coefficients and power loss coefficients for a particular element
- Presently have GCR coefficients for H, He, Li, Be.
- Excited state ionization important for Li GCR
  J. P. Allain et al., Nuclear Fusion **44**, 655 (2004)
- Working on B, C, Ne
- More HPC resources in the next 3-5 years will enable GCR data for Mo (Z=42), Xe (Z=54), W (Z=74)

Methods

- R-Matrix with Pseudo-States (RMPS) and Time-Dependent Close-Coupling (TDCC) methods are used to solve the Schrodinger and Dirac equations.
- RMPS is a time independent basis set method that starts from a structure calculation. Diagonalization of large, symmetric Hermitian matrices for all eigenvalues and eigenvectors are required. Current codes use up to 80,000 cores. (~100,000 lines of code split into a number of separate programs)
- TDCC method solves the Schrodinger or Dirac equation on a multi-dimensional lattice using explicit or implicit propagators. (< 10,000 lines of code)
- The dimensionality of the lattice is related to the number of active electrons.

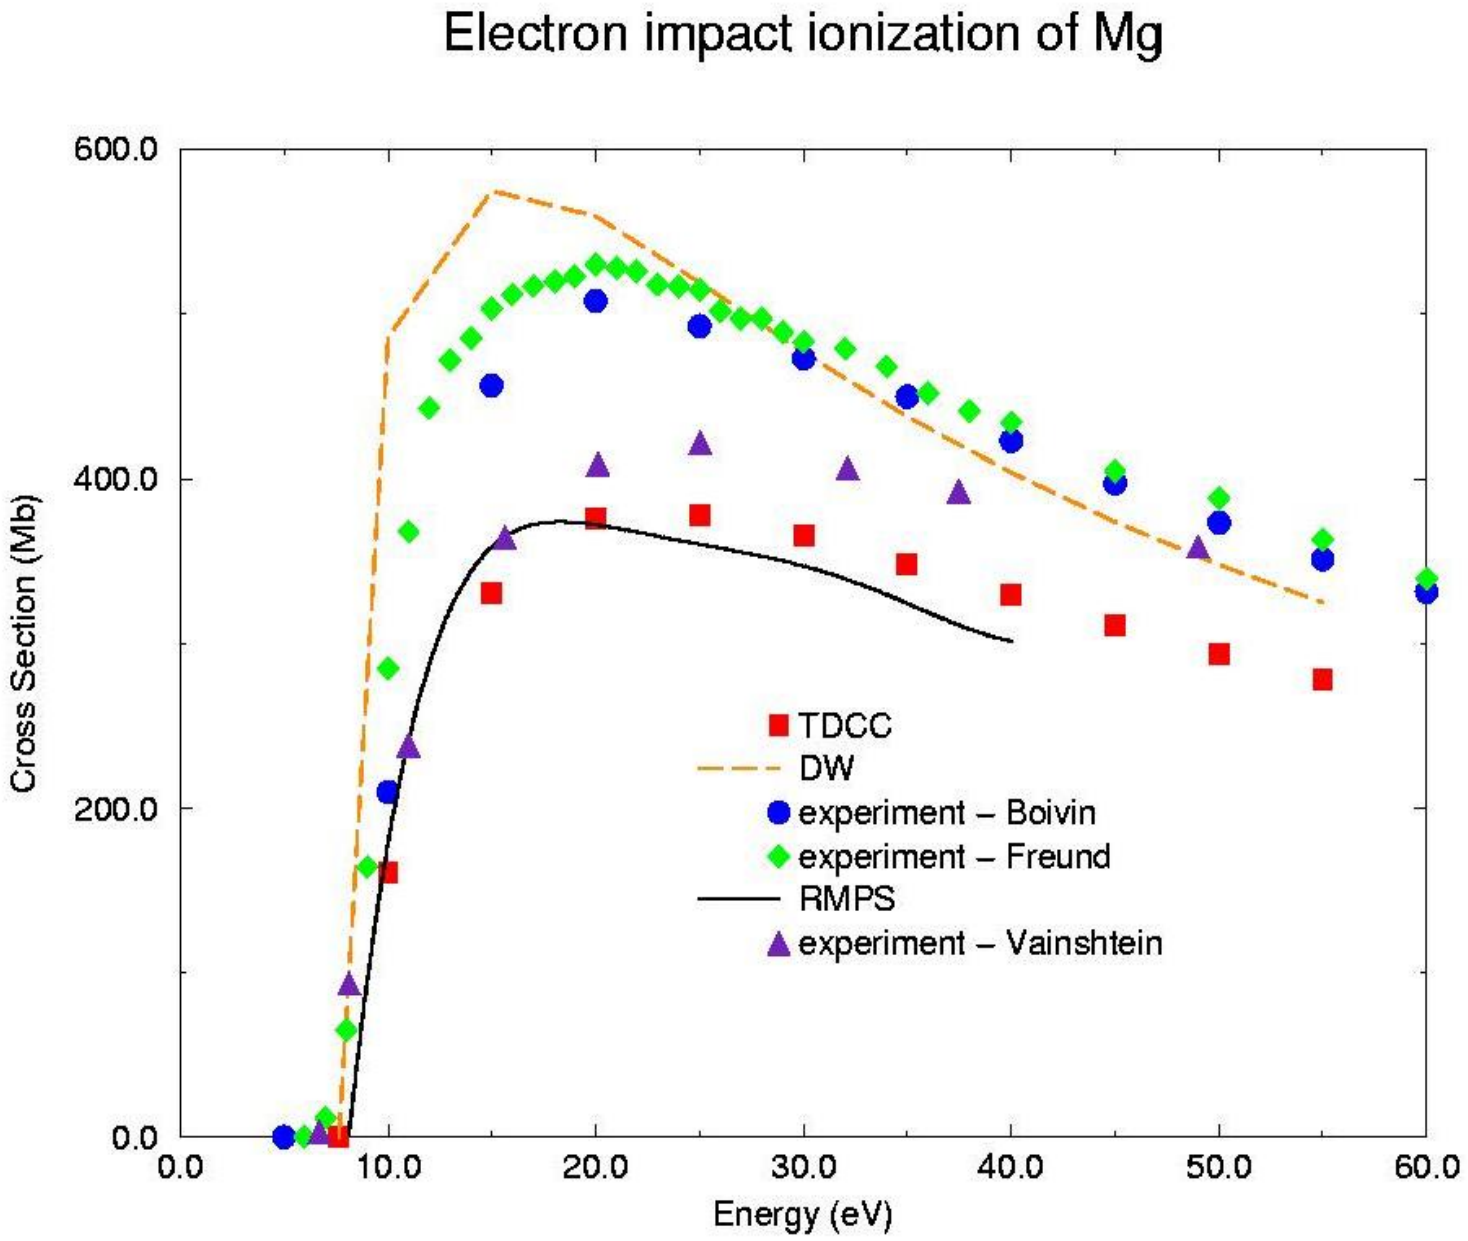
Electron impact ionization of Mg
600.0
400.0
200.0
0.0
Cross Section (Mb)
0.0
10.0
20.0
30.0
40.0
50.0
60.0
Energy (eV)
TDCC
DW
experiment – Boivin
experiment – Freund
RMPS
experiment – Vainshtein

- For electron-impact single ionization of Mg, TDCC-2D used 64 processors on a 384*384 point lattice for 160 different energies and LS partial waves
- For electron-impact double ionization of Be, TDCC-3D on a 192*192*192 lattice used 13,824 processors for 30 incident energies and partial waves. Larger lattices will require > 100,000 processors
- For bare ion double charge transfer, e.g. for alpha + helium, TDCC-6D on a $(100)^6$ lattice, 1,000,000 cores will be needed

- Future RMPS calculations on W, Mo will require concurrent diagonalization of matrices of order 200,00-400,000 across 50-100 JΠ symmetries which will require processor counts in the 100-500k range, outer region of the RMPS code is 'embarrasingly parallel' and could use up to 1,000,000 processors
- The updated parallel codes are currently being applied to Carbon, Neon for electron-impact excitation and ionization
- RMPS and TDCC are independent methods, which provides useful benchmarking in absence of experiment

- Packages used include LAPACK, ScaLAPACK and NAG parallel FFT, NCAR graphics
- Need a Scalapack routine for diagonalizing non-hermitian matrices (e.g. for studies of dielectronic recombination in an external electric field)
- Work has been underway to improve memory needed per core in the RMPS codes by replacing statically allocated arrays with dynamically allocated arrays, a minimum of 2Gb per core is required. The recent introduction of 48Gb nodes on Carver is very welcome (more memory the better)
- Low latency is important for the TDCC codes

- With 50X the HPC resources:
- Electron impact excitation and ionization of W, W^{+}, W^{2+}, Mo, Mo^{+}, Mo^{2+} would be possible with the RMPS method
- More extensive investigations of ionization from excited states of atoms up to higher n-shells with RMPS and TDCC (important for density effects)
- Studies of double and single ionization using the RMPS/TDCC methods
- Double charge transfer in alpha particle collisions with He using the TDCC method

- The introduction of hybrid architectures that utilize accelerators e.g. GPU's will present challenges to the TDCC and RMPS codes
- We have been experimenting with hybrid MPI+ OPEN-MP for the TDCC method using small numbers of cores per node, worked but no advantage over pure MPI so far
- MPI_COMM_SPLIT has been used in RMPS codes to implement multi-level parallelism
- For future nodes containing 10's-1000's of cores, we plan to experiment with CUDA fortran for GPU's. It is also important that our codes are portable, not specific to one variety of GPU or architecture.

- We expect our storage requirements on HPSS to be of the order of 100TB in the next 3-5 years

e.g. to store collision data for all ion stages of W

- High bandwidth in and out of NERSC will be important. Parts of an RMPS calculation can be ran at different centers as the calculation is naturally broken into a number of steps

	Present	3-5 years
Total computational time	10,000,000	100,000,000
NERSC time	4,000,000	40,000,000
Number of cores in typical production run	2000 to 50,000	10,000 to 1,000,000
Wallclock hours of Typical run	1 to 24	1 to 24
Minimum memory per core	2GB	2GB
I/O from a running job	1TB	20TB
Off-line storage	20TB	100TB